

Scheduling and Ahead Markets

Design Options for post-2025 NEM

Creative Energy Consulting Pty Ltd

June 2020

Executive Summary

Introduction

Creative Energy Consulting (CEC) has been engaged by the Australian Energy Council (AEC) to provide advice on the design and development of ahead markets in the National Electricity Market (NEM). These markets do not exist in the current NEM design but the Energy Security Board (ESB) is considering their introduction as part of its review into the post-2025 design of the NEM, whose goal is to develop:

“A long-term, fit-for-purpose market framework to support reliability, modifying the NEM as necessary to meet the needs of future diverse sources of non-dispatchable generation and flexible resources including demand side response, storage and distributed energy resource participation.”

This is a wide-ranging review and the ESB has established a number of separate “market design initiatives”, each with a dedicated team, to explore a different part – or potential part – of the NEM design. This report focuses on the issues and concepts being discussed in the “scheduling and ahead markets” initiative, whilst recognizing that solutions may be found in other areas.

Summary of Conclusions

This paper’s conclusion are summarized in the table below.

Issue	How ahead market proposes to address this issue	Our conclusions
Pre-dispatch bids are non-firm, and a withdrawn bid close to real-time can create security problems which it is difficult for AEMO to deal with at short notice.	Ahead markets can be introduced that will financially discourage, or physically prohibit, generators making such rebids.	Pre-dispatch rebids only create scarcity of services that are not priced in the spot market (eg system strength). So the solution is to price these services, not to discourage rebids.
The current pre-dispatch process – based around self-commitment – might not produce efficient and secure schedules in the future, when new technologies and weather impacts make the scheduling process more complex	An ahead market could incorporate more sophisticated schedulers, including central commitment as is seen in US markets.	Central commitment <i>could</i> be introduced into the pre-dispatch process. Ahead markets are not needed for this. However, self-commitment is a flexible and transparent architecture which is likely to better adapt to future scheduling challenges than a centralized approach.
AEMO is finding the process of scheduling interventions to maintain security progressively more difficult, particularly as security is increasingly dependent on generator commitment decisions	Ahead markets, coupled with a new “UCS” process, will allow AEMO to manage these challenges.	The proposed UCS process can give AEMO the tools it needs. Introducing ahead markets as well is unnecessary.

What are Ahead Markets?

Ahead markets, as envisaged by the ESB, have three essential characteristics:

- Trading takes place over the pre-dispatch timescale of up to 24 hours (or so) ahead of real-time;
- Trading is in the same products as in the spot market: that is energy and market ancillary services; and
- AEMO is the market operator.

The NEM does not currently have, and has never had, such ahead markets.

The following features need to be specified in an ahead market design:

- Whether the forward products are “financial” or “physical”;
- Who can participate, and whether this is optional, strongly-encouraged or mandatory; and
- The structure of the ahead market bids and the mechanism for clearing these.

The ESB has described four ahead market options in its March 2020 paper but has not provided design clarity in these areas. This has made it difficult to understand, analyse and evaluate the ESB options.

Why Introduce Ahead Markets?

The NEM does not have ahead markets and seems to have operated successfully to date, in terms of efficiency, security and reliability. So why are ahead markets being proposed for inclusion in a post-2025 NEM design? The ESB identifies three areas of concern that it considers might be addressed by ahead markets:

- a lack of “firmness” in pre-dispatch bids;
- a potential inability of the existing pre-dispatch process to schedule generation and demand response effectively; and
- the difficulty that AEMO may have in managing security effectively through market intervention.

These concerns emerge from doubt about the effectiveness of the current dispatch and pre-dispatch processes. For that reason, this report has first studied and explained how these current processes work.

In a best-practice approach to design reform, three hurdles need to be cleared:

- the issues to be addressed need to be real and material;
- the proposed redesign must be able to address those issues; and
- alternative, simpler approaches to addressing them must be explored and ruled out as less effective.

For the reasons discussed below, we are not persuaded that these hurdles have been cleared, in any of the perceived areas of concern.

Firmness of Pre-dispatch Bids

It is a feature – and strength – of the current NEM design that bids are able to be changed at any time during the pre-dispatch period, right up until just before dispatch. This allows generators to respond to changing forecasts. If a day is turning out hotter than expected, demand forecasts rise, price forecasts follow, and generators respond by bidding in more capacity.

The opposite is true if demand forecasts fall and generators instead bid some capacity *out*. But this can sometimes cause security problems; if those units were also providing other services, such as “system strength”, that, whilst vital to security, are not priced or paid for. It is this kind of scenario, frequently seen in South Australia currently, that seems to have led to ESB concerns around the “non-firmness” of pre-dispatch bids.

The obvious solution is to find a way to price such essential services – and indeed this is being investigated elsewhere in the ESB processes. With this fix, any scarcity would lead to high forecast prices for these services and so, just as with energy, *more* rather than less of the service will be provided. In short, the problem is caused by missing prices, not rebidding. And because the problem has been misdiagnosed, we are also looking at the wrong remedy: trying to “firm up” pre-dispatch bids.

In any case, the proposal has not been able to explain how ahead markets would make pre-dispatch bids firmer, even if this were desirable. It hopes that generators would take on greater forward commitments in ahead markets and so need to bid more capacity to “defend” this position. Or, if that does not work, the ahead market could be made physical, forcing plant to run in accordance with the “ahead schedule”. But the former seems unlikely, and the latter can be achieved in today’s NEM design, simply by AEMO issuing directions where necessary.

In summary, pre-dispatch “non-firmness” is a strength, not a weakness, of today’s NEM. But, even if it were a problem, ahead markets would not be the way to address it.

Scheduling and Commitment

Scheduling is becoming more complex. It was historically focused on committing slow-start units to meet weather-affected demand but is now moving to include more storage and demand response in the presence of weather-affected *generation*.

So it is plausible that the current pre-dispatch process, whilst effective historically and currently, might not remain fit-for-purpose over the long term. And, perhaps, to address this, the existing “self-commitment” design – where generators are responsible for deciding when to start-up and shut-down their own plant – could be replaced by a “central commitment” approach, in which these decisions are made by a central algorithm, designed and operated by AEMO. Certainly, there is no harm in contemplating and assessing such a change as part of the post-2025 design review.

But the ESB has instead decided to seek solutions to this issue in the form of ahead markets. This is a mistaken approach, for two reasons. Firstly, the question of how to design such a new scheduling algorithm still needs to be addressed, whether this algorithm is part of the pre-dispatch process or a new ahead market clearing process. Secondly, pre-dispatch provides a full picture of the scheduling problem to be solved, because participation is both mandatory and physical. It has not been explained

how to can achieve similar participation in an ahead market; or how, without this, scheduling can be done effectively.

So, again, ahead markets do not have the potential to address the concerns.

AEMO Intervention

AEMO has the responsibility and the tools to intervene in the market to ensure system security where the market processes, of pre-dispatch and dispatch, would fail to achieve this by themselves. Historically, AEMO has needed to intervene only occasionally and has relied on manual, *ad hoc* methods to decide when and how to intervene. The ESB anticipates that such intervention is likely to become more frequent and complex in the future, for similar reasons to the increasing complexity of scheduling, noted above.

So what AEMO needs is new scheduling tools – analogous to those developed and used by generators. And, in fact, the ESB is proposing this: that AEMO develops a new “unit commitment for security” (UCS) process. This proposal is supported in principle, although some important details need to be worked through.

Nevertheless, the ESB considers that there is still a potential role for ahead markets, by providing important input to the UCS process. But, in fact, all the necessary information would be provided by pre-dispatch. Once again, ahead markets are superfluous.

Conclusions

The introduction of ahead markets would be a major change to the NEM design, with consequent impacts on market participants and customers. In proposing such a reform, the ESB needs to make it clear what issues it is seeking to address, how ahead markets would address them, and why alternative, less disruptive, reforms could not. Rather than introducing ahead markets, concerns around:

- *pre-dispatch firmness* should be addressed by changes to the spot market;
- *scheduling* could be addressed through changes to pre-dispatch
- *AEMO intervention* would be addressed by the ESB’s proposed UCS process

The NEM has no need for ahead markets and there is no reason to consider them further in the post-2025 design project.

TABLE OF CONTENTS

1	INTRODUCTION	1
1.1	ENGAGEMENT AND SCOPE	1
1.2	CONTEXT	1
1.3	APPROACH.....	3
1.4	STRUCTURE OF THIS REPORT.....	4
2	DISPATCH AND THE SPOT MARKET.....	6
2.1	INTRODUCTION.....	6
2.2	CONCEPTS	6
2.3	ISSUES ARISING.....	10
2.4	POTENTIAL SOLUTIONS.....	15
2.5	CONCLUSIONS	22
3	SCHEDULING AND PRE-DISPATCH.....	23
3.1	INTRODUCTION.....	23
3.2	CONCEPTS	24
3.3	ISSUES ARISING.....	31
3.4	POSSIBLE SOLUTIONS	35
3.5	CONCLUSIONS	44
4	AHEAD MARKETS.....	45
4.1	INTRODUCTION.....	45
4.2	CONCEPTS	46
4.3	ISSUES.....	55
4.4	POTENTIAL SOLUTIONS.....	61
4.5	CONCLUSIONS	64
5	ESB OPTIONS	66
5.1	INTRODUCTION.....	66
5.2	ESB OPTION 1	66
5.3	ESB OPTION 2	68
5.4	ESB OPTION 3	70
5.5	ESB OPTION 4	74
5.6	US MARKET DESIGNS	75
5.7	CONCLUSIONS	77
6	OVERALL CONCLUSIONS	78
6.1	OVERVIEW.....	78
6.2	FIRMNESS OF PRE-DISPATCH BIDS	78
6.3	SCHEDULING EFFECTIVENESS OF PRE-DISPATCH.....	79
6.4	AEMO INTERVENTION.....	80
6.5	US MARKETS	80
6.6	FINAL CONCLUSION	81
	APPENDIX: OVERSEAS MARKETS.....	82

1 INTRODUCTION

1.1 ENGAGEMENT AND SCOPE

Creative Energy Consulting (CEC) has been engaged¹ by the Australian Energy Council (AEC) to provide advice on the design and development of ahead markets in the National Electricity Market (NEM). These markets don't exist in the current NEM design but the Energy Security Board (ESB) is considering their introduction as part of its review into the post-2025 design of the NEM. The AEC is looking to:

- Better define ahead-markets;
- Understand what issues (or opportunities) they attempt to address;
- Understand the implications were they to be introduced into the NEM; and
- Identify other approaches to these issues.

The AEC's objectives are described in the request for proposal for this project as follows:

"The AEC is seeking initially a clear definition and description of what ahead markets are and what would need to be introduced into the NEM to achieve them. From this definition, a qualitative assessment should be made of the potential overall benefits, costs and risks that would result should the NEM shift to this design either as a whole or by the implementation of individual aspects. This will aid the AEC in its understanding of the concept, how competitive businesses might operate within it, and the transitional challenge from the existing NEM. Consideration should be given to the physical, economic and commercial dimensions."

This report describes CEC's approach to, and findings from, this task.

1.2 CONTEXT

Ahead markets are being considered as a part of the ESB post-2025 design review. The review was requested by the COAG Energy Council, with a goal of advising on:

"A long-term, fit-for-purpose market framework to support reliability, modifying the NEM as necessary to meet the needs of future diverse sources of non-dispatchable generation and flexible resources including demand side response, storage and distributed energy resource participation."

This is a wide-ranging review and the ESB has established a number of separate "market design initiatives" (MDIs), each with a dedicated team, to explore a different part – or potential part – of the NEM design.

¹ CEC has subcontracted MarketWise Solutions to assist with this engagement

There are, at the last count, seven MDIs, four of which have direct relevance to this project:

- Essential System Services
- Scheduling and Ahead Markets
- Two-sided Markets
- Coordination of Generation and Transmission Investment (COGATI)

For the time being, these initiatives are progressing in parallel. Later this year, the findings of the MDIs will be brought together and some candidate designs for the NEM overall will be developed by stitching together the ideas and options from the MDIs.

This is an unconventional approach to market design. It is difficult for the individual teams to develop sensible options in isolation and, equally, it is difficult for stakeholders to anticipate and infer how the individual elements might interface and interact in a finished design.

CEC has been engaged in this project since March and, over this period, the ESB work has been proceeding in parallel, with new concepts or details coming to light in ESB publications and workshops: in particular:

- The publication of the ESB papers to the March 2020 COAG, in particular a paper on “system services and ahead markets”²
- A presentation to the industry working group for the ESB project³
- A webinar to present and discuss further details of one option⁴ outlined in the March ESB paper⁵.
- A future webinar⁶ to present and discuss further details of the other options that were outlined in the ESB paper.

There have also been some publications and workshops relating to the Essential System Services and Two-sided Market MDIs over this period.

This dynamic situation has made it difficult to meet the objectives of the AEC, for several reasons:

- The design definitions for the ahead market options that the ESB has published are generally incomplete, unclear, or lacking in coherence; indeed, it appears to be the ESB’s strategy to release ideas early, before the design or implications have been fully developed.
- Whatever assessment CEC makes of the *currently-published* designs is likely to be quickly “overtaken by events” as the designs are amended, improved or clarified;

² “System Services and Ahead Markets”, ESB paper to COAG EC, March 2020

³ Technical Working Group Webinar, 7th April 2020

⁴ the Unit Commitment for Security (UCS) concept

⁵ Technical Working Group Webinar, 14th May 2020

⁶ scheduled for 15th June 2020

- Functional overlaps between the different MDIs means that the eventual design of the ahead market element, and its impact and effectiveness, will depend upon the results from other MDIs.

Fundamentally, an ahead market is a “shell” of a concept. Any market must be defined in terms of participation (who buys and who sells), products (what is traded) and timing (when does the trading occur). The first two aspects are being considered in two different MDIs: the “two-sided market” and “essential system services” workstreams, respectively. The ahead market MDI considers only the third aspect.

1.3 APPROACH

In the light of this context, the approach taken in this report is to look at these potentially new or amended NEM markets in the round, rather than just considering the market timing “ahead” issue in isolation. So fundamental concepts around market design are first defined and discussed, because these concepts do *not* become stale: any design proposals that the ESB presents now or in the future must incorporate and reflect these concepts. A solid conceptual grounding provides a framework for analysing and assessing market design proposals, current and future.

Similarly, the issues to be addressed by these market reforms, and the impacts they might cause, are considered conceptually rather than empirically and quantitatively. This is different from a usual market design process which will start with identifying *actual* issues arising in the market and then consider alternative design reform options to address those specific issues. The more abstract approach has been taken here because:

- The ahead market concept is itself rather abstract: since it is not even clear what services might be traded on this market;
- Any issues in the market may be dealt with by options developed in *other* MDIs: for example, today’s problem of frequent intervention by the Australian Energy Market Operator (AEMO) in the market may well be addressed by introducing new *essential system services* into the market design: so ahead markets may not be necessary to address current issues; or, alternatively, they may be needed to address *new* issues that might emerge when these other design reforms are introduced.
- The ESB goal is for the NEM design to be fit-for-purpose for the long-term: given the uncertainty around the development and deployment of emerging and new technologies over this time frame, the design must be robust against a range of alternative futures, rather than simply addressing a current issue caused by existing technologies.

The conceptual framework developed is then used to analyse and assess the latest options put forward by the ahead markets MDI team. It is hoped that the AEC and its members will also find the conceptual discussion useful for examining future proposals from the ESB.

1.4 STRUCTURE OF THIS REPORT

The next three chapters of this report develop this conceptual approach for 3 major areas of the NEM design:

- Dispatch and the spot market
- Scheduling and pre-dispatch
- Ahead markets

The relationship of these processes to each other, and to other NEM processes, is illustrate in figure 1, below.




Figure 1: Overall Market Architecture

The first two areas encompass much of the existing design of the NEM. Thus, the current design is used to explore and illustrate the various embedded concepts, which then become important when considering ahead markets. Ahead markets are new to the NEM, although they are common in overseas markets. But they are necessarily defined in terms of the concepts that exist, or could exist, in dispatch and pre-dispatch (PD): the products and services that are traded; the bid and offer structures and participation processes; and the clearing algorithms, prices and settlements processes. So, the conceptual examination of dispatch and PD sets us up well for a coherent description and discussion of ahead markets.

Each of these chapters is divided into three main subsections:

- *Concepts*: describing these theoretically but illustrating them in terms of the current NEM design;
- *Issues Arising*: exploring current or potential shortcomings in NEM operation that the ESB has expressed concerns about in its design project;
- *Possible Solutions*: identifying possible reforms to the NEM design, generally *not* requiring the introduction of ahead markets, to address these issues.

Each “possible solution” is described and discussed at a high-level and its potential strengths and weaknesses assessed. It is not the intent to reach a definitive recommendation as to whether a “possible solution” should be incorporated into – or excluded from – the NEM post-2025 design. However, it is generally recommended that these potential solutions should be considered by the ESB project and compared to the ahead market proposals for efficiency and effectiveness.

The final chapter examines the four ahead market options that the ESB has presented at the time of writing. For each option, the report:

- *describes the option* in terms of the concepts introduced in the earlier chapters. This can be difficult where the description is unclear, usually because the terms and concepts it employs have not been defined,
- *assesses* its likely effectiveness in addressing the issues previously identified;
- *infers* the possible direct impacts on AEC members: eg in developing the capacity to trade in these new ahead markets

As already noted, the shelf life of this last chapter may be limited, given that the ESB’s proposals are constantly being revised and updated. However, the conceptual framework developed in the preceding chapters should retain its relevance, and provide the analytical tools needed to understand and assess future proposals as they emerge from the ESB work program.

2 DISPATCH AND THE SPOT MARKET

2.1 INTRODUCTION

All roads lead to dispatch. For all its labyrinthine complexity, the NEM design has a straightforward ultimate goal: to ensure an efficient and secure dispatch of generation, thus providing a reliable and economic supply of electricity to consumers. So it is impossible to develop and assess other demand elements without the fundamental underpinning of clear dispatch concepts: what it does, why and how.

There are two facets to dispatch, which makes it both elegant and subtle. It presents as a traditional engineering problem, which system operators have grappled with for decades, long before deregulated energy markets appeared on the scene: how to dispatch the generation fleet in a least cost way which supplies electricity demand whilst ensuring system security. Engineers have over the years, and with the help of increasingly powerful computers, developed sophisticated ways to formulate and solve this problem.

But at the same time dispatch is the platform for clearing spot energy and ancillary services, setting the spot prices which then ripple through the scheduling, planning and investment timescales. So it is the financial heart of the NEM, about which other processes and markets revolve.

AEMO's twin roles reflect these two facets: it is the *system operator*, responsible for maintaining system security and reliability; but it is also the *market operator*, responsible for providing the platforms that allow the spot market and other markets to function effectively. In these roles, AEMO needs to be "hands-on" and "hands-off", respectively, and it is the tensions created by this paradox that seem to be at the root of the issues that the ESB is concerned with, and which it is considering introducing ahead markets to address. This philosophical ambivalence or contradiction around AEMO's role seems to be creating some associated uncertainty and complexity in the ESB's design process and proposals.

For these reasons, this report starts by examining the dispatch process, carefully and clinically. To define the underlying concepts that it makes use of, particular in terms of the "spot market" facet. Because, to be effective, any ahead markets must complement and enhance the existing spot market.

2.2 CONCEPTS

2.2.1 Defining "Dispatch"

In its March ahead market paper, the ESB defines dispatch as the "process of scheduling, committing, and providing targets in the electricity market" which seems to cover both dispatch and PD. This report separates out these two timescales. So *dispatch* here refers only to real-time processes:

- the clearing of generator bids and offers using the NEM dispatch engine (NEMDE);
- the issuance of dispatch targets based on these cleared amounts, which generators are required to follow
- the derivation of spot prices used for settlement

Actions taking place *ahead* of real-time are discussed in the pre-dispatch chapter.

2.2.2 Physical Auction

NEM dispatch is an *auction*⁷ in that generators submit offers to be dispatched and AEMO selects the cheapest offers first. Auctions generally⁸ have three objectives:

- Select bids and/or offers to be cleared such that supply matches demand;
- Optimise the benefits from trade: eg by selecting the lowest price offers and/or the highest price bids
- Set a clearing price at which the cleared bids and offers are traded

A *clearing price* requires that:

- Every cleared offer has a price no higher than the clearing price, and every uncleared offer has a price no lower than the clearing price;
- Every cleared bid has a price no lower than the clearing price and every uncleared bid has a price no higher than the clearing price

Setting the spot price at a clearing price means that every generator is dispatched in accordance with its offer at a price X which, in effect, say: “only dispatch me if the spot price is at least X”.

Where multiple products are traded simultaneously, there is a separate clearing price for each product and the clearing price principles are generalized. In this case, the volume-weighted-average (VWA) clearing price across the products cleared must, for each auction participant, exceed the VWA offer price, with the weighting relating to the quantities cleared for each product. So a “loss” on one cleared product would be offset by a “profit” on another.

It is not always possible, or even necessarily desirable, to establish and use clearing prices. However, where auction prices do not satisfy the clearing price principles, the efficiency of outcomes may be impaired. For example, the NEM uses regional energy pricing, rather than nodal pricing, meaning that generators can be *constrained on or off*; ie dispatched in a way inconsistent with their offers and in contravention of the clearing price principles described above.

The dispatch auction is *physical* in that each generator must generate in accordance with the dispatch instruction which follows from the cleared offer, and may face statutory penalties if it unreasonably fails to conform with this instruction. To participate in the dispatch auction, then, one must operate a corresponding physical plant and bid in accordance with its operating capabilities.

⁷ strictly a reverse auction, since it is sellers rather than buyers who participate

⁸ not all auctions though. For example, some have a “pay-at-bid” structure instead of using clearing prices.

Figure 2, below, illustrates the linkages between the physical and financial in the dispatch process.

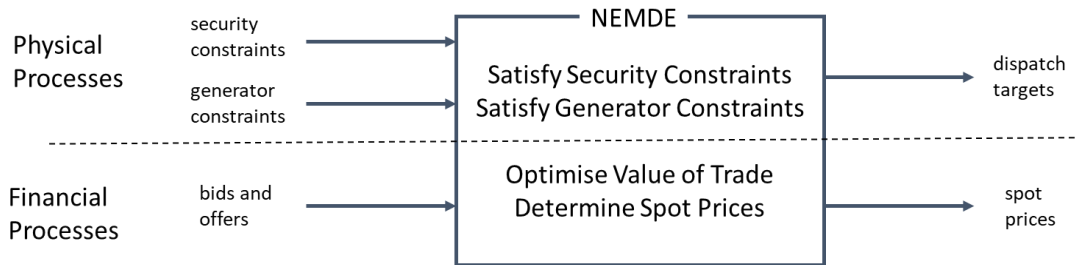


Figure 2: Physical and Financial Elements of the Dispatch Process

2.2.3 Gross Market

Large generators⁹ are required to operate in conformance with dispatch instructions, and the only way to receive such instructions is to participate in the dispatch auction. Even those smaller (non-scheduled) generators not required to be active in the auction must still trade physical energy at the price determined by dispatch¹⁰. The demand side similarly passively buys at the spot price. Given the high level of mandatory participation, the NEM dispatch auction is commonly referred to as a *gross market*.

The requirement for active participation gives AEMO high visibility, and control, of the intentions and capabilities of generators and this allows it to maintain system security, as discussed next. However, as the passively-participating portion of the market (the demand side and non-scheduled generation) becomes larger, more volatile and more responsive to spot prices, new rules may be needed to give AEMO greater visibility and control of this segment. This is referred to as a *two-sided market*. This is being considered by the ESB post-2025 design project, but is outside the scope of this report.

2.2.4 Secure Dispatch

AEMO must ensure that generation dispatch is *secure*: ie that power system security is able to be maintained. This is done by including constraints within NEMDE which the dispatch outcome must comply with. The constraints are formulated by placing *controllable variables* (dispatch targets) on the left-hand-side (LHS), and constants or formulae that are independent of these controllable variables on the right-hand side (RHS). Constraints can take two essential forms, depending upon whether a “ \leq ” or “ \geq ” inequality relates the LHS and RHS of the constraint.

⁹ scheduled and semi-scheduled generators

¹⁰ except for a few ‘non-market’ participants

Security constraint can be thought of as ensuring that sufficient supplies of *essential system services* are provided in dispatch. The form of inequality reflects whether it is the supply side or the demand side that is controllable by AEMO. If the *supply* side is controllable, then we have:

$$\text{Controllable supply} \geq \text{uncontrollable demand}$$

System services which are controlled in dispatch and are formulated in NEMDE in this way are referred to as *market ancillary services* (market AS). Currently only frequency control ancillary services (FCAS) are managed in this way, but additional market AS might be introduced in the future, as discussed further below. AEMO also procures other system services, but ahead of real-time and outside of the dispatch process. These are referred to as *non-market AS* and are discussed in the pre-dispatch chapter.

The demand quantity might be quite volatile, in which case it may need to be recalculated or reforecast prior to each dispatch run.

The second form of NEMDE constraint is employed where it is the *demand* side that is controllable through dispatch. In this case, the constraint takes the form:

$$\text{Controllable demand} \leq \text{uncontrollable supply}$$

For example, such a constraint would be used to prevent transmission flows (the *demand* for transmission) exceeding transmission capacity (the *supply* of transmission capacity)¹¹. The supply quantity might also be volatile and need to be regularly recalculated between dispatch runs.

Another example of this form of NEMDE constraint is the so-called *system strength constraint* used currently to manage voltage stability. The constraint might take the form:

$$\text{Aggregate Output of non-synchronous generation} \leq f(\text{combination of synchronous units on-line})$$

So in this case the non-synchronous output is creating a demand for this system strength service, which is supplied by on-line synchronous units. NEMDE can constrain the demand for system strength by reducing the dispatch of non-synchronous units.

The $f()$ function on the RHS is a complex function that is derived by AEMO through system studies. It is not controllable in dispatch because the factors it depends upon – unit commitment – are not controllable. Therefore system strength is not a market AS¹².

Because the provision of, or use of, these essential system services is specific to the size, technology or location of each generating unit, dispatch needs to operate at this level, to know precisely what the supply of, and demand for, essential system services will be under alternative dispatch solutions and ensure that supply always covers demand. Thus separate offers for each unit must be made; each unit then receives its own dispatch target (for energy and market AS) and must conform with these. Without this level of granularity, it would not be possible to ensure system security through the dispatch process.

¹¹ transmission capacity *can* be controlled – eg by switching lines in or out – but not in dispatch. This is done ahead of real-time, outside of the dispatch process.

¹² in fact, it is not even a *non-market AS* currently, as discussed in the next section.

2.2.5 Shadow Pricing and Spot Pricing

In NEMDE, all constraints are expressed in a *linear* form: the LHS is a simple linear combination of the controllable variables of energy and market AS dispatch targets. The RHS can take any form, because it is calculated outside of – and prior to – the NEMDE solving process. Furthermore, controllable variables are continuous within a specified range: eg a generator offering 100MW of energy can be dispatched for any real number of MW up to this amount¹³. As a result of these characteristics, the NEMDE problem is a linear program (LP) and can be solved quickly and reliably.

With this formulation of the secure dispatch problem, *shadow prices* are created for each constraint that binds in a NEMDE run. These shadow prices reflect the marginal value of the provision – or the marginal cost of the use - of the associated essential system service who need is reflected in the constraint. For mathematical reasons, they are also clearing prices, in the sense defined above. So, these shadow prices can potentially be used to set spot prices for the system service that the constraint relates to. For example, in the current market design, FCAS spot prices are based on the shadow prices of the associated NEMDE constraints.

2.3 ISSUES ARISING

2.3.1 NEMDE Incompleteness

A problem arises where it is *not* possible to formulate a set of constraints for NEMDE, of the appropriate form, such that compliance with these constraints is sufficient to ensure system security. Logically this might be because:

- It is impossible to formulate *any* set of constraints that ensure secure dispatch; or
- The constraint set can be formulated, but some of these constraints do not include any variables that are controllable by dispatch or;
- There are controllable variables, but the LHS of the constraint cannot be expressed in the necessary linear form for NEMDE; or
- The controllable variables can only take discrete values (eg integer) rather than continuous.

The first is an engineering problem which is outside the scope of this report. In practice, as we have seen with system strength constraints, even the most complex security issues can be formulated as constraints, even if in a rather approximate or heuristic form. At worst, the most problematic security problems can be avoided simply by denying permission to connect the plant that causes the problem.

The second problem might be fundamental or contingent. For example, it is fundamentally impossible to make the commitment status of a slow-start plant a controllable variable, because the start time is far longer than the 5-minute dispatch interval. On the other hand, a problem of fast frequency response

¹³ although in practice dispatch targets are specified to the nearest MW

(FFR), say, being uncontrollable simply reflects the fact that the appropriate arrangements for defining, offering and dispatching FFR have not been implemented in the current NEM design.

Any convex mathematical formula can be locally approximated to an equivalent linear formula for use by NEMDE. Indeed, this is done routinely by AEMO for non-linear security constraints: eg transmission stability constraints. However, non-convex constraints can be problematic.

Finally, some controllable variables are naturally integer: for example, the commitment status of a fast-start plant. These variables could be dispatched in principle, but not using an LP as the clearing engine. A more sophisticated *mixed-integer* program would be needed.

2.3.2 *Missing Spot prices*

As noted, any system service that is controlled in dispatch will have an associated shadow price that *could* be used to set a spot price for that service; whether or not to price the service in this way is a market design decision. If the controllable *supply* of a system service is *not* priced, it might be under supplied, leading to a violation of the associated NEMDE constraint and an insecure system. In that case AEMO would need to intervene to ensure the service is provided. An example of this is a constrained-on generator, which effectively provides a system service of virtual transmission capacity to a load pocket¹⁴. Shadow prices for this service are calculated in dispatch¹⁵ but are not currently used in dispatch settlement¹⁶.

On the other hand, if the controllable *demand* for a system service is not priced, the excess demand for that service *can* be managed by AEMO in dispatch. So this does not create a security issue. However, because there is no clearing price payable for use of this service, dispatch efficiency might be reduced. An example of this is that intra-regional transmission capacity is not priced in dispatch and so the resulting congestion is managed by AEMO constraining-off generation in export-constrained zones. This is done with an administrative algorithm which might be less efficient than if demand had been rationed by applying the shadow price.

2.3.3 *Resource Inadequacy*

At times, there may be no feasible solution to the NEMDE constraint set, due to inadequate resources being offered to the market. This is distinct from the missing spot price problem: prices are paid, but supply is inadequate no matter how high the spot price is. This inadequacy might be simply down to insufficient physical assets having been constructed or available, or might be due to other NEMDE constraints - eg transmission constraints - limiting their output. Where there is resource inadequacy, a secure solution can usually be found by shedding load, but at the expense of reliability, obviously.

¹⁴ a large load, remote from the regional reference node, which has limited transmission capacity serving it

¹⁵ this would be manifested in high nodal energy prices

¹⁶ In some cases, a TNSP is responsible for ensuring that this “virtual capacity” service is provided and would contract with a generator to provide this, through a so-called *network support agreement*. The contracted generator will then offer this service into dispatch, notwithstanding that it is not paid *in the spot market* for providing it.

2.3.4 Dispatch Myopia

NEMDE is set up to find a secure and economic dispatch that supplies the five-minute ahead forecast demand. It is not concerned with meeting demand, and achieving security, *beyond* this horizon and may have chosen a different, but more secure or more efficient, dispatch if it was.

For example, consider a ramp rate problem, where demand increases rapidly over an hour or two leading up to the evening peak. Many generators have limited ramping capabilities, and these are expressed in offers and taken into account in NEMDE. A high ramp rate in demand might be met only if a large part of the generation fleet is ramping simultaneously. On the other hand, a myopic dispatch might have already ramped up much of that fleet to full output prior to the demand ramping, making it difficult, or expensive (eg in having to start up expensive fast-start plant to provide additional ramping capability), to meet the ramp when it comes.

An example of this is illustrated in figure 3, below.




Figure 3: myopic dispatch outcome

In this simple example, a myopic dispatch has left cheaper baseload generation at full output during the midday trough in demand and deloaded more expensive mid-merit instead. This obviously give a least cost solution over the trough. But the run-up in demand towards the peak is then too fast for mid-merit to follow, so expensive peaking plant must be dispatched as well to follow the load.

If dispatch had been able to look further ahead during the trough, it might have anticipated this. An alternative dispatch, which avoids deployment of peaking generation, is shown in figure 4, below.




Figure 4: optimized dispatch

By deloading the baseload somewhat over the trough, the combined run-up rates of the mid-merit and baseload are now sufficient to follow the demand ramp, avoiding the need for peaking generation; albeit at the cost of substituting some baseload output for more expensive mid-merit output. So this alternative dispatch might be cheaper overall. The possibility of a less myopic dispatch engine is considered in section 2.4.4.

Figure 3, above, illustrates not just a myopic dispatch but also a *myopic bidding strategy*: by the mid-merit generators in particular. Their strategy leads them to operating below capacity at the start of the peak period when spot prices are likely to be high. They could avoid this by bidding at a price *below* the baseload bids at the end of the trough period, leading dispatch to load them higher and deload some baseload generation. This is illustrated in figure 5, below.

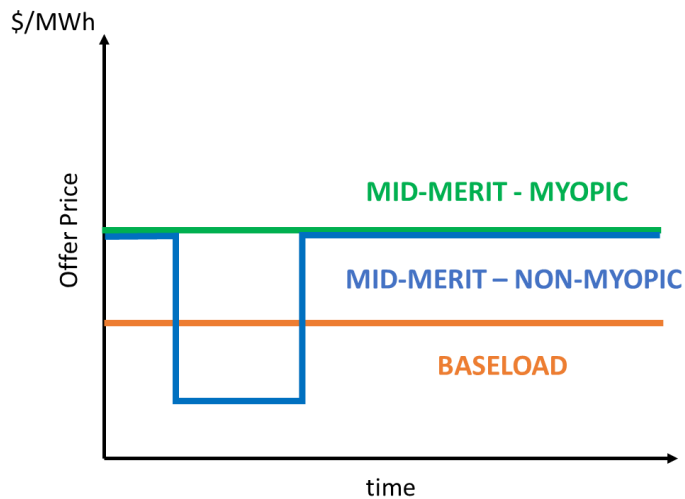


Figure 5: myopic and non-myopic bidding

This would lead to a dispatch outcome similar to figure 4. So *non-myopic* bidding is potentially an alternative solution to this problem of myopic dispatch. Bidding strategies are formulated during the PD process, so this non-myopic bidding is discussed further in the PD chapter.

2.3.5 *Thirty-minute Settlement*

Currently, energy is settled on a thirty-minute basis¹⁷. But dispatch is run, and shadow prices created, each five-minutes. Therefore, shadow prices cannot be used directly to set spot energy prices, and these are instead based on half-hour averages of the five-minute shadow prices.

This issue will be addressed when the five-minute settlement (5MS) rule change is implemented. The reason for mentioning it here is that current dispatch behaviour and outcomes are based on thirty-minute settlement. Thus, it may be hard to know which difficulties seen in the NEM currently flow from thirty-minute settlement – and thus will be addressed with 5MS – and which have other causes and so will continue. This makes it hard to assess what additional NEM design changes are appropriate for post-2025.

2.3.6 *Fast-start Plant Dispatch*

Unit commitment is an *integer* decision: a unit is either started or not started. In principle, commitment of a unit is a controllable variable in dispatch for units that can be started up within the 5-minute dispatch interval. However, as noted above, some integer programming capability would be required in NEMDE to do this efficiently.

Notwithstanding this, NEMDE *does* currently make commitment decisions for *fast-start plant* (FSP), which are in fact defined as units that can start-up with 30 minutes, rather than five. A committed unit receives a start-up instruction which it must follow. However, because dispatch is myopic – looking only 5-minutes ahead – and does not have full integer programming capabilities, these commitment decisions cannot be made efficiently. Instead, NEMDE naively commits the FSP simply based on it being in-merit at the five-minute-ahead point. As a result, it is common for the relevant generator to cancel the start – by rebidding unavailable – if (based on looking out over a further horizon) it expects it to be unprofitable.

If the start is not cancelled, the relevant unit is locked into a *fast-start inflexibility profile* (FSIP), which dictates how it must be dispatched over the next half-hour or so, irrespective of the economics. The FSIP reflects the operating limitations of the plant.

Together with thirty-minute settlement, this FSP commitment process can create volatility and uncertainty in dispatch and make it more difficult for AEMO and market participants to forecast dispatch outcomes over the next 30 minutes or so.

¹⁷ in contrast, FCAS is settled on a five-minute basis, so this issue doesn't arise

2.4 POTENTIAL SOLUTIONS

2.4.1 New Market Ancillary Services

As noted above, AEMO may have established – from an engineering perspective – a particular constraint that must be satisfied to maintain system security, but be unable to incorporate this into NEMDE because the relevant variables are not currently controllable in dispatch, are not continuous or cannot be formulated in a NEMDE constraint with a linear LHS.

Synchronous inertia is a good example of a system service which is missing from NEMDE. It has generally not created security issues historically because it is provided automatically when synchronous generation is online to provide energy and FCAS. However, as synchronous generation is being progressively displaced by new non-synchronous generation, security risks are emerging, which could potentially be addressed if inertia could be introduced as a new market AS¹⁸.

Inertia constraints would naturally be linear, taking the form:

$$\text{Sum of unit synchronous inertias} \geq K$$

However, because inertia depends upon commitment status, it is neither continuous (it is all or nothing) nor controllable in dispatch. So it is not immediately clear how it could be made into a market AS.

Comparison with the existing FCAS service is instructive. FCAS supply similarly depends upon the commitment status of the relevant unit. However, it also depends upon the “commitment status” of the unit governor system that provides frequency response and so the FCAS service. For most units, governor units can be switched on or off in real-time, so FCAS can be dispatched by instructing the governor to be switched on: this is referred to as *enablement*. Enablement is, in effect, an integer decision, but the FCAS dispatched can be effectively treated as a continuous variable because:

- The amount of FCAS supplied will depend also on the dispatch energy output, so this gives some degree of continuous control;
- FCAS from each of these enabled units is typically small relative to the total requirement, so the provision is not too “lumpy”
- It doesn’t impact system security if an excess amount of FCAS is dispatched: eg because of this lumpiness.

Indeed, this means that it doesn’t actually matter if a unit governor system can be switched on or off in real-time. Where this is *not* possible¹⁹, the governor can be left on continuously. It might then be providing FCAS service continuously, but is nevertheless only paid for this when it is dispatched and enabled.

Similarly, the essentially on-or-off nature of physical inertia provision need not prevent it being treated as a continuous controllable variable, placed on the LHS of the NEMDE constraint. When dispatched, it

¹⁸ Alternatively, system security can be maintained by limiting demand for inertia, using a “controllable demand < uncontrollably supply” constraint form, and this has been incorporated into NEMDE when required

¹⁹ which was the case, at least initially, on some older units

could be paid a spot price based on the shadow price of the binding constraint. Of course, the physical inertia is provided continuously. But, as with FCAS, it would only be paid for when it is “enabled”.

So, for example, a synchronous unit with 1000MWsecs of synchronous inertia could offer this amount across several price bands, similar to energy, and NEMDE could enable some or all of this, depending upon the need for inertia and the offer prices of competitors. Generators would be paid the spot price (based on the shadow price of the inertia constraint), only for the inertia that is “enabled” by NEMDE.

This concept should be seen in a possible future context where there may be alternative sources of inertia, or effective substitutes for inertia, other than from slow-start synchronous units: for example:

- Fast-start synchronous units²⁰ ;
- synchronous compensators
- virtual inertia: eg provided by grid-forming inverters
- fast frequency response: which might be a substitute for inertia in terms of system security

So the relevant NEMDE constraint equation might take the form:

$$A \times \text{synchronous inertia} + B \times \text{virtual inertia} + C \times \text{FFR} \geq K$$

NEMDE would dispatch a combination of these three services to satisfy the security constraint. Since the services are substitutes, they would be paid a common spot price: adjusted by the A/B/C coefficients.

2.4.2 Pricing Uncontrolled Services

An uncontrolled service – appearing on the RHS of a NEMDE constraint - can potentially be priced and paid (or charged) in dispatch, using the shadow price of the NEMDE constraint that it appears in. Indeed, this is how load is priced in the current NEM design.

So, for example, synchronous inertia could be rewarded whilst remaining *uncontrolled*. In this case, the equation above would be reformulated as:

$$B \times \text{virtual inertia} + C \times \text{FFR} \geq K - A \times \text{synchronous inertia}$$

Since the variables on the LHS and RHS represent substitutes, it makes sense for them to be paid the same price, whether they are controllable or not. Whilst owners of synchronous units cannot respond in real-time to the inertia spot price, they can nevertheless respond to *forecasts* of the inertia price in PD and commit units accordingly during the PD process. Of course, in making such commitment decisions, generators would consider forecast prices for energy and FCAS as well as inertia and look for commitment options that co-optimize profitability across all these services. For example, the inertia spot price would need to be high enough to offset any costs of operating at a loss in the energy market.

For payments to be made to providers of uncontrolled system services, the RHS of the constraint must be linear in the corresponding variables. For inertia, the RHS depends on a simple sum of the supplied inertia quantities, so the contribution of each service provider to the RHS value is clear and payment can be made accordingly. For system strength, on the other hand, as currently formulated, the RHS

²⁰ eg in hydro power stations

currently depends upon a complex function of the on-line status of synchronous generators, based on combinations of units:

$$\text{Aggregate output of non-synchronous generation} \leq f(\text{synchronous unit statuses})$$

In this case, it is not clear how payment for the system strength service should be defined. For example, with three synchronous units X, Y and Z on-line, the RHS might take a value of 1000MW. So, perhaps, the three units should, together be paid for 1000 units of system strength: ie paid 1000 times the shadow price of this constraint. But how should this lump sum be allocated between the three units? How much would unit X be paid, say? This is unclear and could be quite contentious. A generator considering committing a unit based on forecast system strength prices in PD would not have the same clarity or certainty around how much it would be paid in dispatch as was the case for inertia providers.

It is possible that, in the future, AEMO engineers will work out a way to define a new “system strength service” whose requirements can be expressed in a NEMDE constraint with a linear RHS formulation: eg

$$\text{Total Non-sync output} \leq K \times (\text{sum of system strength service amounts})$$

In this case, each synchronous unit could be paid the shadow price of this constraint on the amount of this new system strength service that it provides. Or, plausibly, this system strength might be treated as pseudo-controllable, as system inertia was in the previous section, and so be bid and dispatched, being placed on the LHS of the NEMDE constraint equation.

So, in summary, it might be feasible to value and pay new system services in dispatch, but only if these can be formulated in NEMDE constraint using a linear expression on *either* the LHS or RHS of the relevant security constraints. This is illustrated in figure 6, below.

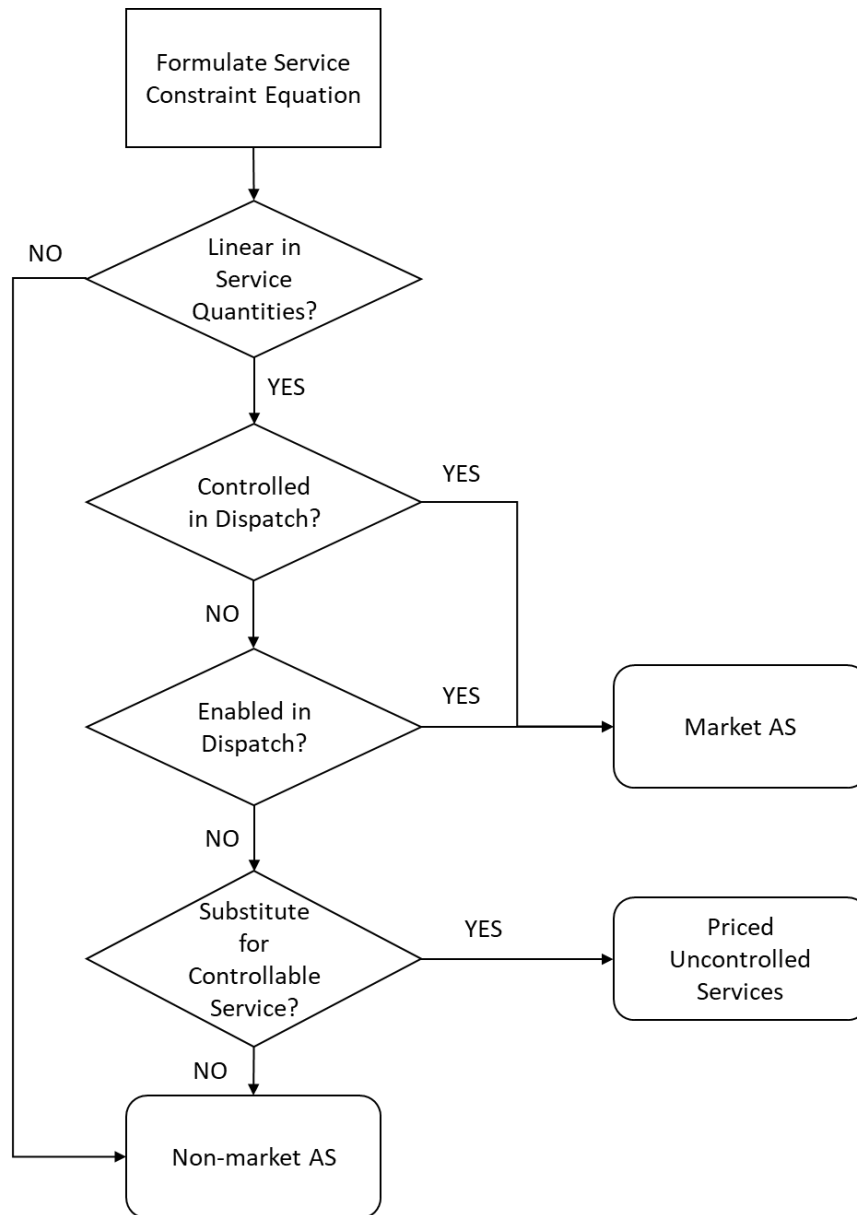


Figure 6: Spot Pricing of Ancillary Services

2.4.3 “COGATI” Issues

In the context of considering the pricing of essential system services, a brief discussion of COGATI issues may be worthwhile. The COGATI MDI is considering the possible introduction of nodal energy pricing for at least a part of the market. One way of thinking about a nodal energy price is as a combination of a regional energy price and spot transmission prices. The latter reflect the value of transmission in the spot market and will be zero when transmission is uncongested and positive when transmission is congested; thus, these spot prices are typically referred to as *congestion prices* and this terminology will be used. However, in the context of pricing essential system services, it is useful to think of *transmission* as the essential system service and the congestion price as being the shadow price for that service.

As discussed in the previous sections, some adjustments to the definition of inertia or system strength services, or how they are formulated in NEMDE, would be needed to incorporate them as new market AS in the NEM design. However, in the case of transmission/congestion, these services are *already* formulated into NEMDE and shadow prices are already being calculated in every dispatch run. So, the issue for COGATI is *whether* – not *how* – to price congestion. This is not to say that pricing congestion is a straightforward design issue. It raises several important and difficult secondary issues that would need to be addressed. Who should receive the revenue from congestion pricing? What commercial risks are created and how can those risks be hedged? Should pre-existing market participants be grandfathered against these new charges? And so on. These are the issues that the COGATI MDI is concerned with.

It is beyond the scope of this engagement to analyse these COGATI issues or to offer potential solutions. However, it is worthwhile identifying the conceptual and practical overlap between pricing congestion and pricing other system services. Pricing, say, system strength or inertia might raise similar issues that need to be addressed. Indeed, when it comes to congestion, the delineation between transmission and other system services may be unclear: a system strength constraint and a transmission constraint may take the same generic form and it might not even be meaningful or feasible to distinguish between them. So the introduction of each new market AS would raise a “mini-COGATI” of its own. The market design problem is not *just* about the formulation and pricing.

2.4.4 Multi-step Dispatch

To address the problem of dispatch myopia, the dispatch calculation could use a longer horizon (up to one hour ahead, say) by using multiple, linked timesteps. The objective of NEMDE would then be to find the lowest cost (or highest value) dispatch solution that meets the demand forecast and satisfies the security constraints, across the entire *study period*. This would imply *co-optimising* across *all* of the timesteps in the study period. The solution might well give a *higher* cost dispatch in the first timestep, compared to the solution that the current myopic dispatch finds, justified by it positioning dispatch to then find *lower* cost solutions in the subsequent timesteps.

For example, suppose that a short-duration battery were able to submit its current amount of stored energy into this multi-step dispatch, together with its normal offer price. A myopic dispatch would have no knowledge of the energy constraint and simply dispatch the battery in merit, according to its offer. A multi-step dispatch might potentially withhold that battery output from the first timestep - dispatching a more expensive offer instead - because it anticipates that the battery’s energy has a higher value in a later timestep and it “dispatches” it then instead.

The multi-step dispatch would employ NEMDE constraints as now, although they would separately apply to each 5-minute timestep within the longer problem, reflecting the forecast demand or system conditions in each period. Consecutive dispatch targets would be linked through ramp rate limits. However, additional generator information – such as energy constraints – could be submitted in offers, and additional constraints would be included in NEMDE to process these.

A multi-step dispatch might be better at committing FSP, given it is able to consider its economics over a fuller period of its operation rather than just the first five minutes. Commitment efficiency might even be improved further by allowing FSP to submit start-up costs etc and improving the integer

programming capabilities of NEMDE. However, this extra complexity may introduce some concerns around transparency and spot pricing. These are described in the pre-dispatch chapter, where central commitment of slow-start plant is also considered.

The multi-step dispatch solution would provide a set of dispatch targets covering the study period. However, generators would not *actually* be dispatched in this way. Rather, they would be dispatched in accordance with the targets for the *first* timestep only. The multi-step dispatch calculation would then be repeated five minutes later, with the results for the first timestep from that run then being used for the next set of dispatch targets, and so on.

This is illustrated in table 1, below, assuming a shorter 30-minute study period for simplicity of exposition. The unit X presented has ramp rate limits of 2MW/minute up and down, so can only change by +/-10MW over each dispatch timestep. Each dispatch run produces a set of 6 consecutive dispatch targets for the unit, covering the next 30-minute period following the run.

Target time	16:00 run	16:05 run	16:10 run	16:15 run
16:05	100			
16:10	110	110		
16:15	120	120	110	
16:20	130	130	110	100
16:25	140	140	110	90
16:30	150	150	110	80
16:35			110	70
16:40			110	60
16:45				50

Table 1: dispatch targets (MW) for unit X in consecutive runs of a multi-step dispatch

At the time of the 16:00 and 16:05 runs, a high peak demand is being forecast and NEMDE is ramping up the unit’s output as fast as possible. However, the demand forecast is then lower for the 16:10 run and further lowered for the 16:15 run. The dispatch targets from these later runs reflect this. Only the first target in each run – highlighted in red – is issued to the generator for it to follow.

Each dispatch run would also produce a corresponding set of spot prices for the study period. Similarly, only the spot price for the first timestep of each run would be used in settlements.

For each run, the sets of dispatch targets and spot prices are self-consistent, in the sense that the spot prices would be clearing prices. So if a generator were dispatched in accordance with this solution, it would be paid an amount which equalled or exceeded its offer price. But only in *aggregate* over the period, not necessarily within each timestep considered separately. As illustrated by the example above, demand forecasting errors will lead to changes to this dispatch. So, unit X is *actually* dispatched

{100, 110, 110, 100} over the first four periods after 16:00, whereas the 16:00 run produced targets of {100, 110, 120, 130}. The clearing prices produced in this 16:00 run will be consistent with the latter set of targets but not the former.

Therefore, with multi-step NEMDE, there is a risk that a generator might be dispatched “at a loss” (ie below its offer price). The unit X might have been theoretically “made whole”, in the 16:00 run, through high spot prices and high output in the 16:30 dispatch period, say, to offset lower spot prices in the earlier timesteps. However, the change in forecast means that these high prices never eventuate and the generator bears the consequent loss. These risks could perhaps be mitigated by having some ahead clearing occurring alongside the multi-step dispatch. This design option is discussed in the ahead market chapter²¹.

This risk of a loss has similarities to a slow-start generator committing based on high forecast prices in PD, only to see lower than expected demand and low outturn spot prices. This issue is discussed further in the pre-dispatch chapter. So, risks of this kind are inevitable when decisions are made based on forecasts. However, there is a difference between losses incurred because of a *generator’s* decision in PD, compared to losses caused by a *NEMDE (AEMO)* decision in dispatch. The latter might be considered less acceptable than the former.

As discussed in section 2.3.4, dispatch myopia only leads to inefficiency when combined with bidding myopia. Put another way, the efficiency benefits from a multi-step dispatch might be obtained instead by ensuring or encouraging non-myopic bidding in the PD process. Similar risks arise for generators if forecasting errors lead them to bid in a way that they regret in hindsight. But at least bidding risks are under their control, rather than dependent on a multi-step NEMDE process.

In summary, this multi-step dispatch option needs to be compared against alternative PD designs which may address the same “myopia” issue: and possibly more effectively or transparently.

2.4.5 Two-sided Market

For completeness, it is worth considering the design option of a two-sided market, where some or all load must be bid into dispatch, by retailers, similar to how generators bid currently. Since this would incorporate the load of small consumers, it would make sense for this reform to encompass small²² generators too. There is also the question as to whether semi-scheduled generation should be responsible for forecasting their five-minute output in dispatch.

Similar to adding new ancillary services, a two-sided market would introduce new controllable variables (ie customer load) into NEMDE, and so improve the efficiency, reliability and security of dispatch: assuming, of course, that retailers can conform adequately to their dispatch instructions.

²¹ See section 4.4.4

²² eg currently non-scheduled

2.5 CONCLUSIONS

Weaknesses emerge in the dispatch design when the engineering exigencies of maintaining system security are not fully complemented by market frameworks to define, price and procure the necessary services from market participants. Meaning that the two facets of dispatch, described in the introductory section, are not in harmony and balance. Any anomalies will mean that the market does not provide the services required, creating the need for the system operator to intervene either to requisition additional supply of, or to administratively ration the demand for, these services.

These issues are being considered by the three ESB MDIs concerned with dispatch design: “essential system services”, “two-sided markets” and “COGATI”. It is unclear what issues might be left for PD and ahead markets to deal with once these other MDIs have completed their tasks. We will nevertheless press on into the unknown, considering PD next.

3 SCHEDULING AND PRE-DISPATCH

3.1 INTRODUCTION

Scheduling occurs ahead of dispatch, as generator and storage operators position their plant to be ready to operate as needed in the spot market. Like dispatch, scheduling can be considered from both an engineering and an economic perspective. In traditional regulated utilities, it is a long-standing engineering problem, and scheduling tools and algorithms have been developed over a long period, growing in sophistication and complexity over the years as the computers needed to support them have become more powerful.

But in a generation market, scheduling essentially becomes a commercial proposition for market participants. Each firm aims to best deploy its assets so as to maximise its profitability, whilst managing associated risks, in the spot market. And since each firm does this scheduling independently and autonomously, there is an implicit need to coordinate these schedules.

PD provides a platform for this coordination. And in this process, AEMO is again torn between its market operator and system operator roles. It provides the PD platform in which this coordination plays out, whilst always being ready to intervene in its system operator role when things appear to be heading towards an insecure dispatch outcome.

Like market participants, AEMO needs to prepare ahead of dispatch to be able to deploy its interventions effectively. So AEMO must undertake a kind of scheduling process itself, similar in some ways to those of market participants, albeit with a quite different objective of system security rather than profit maximisation. But AEMO is nevertheless, in a sense, a participant within its own platform, a potentially confusing and conflicting state of affairs. Once again, these separate AEMO roles must be carefully identified and disentangled.

Dispatch is an automated process operating in real-time, which limits the opportunity for innovation and variety in market design. Every market around the world employs essentially the same dispatch process, in an engineering sense, although the spot markets built around this core may vary somewhat. In contrast, scheduling takes places over a more extended timescale – from day-ahead into real-time – and inevitably involves a mixture of manual and automated processes. Thus market designers around the world have had the opportunity to come up with a variety of ways to design scheduling processes. And, as with dispatch, the tensions between “hands-on” and “hands-off” roles for the market/system operator play out in the varying philosophies and approaches.

The NEM PD process is probably at the “hands-off” end of this spectrum, at least in comparison to US market designs²³. But this has created concerns, for the ESB at least, that dispatch security could be compromised without AEMO at the helm during this key period leading up to dispatch. So the options that the ESB is considering – specifically the ahead market ideas – have the underlying objective of giving AEMO more visibility and control of the scheduling process.

²³ See the Appendix.

3.2 CONCEPTS

3.2.1 *The Pre-dispatch Engine*

At the heart of the PD process is the PD engine. This is designed to emulate the dispatch engine. That is, it takes in the same offer structures and employs the same NEMDE constraints²⁴. However, it operates over different time horizons, going from real-time to the next day, in 30-minute rather than 5-minute timesteps.

While the PD engine clears over a study period covering multiple consecutive dispatch intervals, it does not co-optimize across these multiple timesteps in the way envisaged for a multi-step dispatch engine, discussed in section 2.4.4 above. Rather, it has the same myopic approach as dispatch, solving one just one timestep at a time, with one step linked to the next only through ramp rate limits constraining the difference between consecutive PD targets.

Also like the dispatch engine, the PD engine does not make commitment decisions (except for fast-start plant), with these implicit in the bids that market participants submit to PD, as discussed below. Unlike dispatch, it *does* include energy constraints. However, these are employed myopically, with energy simply being used until it is exhausted, rather than deliberately being conserved for use in the highest price periods.

3.2.2 *Good Faith*

A key requirement in PD is that bids are submitted in *good faith*. This means, essentially, that these bids are genuine *forecasts* of how the generators expects and intends to bid into dispatch itself. Like any forecast, these can change when new information becomes available, and the good faith principle requires that PD bids are updated accordingly, as soon as practical. Conversely, if no new information arises – and so the forecasts (eg weather and demand) relied upon in making the original bid turn out to be 100% accurate – then the bid is not permitted to change. In this sense (and *only* in this sense) a PD bid has a kind of firmness.

The term bid (or offer) itself is somewhat misleading, in that a bid implies an associated market or auction and some financial or physical consequences if the bid is cleared. But no *clearing* happens in PD. It is purely for information only. So, it would be more accurate to refer to a PD bid *forecast*. A bid forecast is not a bid, any more than a load forecast is load. Nevertheless, I will refer to PD “bids”, as this is the accepted terminology. This distinction is key when considering ahead markets, where genuine bids *are* involved.

3.2.3 *Gate Closure*

It is perhaps worth mentioning that, although bids are used in the dispatch process, bidding itself is part of the PD process, *not* the dispatch process. Bids can be changed (ie rebids submitted) at any point up

²⁴ although the formulation of these constraints may be different. For example, in dispatch, NEMDE employs so-called *feedback* constraints, which utilize the *headroom* on transmission lines, calculated by comparing the maximum secure flow on the line to the latest *measured* flow from SCADA. Because PD operates ahead of time, such measurements are not available and different formulations have to be used

until real-time, in accordance with good faith requirements. The latest bid or rebid submitted is then input into NEMDE in the dispatch process. So the bid forecasts in PD do eventually become *actual bids*, but only in the dispatch process.

The point at which bid forecasts are locked in and become actual bids is known as *gate closure* and this represents the interface between the PD and dispatch processes. The NEM – unlike some other market designs – has gate closure as *close* to real-time as is practical, from an IT perspective.

3.2.4 Scheduling

Scheduling refers to decisions and actions taken in advance of real-time to maximise opportunities and profitability in real-time dispatch²⁵. Key scheduling decisions and actions in this context are:

- Commitment and decommitment of slow-start plant
- Cycling of short-duration or daily storage
- Calling of customer demand response
- Procuring gas or other fuel on (say) day-ahead markets

Some decisions, such as outage scheduling, can be taken months or even years ahead of real-time. However only scheduling decisions taken within PD timescales are considered here: ie between day-ahead and real-time.

As already noted, scheduling decisions are made by market participants, *not* PD, with the latter being presented with these as a *fait accompli* within bids: for example, a generator deciding to start its units at 6am would offer unavailable in PD prior to this and would also offer after 6am in a way that ensured that is dispatched at or above its minimum generation. I will refer to those making these decisions in each generating company as *traders*.

Each firm must make its scheduling decisions independently of its competitors: this is a practical but also (potentially) a legal requirement under competition law. In doing so, it will have the objective of maximizing its own profitability, subject to appropriately managing its financial risks.

In summary, the scheduling carried out during PD is *decentralized*. Each trader orchestrates and coordinates its own resources so as to optimize the position of its firm. Each trader acts independently; there is no single person or algorithm responsible for scheduling the entire generation fleet.

3.2.5 Coordination

Because there is no central scheduler, a process is needed to coordinate the various individual schedules, to ensure they are efficient, reliable and secure. This is the role played by the PD engine. Each trader expresses its schedule through its PD bids and the PD engine then uses these bids to generate *forecasts* of dispatch outcomes: cleared quantities and spot prices. These outcomes will indicate where the overall schedule is poor and provide opportunities for traders to change their PD bids

²⁵ confusingly, the word “scheduled” is a defined term used in the NEM to convey the same meaning as the “controllable” concept discussed in the previous section. So, a *scheduled generator* is one that must offer into the NEM and comply with dispatch instructions. The term is used in this report in the sense of its usual English language meaning of making decisions ahead of time.

to improve their own profitability and (consequently) improve the overall schedule. For example, if insufficient capacity has been committed for some part of the PD period, the PD engine will produce high spot prices that reflect this. This will encourage traders to commit more plant over this period.

Coordination requires iteration. If the first PD run indicates a shortfall and high prices, traders will commit more plant and will alter their bids to reflect this. These rebids are then input into a fresh PD run. There may still be not an optimal amount of plant committed – it might be too much or too little – so more rebids and more PD runs are needed.

Therefore PD needs to be re-run repeatedly and this occurs in the current NEM design through PD being re-run each 30-minutes, with the study period rolling forward by 30-minutes each time.

Iteration is a common feature of many *centralized* scheduling algorithms – those where scheduling decisions are made centrally in the clearing engine rather than separately by traders. Some of these algorithms may even use an internal shadow price to guide and coordinate decisions similar to how this occurs in the PD process. So this PD process can be considered to be analogous to a *mega-algorithm*, incorporating not just the PD engine but all of the traders and the various engines and decisions support tools that *they* are using. A single (notional) algorithm, but *not* a monolithic algorithm designed and operated by a central scheduling entity. Rather it is a kind of ecosystem: combining a central, regulated algorithm (ie the PD engine) and decentralized, non-regulated trader processes.

This mega-algorithm is illustrated in figure 7, below. It is seen that the inputs to this mega-algorithm are the costs and constraints of the individual generators, which is similar to a centralized scheduling algorithm. However, the mega-algorithm’s decentralized architecture allows these costs to remain private within each firm.




Figure 7: the pre-dispatch scheduling architecture

3.2.6 Tracking

Each iteration of PD incorporates new information that has become available since the last run: AEMO updates its forecasts; traders will also update their private information (eg on the technical availability of their own plant) before re-running their schedules. So PD iteration is performing two roles: *tracking* these changes as well as provide an opportunity for the mega-algorithm to converge to an optimal solution. This differs from a conventional, centralized algorithm, which will keep iterating – using a *constant* set of inputs – until a solution has been found.

To use an analogy, the PD process is operating like a driver on a winding road. The driver is constantly adjusting the wheel, to correct his position on the road but also to take the best line into the corners up ahead as these are revealed to him. In contrast, on a dead straight road, the driver is doing the former but not the latter, like a conventional scheduling algorithm.

Effective tracking becomes particularly important when there is more forecasting uncertainty: the metaphorical road becomes more winding. A PD process that does not track well may end up a long way from the optimal solution when it comes to real-time dispatch.

3.2.7 Competitive Dynamics

Generator bids in the NEM are not required to be cost-reflective. In fact, generators are free to bid anywhere between the market price floor and the market price cap. Whilst a generator might be profitable in dispatch even when bidding at cost (because the spot price is set by another generator with a higher bid price) it may be able to improve profitability by bidding at least part of its capacity above cost in order to cause a higher spot price to be set. However, in a competitive market, higher bids are less likely to be called, so each generator will have to trade off a higher price against a loss of market share.

The iterations of the PD process will allow the competitive dynamics to play out. For example, generator A might reduce its offer prices in order to increase its output over a high price period seen in PD, and in so doing undercut generator B who sees a forecast loss of output as a result. Generator B might then respond in the next iteration by dropping its offer price further, and so on. Hopefully, this process will eventually converge prior to dispatch, reaching a kind of Nash equilibrium, where no generator can see a way to increase its profitability by changing its bid.

Again, this is a dynamic that does *not* exist in a conventional scheduling algorithm, where production costs (or bids reflecting them) are simply presented to the scheduler each time it is run, rather than (potentially) progressively revealed as the scheduling process unfolds. This is an important distinction when we are considering alternative scheduling processes.

This is not to say that this form of competition only occurs through an iterated PD process. Even a “one-shot” scheduler will be repeated day after day, which provides a potential platform for a similar dynamic to occur. And it is also not necessarily the case that such iteration promotes competition. In some

markets, such repetition can provide a signalling platform for tacit collusion between large companies²⁶, although there is no suggestion that this occurs currently in the NEM.

3.2.8 System Insecurity

Since PD outcomes are simply forecasts of dispatch outcomes, insecurity in PD prefigures insecurity in dispatch, giving AEMO or traders time to rectify the problem prior to dispatch. How and whether this is done depends on the underlying reason for the insecurity.

As discussed in section 2.3, insecurity might arise for three reasons:

- NEMDE incompleteness
- Missing prices
- Resource inadequacy

If NEMDE is incomplete then, since PD also uses NEMDE, the process outlines above will not by itself reveal any insecurity problem. AEMO would need to run a separate process, using a more detailed model of power system security, to assess and identify insecurity. It would then need to take action outside of the PD process to remedy the insecurity.

If the problem is missing prices, and a consequential shortfall in supply of the associated system service, this will be revealed by a constraint violation in PD and an associated extreme shadow price (eg set at some constraint violation penalty). But, because this shadow price has no financial impact on participants, they will not be motivated to rebid to increase supply of this service and, again, AEMO would have to take action outside of the PD process to address the problem.

Finally, if the problem is simply one of resource inadequacy, the high shadow price will be seen by the market as a high spot price and traders will be incentivized to respond by scheduling more resources to the market, which should help to address the insecurity.

These possible mechanisms for identifying and rectifying insecurity in PD are presented in figure 8, below.

²⁶ For example, between petrol retailers.

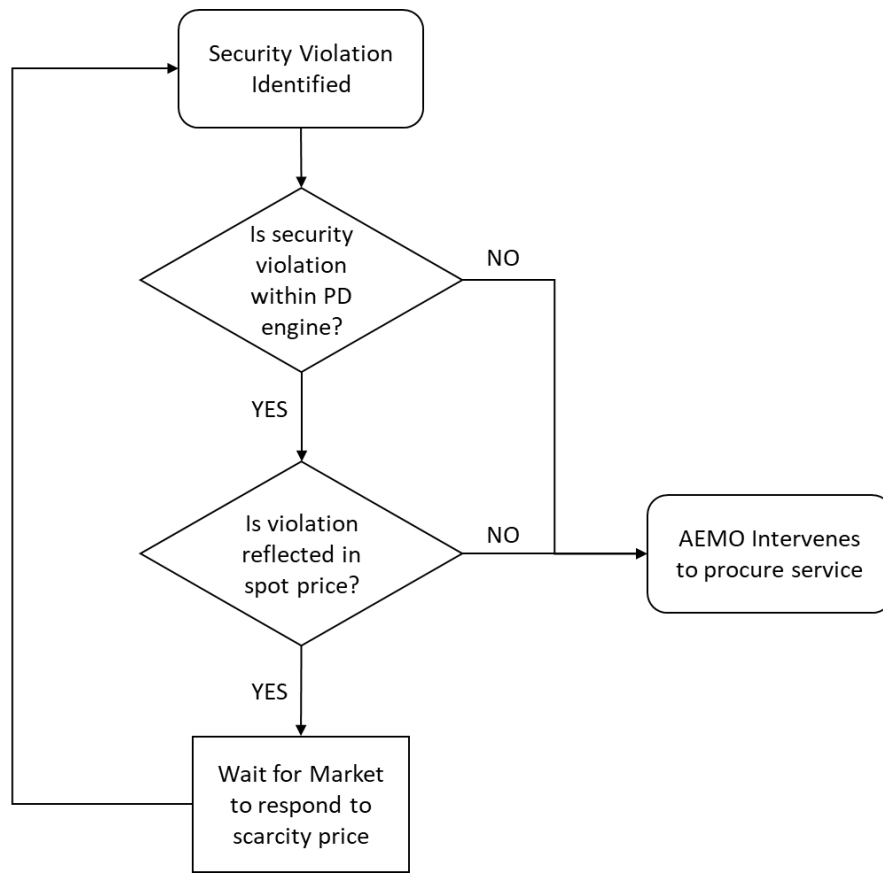


Figure 8: AEMO Intervention

3.2.9 AEMO Scheduling and Intervention

Given the risk of insecurity in PD not being rectified by market responses prior to dispatch, AEMO is given additional tools and powers to address the problem itself by *intervening* in the market. Its current intervention tools include:

- Contracts for non-market AS²⁷
- Contracts for reserve capacity
- Directions

As described in the dispatch chapter, *non-market AS* exist because of NEMDE incompleteness: they are system services that are necessary for security but cannot be incorporated into the NEMDE-based spot market. Instead, a *contract market* is used, with AEMO striking bilateral contracts with the relevant service providers. These contracts provide prices and payments for services, along with operational provisions: how much notice must be given (eg for the AS provider to start-up the relevant plant);

²⁷ The NEM rules would not regard AEMO calling upon non-market AS contracts as “intervention” but I will for convenience include it under this umbrella.

penalties for unavailability; and compensation for other services (eg energy) provided, or not able to be provided, as a result of the AS provision.

For example, AEMO currently procures voltage support as a non-market AS. It contracts with suitably located generating units. Its procurement contracts may allow AEMO to require that these units are committed when they ordinarily would not have been (eg due to insufficient energy prices) and/or to be operated in the appropriate voltage-control mode. A contract would typically require AEMO to pay the plant owner for availability, utilization, direct costs and opportunity costs.

Reserve capacity can be thought of as generation capacity (or a demand response equivalent) which has a levelized cost of energy higher than the market price cap and so which is not viable in the spot market. This is similar to the missing price problem, although in this case the (energy) price is *limited* rather than missing entirely. Under its Reserve Trader (RERT) role, AEMO similarly procures reserve capacity through a contract market. AEMO commonly procures reserve capacity in the form of demand response with industrial customers. The RERT contract would allow AEMO to call upon the counterparty to reduce load, which ensures reliability for other customers in the same way as additional generation capacity would have. Again, payments may be made for availability and for utilization.

Finally, under its directions power, AEMO can require a generator to operate its plant in a way that is inconsistent with its bid, with the generators being compensated for the additional costs incurred. This can be thought of as a generic *regulatory* contract which all generators enter into when they join the NEM. Currently, AEMO regularly directs synchronous units in SA to commit in order to maintain system strength. For directions, payments due to the directed party are defined in the *Rules* – and associated NEM procedures - rather than in a bilateral contract between AEMO and the service provider.

Each of these tools give AEMO powers to schedule and operate plant as though it was its own. In this respect, AEMO uses the forecasts provided by PD to do its *own* scheduling, similar to the market participants. In this sense, AEMO can be considered to be a part of the scheduling mega-algorithm, as illustrated in figure 9, below.

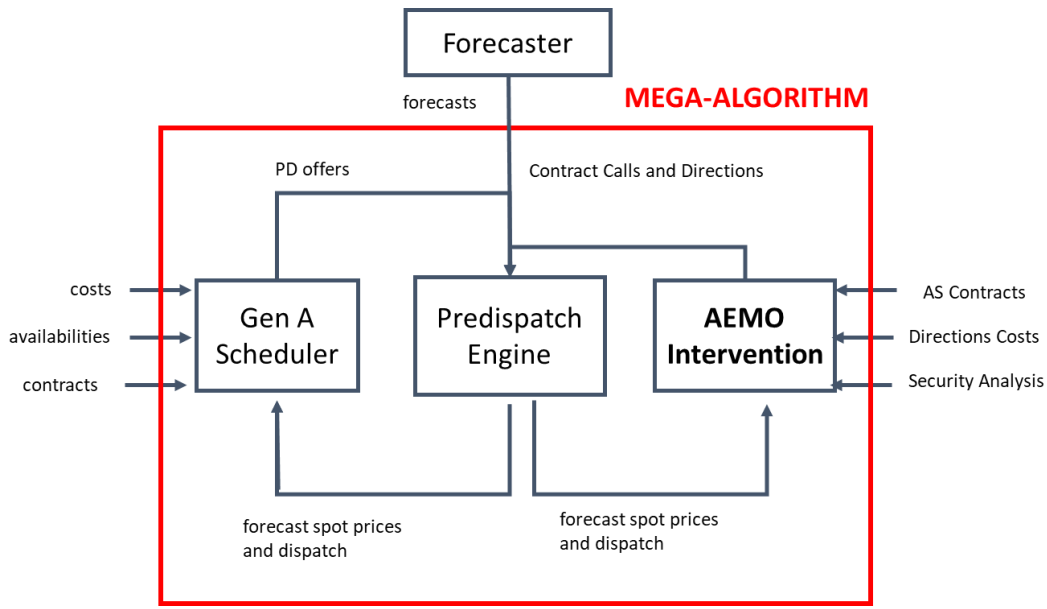


Figure 9:AEMO Intervention as a Component of the Scheduling Mega-algorithm

A key difference, though, is that AEMO has different incentives and objectives to market participants. Rather than seeking to maximise profit, AEMO aims simply to ensure system security, whatever the cost to the market (since AEMO does not bear these costs itself). As a secondary objective, where AEMO has a choice of how it addresses a security problem, AEMO is required to choose the lowest cost option.

3.3 ISSUES ARISING

3.3.1 “Firmness” of Pre-dispatch Bids

A concern raised by the ESB in its March paper is that PD bids are “non-firm”, meaning that they can be withdrawn or altered at any time. This means that PD forecasts of system security cannot be relied upon, because at any time a PD rebid can lead to insecurity. Whilst AEMO has tools to address such insecurity, the concern is that it may not have the time to deploy them once the insecurity has been identified.

However, this suggestion that PD bids might be “firm” or “non-firm” is a category error. *Bids* can be firm or non-firm, but PD “bids” are not actual bids but rather *forecasts*, as discussed in section 3.2.2. It is illogical to say that a forecast is “firm” or “non-firm”. The key is whether the forecast is *accurate* or *inaccurate*. So, there are two ways to interpret this concern:

- The accuracy of PD forecasts needs to be improved; or
- Real bids are needed in the PD period: implying an *ahead market*

Ahead markets are considered in the next chapter. This section will explore further why PD bid forecasts might be inaccurate.

Why might PD forecasts be inaccurate? Firstly, the input forecasts – for weather, consumer demand, semi-scheduled output and so on - that the PD process uses may be inaccurate and so may constantly be revised as real-time approaches. If the PD process works well, it *should* track these changes, giving rise to an apparent problem of poor “PD forecasting” when the real problem is the forecasting of the inputs. Obviously, there would be no point in considering changes to the PD process to address an issue that is external to this process. For example, it is not helpful to hold a bid fixed if changes to forecasts mean that the optimal bid is changed. This is illustrated in figure 10, below.

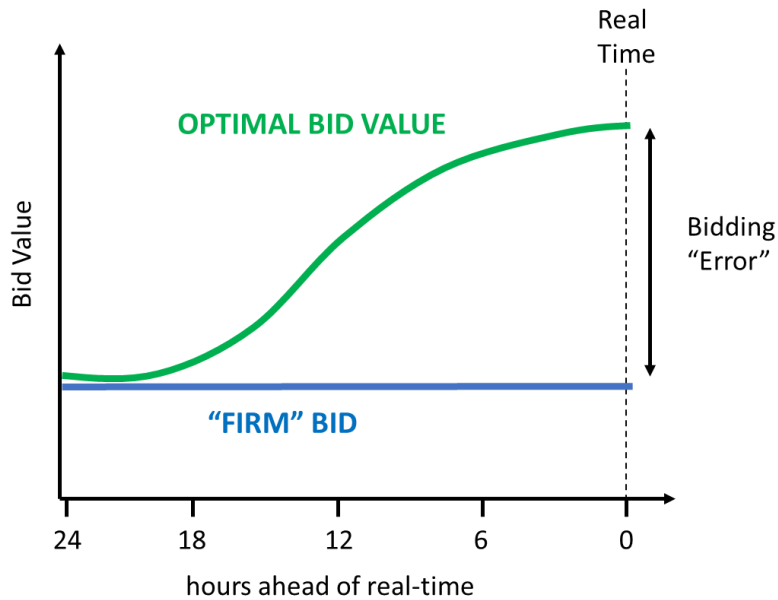


Figure 10: Firming a pre-dispatch Bid

In figure 10, as external forecasts (eg weather and demand) change as real-time approaches, one parameter (eg the quantity of output at a given price, or the amount of pre-charging of a battery, say) of an optimal bid might change correspondingly. “Optimal” might mean for the particular generator or for the market as a whole: as discussed above, competitive dynamics should make these things similar, if not identical.

Because the optimal bid value is changing, any attempt to lock it in at an early stage, for the sake of “firmness” leads to the bid being a long way from optimal when real-time arrives and the bid is used in dispatch. So the appropriate way to improve “firmness” is to improve forecasting so that the green “optimal bid” curve is flattened, *not* to lock-in early bids.

Possibly, there may be problems of *convergence* in the iterative PD process: the mega-algorithm. This might be due to poor convergence in the PD iterations: so PD forecasts are constantly changing from one PD run to the next, even though the input forecasts are unchanged. This is illustrated in figure 11,

below. The situation – and the optimal bid – is the same as in figure 10. But now the bidding “error” is caused by poor performance of the PD process, rather than an attempt to “firm” the bid.




Figure 11: Poor Convergence in Pre-dispatch

Finally, the PD process may be slow to *track* changes in the input forecasts, as illustrated in figure 12, below.




Figure 12: Poor Tracking in Pre-dispatch

So, to see the impact of poor tracking, consider a change in the forecasts that leads naturally (and “optimally” from a market perspective) to an insecure dispatch (because of the NEMDE incompleteness problem). There might be a situation where the forecasts are updated six hours ahead of real-time, say, but because the PD process is slow to track changes, the insecurity is not revealed in the PD forecasts until just 2 hours ahead of real-time, by which time it is too late for AEMO to intervene effectively to restore security. Tracking and convergence are each important characteristics of the PD process, particularly as forecasting becomes more difficult and uncertain.

3.3.2 *Efficiency of Scheduling and Coordination*

Another issue raised in the ESB paper is that the PD process may not be effective for scheduling and coordination: if not now, then perhaps in the future, when new technologies and system services make this problem much more difficult.

This is a different problem to the previous problems of “firmness”, convergence and tracking. Convergence of the PD process might be good, but to the wrong *point*, meaning that scheduling – and hence dispatch – is inefficient or unnecessarily insecure. A problem here would lie with the PD process itself, not with dispatch or with forecast inputs.

It is not clear that this is an issue in the NEM currently. Certainly, there are problems with dispatch security and stability, but these can be explained by the problems in dispatch that were discussed in the dispatch chapter, rather than problems in PD. However, past and current fitness does not *necessarily* imply future fitness, given the substantial changes expected in the generation mix and in demand behaviour over the time period being covered by the ESB review. Indeed, if dispatch problems are addressed – by introducing new AS spot markets and contract markets and by increasing the complexity and sophistication of NEMDE – this might fundamentally change the nature of the scheduling problem: for example by placing greater reliance and dependence on commitment of synchronous generation to provide the new AS.

3.3.3 *Coordination between AEMO and the Market*

As discussed in section 3.2.8, the additional tools provided to AEMO to manage system security essentially give it a scheduling role over the PD period. PD effectiveness relies on coordination between *all* schedulers, including AEMO. However, coordination between the market and AEMO may be poor currently because:

- AEMO is not currently required to participate in PD in the same way as market participants: it has no good faith obligation for example
- AEMO has different objectives than market participants, concerned with security rather than profitability;
- AEMO’s objectives may be unclear to the market - or even to AEMO itself
- AEMO may not have full access to relevant cost information, given that it is scheduling plant owned by other parties.

AEMO's primary objective is to maintain security. However, where there are several options for achieving this, it needs some secondary objectives to guide its choice. These might be to minimise:

- *activity*: ie to only intervene where it becomes clear that dispatch is likely to be insecure and it cannot delay action any longer;
- *cost*: to itself or to the market: but this relies on knowing the costs of market participants;
- *bid cost*: as NEMDE does. There are no bids for non-market system services, although there may be prices in AEMO's contracts, which could be considered as similar to bids.

If AEMO's actions or objectives are unclear or opaque, this might lead to erratic or inexplicable interventions which are difficult for the market to predict, interpret and respond to.

3.4 POSSIBLE SOLUTIONS

3.4.1 Two-sided Pre-dispatch

As discussed in section 2.4.5, two-sided dispatch means that all market participants – and specifically retailers²⁸ - have to submit bids or forecasts into NEMDE, rather than just scheduled generators. Since the bidding process is part of PD rather than dispatch, this inevitably means that PD also becomes two-sided. Questions then arise as to how far ahead of real-time these demand-side bids must be submitted, and what corresponding good faith obligations apply. For example, if good faith PD bids must be submitted at the day-ahead stage, this effectively means that retailers must submit and maintain forecasts of demand and demand-response for the next 24 hours: eg at a 30-minute granularity, similar to scheduled generators currently. Thus, the market – rather than AEMO – would be responsible for forecasting load²⁹, with each MP responsible in relation to its own assets and customers.

Conversely, a good faith forecasting obligation only makes sense, and can only be effective, if there is also a two-sided *dispatch*, since PD bids are forecasts of dispatch bids. Furthermore, if AEMO – rather than the market – was still producing 5-minute-ahead dispatch forecasts as now, there would inevitably be inconsistency between PD bids and dispatch bids and so PD would be less effective.

Notwithstanding PD obligations, each trader already relies on forecasts of its own load and generation in making scheduling decisions, since these will affect the relative risk and profitability of different scheduling options. Traders might make use of AEMO forecasts for this purpose – adapting them as needed to reflect retail market share etc – or might do their own forecasting if this can be done more accurately. So, in a two-sided market, forecasts submitted to PD should, in aggregate, be no worse than the AEMO forecasts used today. In fact, they should be better, because bidding allows the forecast to be spot price dependent: eg demand response can be reflected by bidding a lower load quantity for a higher spot price. AEMO does not currently attempt to forecast price response, and would find it

²⁸ Stronger bidding obligations might also be placed on semi-scheduled generators. Currently these generators are responsible for bidding technical availability whereas AEMO forecasts *resource* availability. As part of a two-sided market design, full bidding responsibility might be placed on the generator.

²⁹ And potentially also resource availability for semi-scheduled generation

difficult to do this without having knowledge of retail contracts that provide for or encourage demand response.

This option only applies to energy. For market AS, AEMO would still decide how much needs to be procured in dispatch to maintain system security, and submit forecasts of these amounts into the PD process as now.

3.4.2 Changes to the Pre-dispatch engine

The current PD engine is set up to mimic NEMDE. However, this is not inevitable. Indeed, since PD operates ahead of real time and over an extended study period, many different scheduling engines are possible. A different PD engine might potentially address the issues around convergence or effectiveness of the PD process that were discussed in sections 3.3.1 and 3.3.2.

One option would be to retain the existing inputs (including bid structures) but change the optimization approach, so that the engine searches for the lowest cost solution over the study period, rather than separately optimizing for each timestep sequentially. This is similar to the idea of a multi-step dispatch discussed in section 2.4.4, and might similarly allow a better coordination of different plant to achieve high ramp rates or better scheduling of storage cycling, say. As discussed in that section, this extended optimization problem remains linear, so fast and reliable solving is guaranteed and shadow prices are produced that can be used as clearing prices.

A more radical option would be to change the bid structures. For example, 3-part bids might be submitted and the PD engine designed to make commitment decisions, rather than these being made by each trader and submitted to PD as a *fait accompli*. However, this is a more difficult problem to solve and requires making integer decisions. A mixed integer algorithm may not be so fast or reliable as the current PD engine, and the shadow prices that are generated are not necessarily clearing prices in the sense of ensuring that bid costs are always covered by spot revenue. This issue is seen in some ahead markets overseas which accommodate 3-part bidding and then have to provide “side-payments” on top of spot price revenue to ensure that bid costs are always covered by market revenue³⁰.

In terms of effectiveness, a more sophisticated PD engine should be able to find lower “cost” solutions – meaning *bid cost* – than the current PD engine. However, as discussed in section 3.2, the PD engine is only one element of the PD scheduling process, and the impact on scheduling effectiveness overall depends upon interactions between the PD engine and the individual traders. Consider how different PD engines might affect scheduling overall, in two areas: storage scheduling and ramping coordination.

Currently, a storage operator might bid its storage into PD in terms of:

- Price: eg fixed bids prices for charging and discharging, or:
- Quantity: sculpting its availability to reflect its preferred schedule

Either way, if the trader has “got it wrong” this will be revealed in the PD outcome. So, for example, if it has offered a price that is too low and the storage is fully discharged before the highest spot prices occur, it might then rebid at a higher spot price. If, on the other hand, it has inadvertently submitted a

³⁰ As discussed in the Appendix

quantity schedule that leads to its storage discharging over a period of relatively low spot prices, it will rebid to reschedule this discharge into a higher-price period. Other traders with storage assets will be similarly rebidding. As PD iterates, it is to be hoped that there is convergence towards a solution that satisfies everyone. But there is, of course, no guarantee on whether and when convergence will occur.

With a more sophisticated PD engine, each trader simply bids its storage characteristics (eg storage capacity) and a structured bid price (eg a minimum price *spread* that must be obtained between charging and discharging) and the PD engine then coordinates and optimizes these bids in a single PD run. Of course, there may still be some iteration as traders adjust their bids, but convergence might occur more quickly; within fewer PD iterations.

Or consider ramping coordination. In a PD run, a plant with ramp rate limitations might be dispatched to minimum generation over the mid-afternoon trough (eg due to negative prices) and then be unable to ramp up to full load in time for the evening peak when spot prices are \$10,000/MWh, say³¹. To address that currently, a trader might then rebid more of the output to negative price bands during the trough, so that it is not deloaded so far over this period. It is then able to reach maximum output for the peak. In doing this, the trader will be trading-off the cost of having a higher output over the negative price period against the benefit of higher output over the evening peak. Again, there will be some iteration as other traders respond to the impact that this rebid has on their profitability. PD will hopefully converge to an efficient solution where the ramping capability of all generators are coordinated so as to be able to follow the demand ramp.

Alternatively, a sophisticated PD engine that cooptimised across all timesteps could do this automatically without the need for a rebid. Ramp-limited generators could simply bid close to cost throughout the day, and the PD engine will automatically increase their output level during the trough, *apparently* out-of-merit, without the need to submit a negative bid. Again, this might allow PD to converge to an optimal solution with fewer iterations.

If scheduling were done by a single, monolithic, black-box scheduling engine – as would be the case in a vertically-integrated utility - one would expect that a more sophisticated algorithm would give more efficient results. But, as shown previously in figure 7, the PD engine is only one component of the PD mega-algorithm. Any change to that component will cause generators to make corresponding and consequential changes to their own scheduling and bidding tools. Plausibly, any increase in sophistication of the PD engine might allow the trader tools to be somewhat less sophisticated³², so essentially transferring functionality to the PD engine but leaving overall functionality and performance unchanged, illustrated in figure 13, below.

³¹ This scenario was discussed in the context of myopic dispatch in section 2.3.4.

³² Because, as discussed in the preceding paragraphs, traders could now submit naïve bids and leave it to the PD engine to schedule storage cycles and coordinate ramp rates.

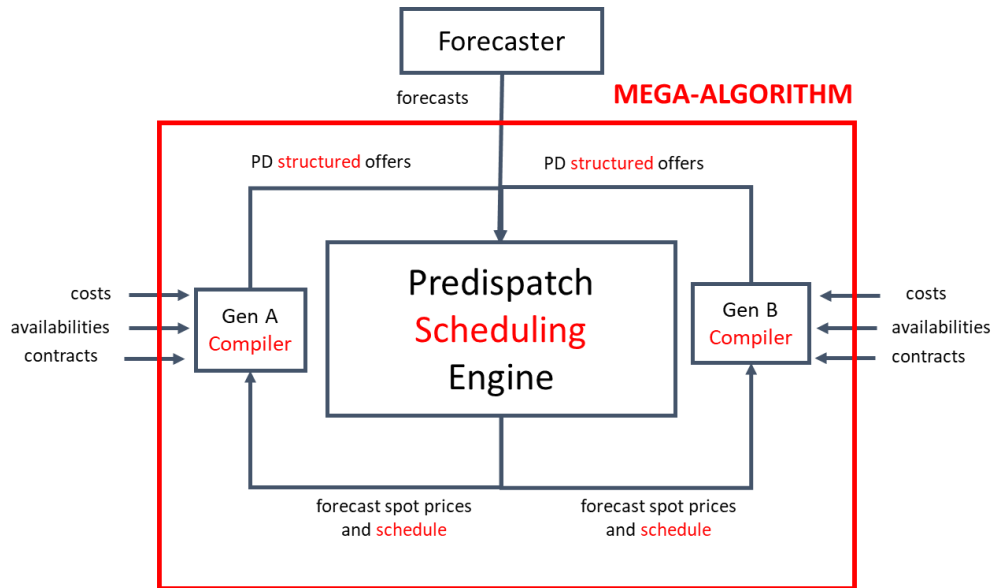


Figure 13: Architecture with a Pre-dispatch Scheduling Engine

On the other hand, possible factors such as a loss of transparency in the PD engine’s operation might create adverse interactions between bidding strategies and the PD engine, making overall performance *worse* than before. So, before any changes to the PD engine are implemented, the scheduling mega-algorithm should be performance tested. This is obviously problematic, given its decentralised nature. Possibly some form of “paper trials” could be undertaken³³.

3.4.3 Aligning the Pre-dispatch and Dispatch Engines

If a more sophisticated PD engine were introduced, this would create inconsistency between the PD and dispatch algorithms. PD bids that are structured differently to dispatch bids could obviously not simply be fed into the dispatch engine. And even if the structures were the same, a PD bid could not simply feed unaltered into dispatch as is the case currently.

For example, as described in the previous section, with a more sophisticated PD engine, the ramp rate limited generator can now bid at cost into PD – rather than negatively – over the trough and nevertheless be scheduled in PD at a higher output during the trough so it can reach its full output for the evening peak. But this generator will then have to bid negative (for part of its output) into dispatch during the trough to follow this PD schedule. This is illustrated in figure 14, below.

³³ Where PD processes are replicated in a test environment.

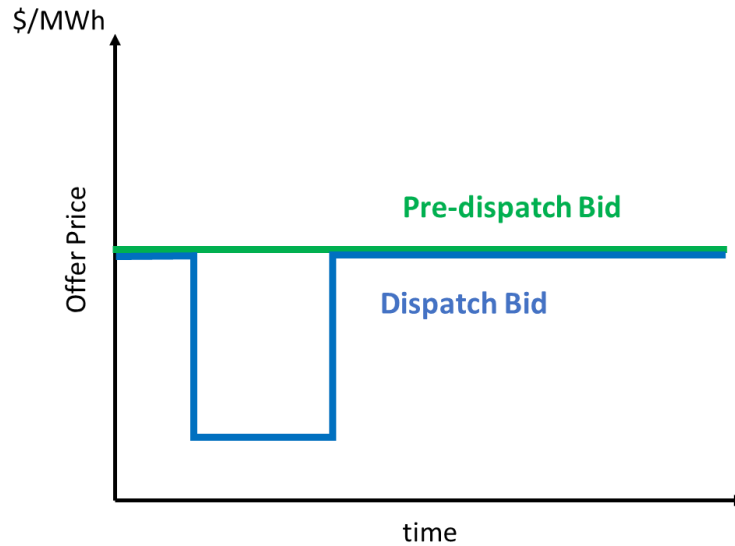


Figure 14: Bidding into incompatible Pre-dispatch and Dispatch Engines

If it bid into dispatch in the same way as PD, the myopic dispatch algorithm would de-load it over the trough and it would then still face the problem of being unable to ramp up to full output for the peak.

Given the need for different bids between PD and dispatch there is both a conceptual issue of how to define good faith bidding in PD and also a practical issue of how to maintain these two separate bid sets: one for PD and one for dispatch.

Thus, even if a more sophisticated PD engine might conceptually seem likely to lead to more efficient scheduling and coordination, this benefit might be lost due to the practical difficulties of aligning PD and dispatch.

Section 2.4.4 considered the possible solution of using a multi-step dispatch in order to address the problem of dispatch myopia. To avoid similar problems arising of having different bids in dispatch and PD, it would seem sensible to amend the PD engine in this scenario, so that it continues to mimic dispatch. Or, conversely, if there were merits in making the PD engine more sophisticated, it might be appropriate to change the dispatch algorithm at the same time, to ensure the two engines continued to be aligned. This raises the question as to whether the study period also needs to be aligned: eg if the PD engine optimizes over a 24-hour period, does NEMDE need to do this as well?

3.4.4 AEMO Good Faith Obligations

As noted above, although AEMO can be thought of as another “trader” in the PD process³⁴, it is not subject to the same good faith obligations as market traders. Currently, AEMO notifies the market when it identifies an insecurity in PD which means it *might* need to intervene. It then notifies the market later, if and when it actually *does* intervene. Between these two notices, there can be “radio silence”.

Presumably, once it identifies the potential insecurity, AEMO quickly investigates the potential interventions to address this problem, and selects the one that best meets its objectives.

Notwithstanding that AEMO does not need to intervene immediately, it has a clear *intention* to do so, unless something changes in the market to otherwise resolve the issue. But because it has no “good faith” obligation, AEMO is not obliged to inform the market of this intention. If AEMO had such an obligation, the market would be better informed about AEMO’s intentions and could respond accordingly. As traders do currently, AEMO could express this intention via PD bids³⁵

With this additional information, market participants will naturally coordinate themselves around this revised PD forecast through the iterated PD process. This is particularly desirable where the original AEMO direction (intention) has a cascade effect that prompts further insecurity and the need for more directions.

The existing system strength situation in SA provides a scenario of how this might occur. Suppose that system strength requirements imply that three synchronous units must be committed to ensure security, but only two are committed in PD. AEMO would notify the market of its intention to direct an additional unit to be committed. When this rebid is incorporated into a new PD run, it will have the effect of depressing the energy price and dispatch volumes available to synchronous plant (other things being equal)³⁶. Since the market can really only accommodate two synchronous units, AEMO’s intervention will likely lead to one of the two originally-committed units rebidding *out* of PD, leaving PD insecure once more: ie down to two units; one market, one directed. AEMO would then need to direct *another* unit, and then (after more PD iterations) *another*, until there are three directed units – and no market-committed units – in PD, restoring security³⁷. Clearly, there may need to be several PD iterations before PD converges to this point. If iterations only occur every 30 minutes, say, it will be vital that AEMO informs the market of its intentions well in advance.

Directions occur in the above example, but similar good faith obligations should apply to all AEMO interventions: eg also to scheduling AS contracts or RERT contracts.

³⁴ Albeit one with the system operator responsibilities of maintaining system security, rather than the usual trader objectives of maximising profitability.

³⁵ this might be by AEMO submitting these intentions to PD directly, or by AEMO informing the trader of the plant that it intends to schedule, with the trader then being obliged to rebid that plant into PD, in accordance with its good faith obligation

³⁶ note that, under the new Rules introduced late last year, there is no intervention “what if” pricing for such an intervention to neutralize the impact on prices of the additional committed unit

³⁷ Similarities to the children’s song “there were ten in the bed” are entirely coincidental.

3.4.5 Clarity of AEMO Scheduling Objectives

As noted in the previous section, coordination between AEMO and the market would be improved if AEMO's scheduling objectives were clarified, leading to greater transparency and predictability of AEMO's actions. AEMO has the twin objectives of, firstly, maintaining system security and reliability whilst, secondly, minimizing the costs of its interventions: both the direct costs (payments made under contracts or directions compensation) and the indirect costs imposed on affected market participants.

A distinction should be drawn here between *spot-priced services*³⁸ and other services. If the insecurity is caused by the shortage of a spot market service, the price of that service would be set at the market price cap³⁹, reflecting that scarcity. Those high prices should encourage greater supply of this service to be offered into PD, hopefully removing the supply gap and associated insecurity. Thus, AEMO should have the objective here of leaving intervention as *late* as possible, to give time for the market to respond and remove the need for AEMO intervention.

On the other hand, if the insecurity is due to a shortfall in *non-spot-priced services*, there will be no such price signal and so little to be gained by AEMO waiting. The market is *never* going to respond, because there is no price for it to respond to. Some rebidding might occur for other reasons, but that would only coincidentally resolve the security problem and is just as likely to make it worse. In this case, the objective should be to minimize the cost of intervention, and so to intervene early if this allows AEMO to reduce the cost of intervention.

3.4.6 Improving Pre-dispatch Convergence

Changes to the PD engine were considered above, with a view to improving the effectiveness and stability of PD. But simpler reforms might achieve this goal. Four possible changes are considered:

- More frequent PD runs
- Fewer restrictions on bids and rebids
- Multiple PD scenarios

These are explained in turn below.

If PD is slow to converge, in terms of the number of iterations, then an obvious solution would be to iterate PD more frequently. Of course, a full iteration of the PD process requires rebidding by traders, as well as a rerun of the PD engine itself, so there is no point in rerunning the PD engine before traders have had sufficient time to rebid. Manual rebids inevitably take considerable time, but the rebidding process could be substantially speeded up if these rebids are automated (using computer algorithms alone), with no or limited manual intervention and interactions. We are starting to see such *autobidders* being used by traders and this trend is likely to continue.

With autobidding, PD iterations could occur with high frequency. Only the solution speed of the PD engine and the autobidders is a limitation. PD might plausibly iterate every minute say. Note that this is

³⁸ services that are priced in dispatch, whether they are market AS on the LHS of NEMDE constraints or priced uncontrolled services on the RHS of NEMDE constraint

³⁹ it is assumed that market price caps would be set for newly-priced services just as, for example, market price caps are set for FCAS services currently

about helping the PD *mega-algorithm* to converge. It is not about tracking external changes (eg in weather forecasts) because of course these would not be updated as frequently as this. Autobidding combined with frequent PD iterations might give more opportunity for PD to converge, but such convergence is not guaranteed. Algorithmic bidding is common in stock markets, for example, and interactions between these bidders can lead to unexpected, and severe, trading outcomes, known as “flash crashes”. In stock markets, this simply causes adverse financial consequences for traders. In the electricity market, if it led to insecurity, the consequences could be much more severe.

Alternatively, convergence could be improved by reducing the inherent instability of PD outcomes. One thing that can occur often in PD and dispatch is that spot prices can jump – or fall – substantially in consecutive iterations: eg from \$1000 to \$10,000. Whilst this might reflect deliberate trading strategies, it is more likely simply to reflect existing restrictions on bid structures. Only 10 prices can be bid and these are not permitted to change over the PD period: only the output quantities that these refer to can vary. This restriction was imposed in the NEM design primarily because of limitations on computing power which are unlikely to remain in 2025, given the advances in IT. So it might be feasible to relax this rule: to allow 100 price bands, say; or to permit rebidding of bid prices in PD. Still subject to the good faith principle, of course⁴⁰.

PD bid quantities are currently specified for each 30-minute trading interval. With the move to 5MS, bids will be able to be specified on a 5-minute basis. PD could then be usefully run with a 5-minute timestep to reflect the impact of this 5-minute granularity⁴¹.

Currently, each PD run is based on a single, central forecast conditions expected in real-time. Whilst there are some *sensitivity runs* around this, these are fairly simplistic, in terms of fixed MW quantities above or below the central forecast. Furthermore, these sensitivities only show how PD outcomes would change if there were no consequential change in bids. But, realistically, as forecasts change, rebids *would* be prompted.

For example, if the wind were to drop around the time of the evening peak, this would be progressively anticipated (one would hope) through updated PD forecasts as real-time approached. Traders might respond to these forecast changes by rebidding to commit more plant or charging up additional storage in anticipation of the increased peak requirement. But the PD sensitivities used today would not be able to model or reveal any of these contingent activities.

This could be addressed by running PD with several sets of PD forecasts: for example, a 90% POE and 10% POE forecast to go with the existing central (50% POE) forecasts, as illustrated in figure 15.

⁴⁰ which, admittedly, might be more difficult to monitor and enforce with these more detailed bids

⁴¹ 5-minute dispatch *is* carried out currently, but only for one hour ahead and this is an informal process not specified or required under the Rules




Figure 15: Pre-dispatch Forecasts using Three Scenarios

Traders would then be required to submit *three sets* of PD bids, one relating to each forecast scenario. Good faith obligations would continue. So, for example, the 90% POE PD bids would reflect how the trader would expect and intend to bid should those 90% POE conditions eventuate. This is easy to define in principle but might be difficult to monitor and enforce in practice, since it is unlikely, by definition, that these conditions will eventuate. This multi-scenario PD process would substantially improve visibility for market participants and AEMO, helping them to manage these tail risks. For example, AEMO might decide that it needs to intervene if and when any *one* of the PD tail scenarios is insecure, to give it an 80% confidence (say) that security is going to be maintained⁴².

Each of these potential solutions might improve the tracking and convergence properties of the PD process and the multi-scenario approach could also improve visibility by providing some additional “peripheral vision” of unlikely but possible outcomes. On the other hand, they would all entail higher costs, particularly the trader costs of managing and maintaining complex bids and/or multiple bid sets. They might also expose bidders to inadvertent breaches of good faith obligations: both because the nature of “good faith” might itself become unclear when there are multiple scenarios, and because of the uncertain behaviour of a complex, autobidding ecosystem and the potential need for manual trader interventions to correct “flash crash” errors. So the costs and benefits of these options need to be carefully considered and compared.

⁴² This, of course, begs the question as to what AEMO’s decision criteria should be, which goes back to the earlier discussion on AEMO’s scheduling objectives. Should AEMO be 80% confident, or perhaps 99% confident, that dispatch will be secure and reliable? The PD scenarios might be chosen to reflect AEMO’s criteria

3.5 CONCLUSIONS

Considered in its entirety, the PD process is a sophisticated, organic scheduling process which is likely to be superior in its performance, robustness, transparency and adaptiveness to any “black box” centralized scheduling algorithm that a system operator could come up with. It is, perhaps, not always recognized as such because of the simplicity of the PD engine that lies at its heart. But the PD engine design is in fact powerful in that it mimics the dispatch algorithm and so largely eliminates the seams between PD and dispatch that would be inevitable if a more complex and sophisticated PD engine design were used. Complexity, instead, lies hidden in the trading systems of market participants, who are directly motivated to develop and fund the sophisticated processes needed to achieve their scheduling objectives.

The PD process is built on the foundation of dispatch, and a building is only as strong as its foundations. Changing the PD process cannot by itself fix flaws in dispatch: in particular, where there are missing markets in some essential system services, their provision will be ignored by the profit-maximising schedulers in the PD process. Either these gaps in dispatch design must be filled, or AEMO must expand its role of arranging for these essential services to be provided.

Like market participants, AEMO needs to schedule its delivery of these services. In this process, market participants and AEMO face similar problems of making decisions in the face of uncertainty and of coordinating effectively with other schedulers. So AEMO’s difficulties around “visibility” and “firmness” during the PD process are not unique to it, but are faced everyday by market schedulers too. So AEMO and generators may have a common interest in identifying any changes to the PD process that could improve its effectiveness and robustness. This should have been the starting point for the ESB’s design investigations. New ahead markets can, at best, complement the PD process and may, instead, compromise or undermine it.

4 AHEAD MARKETS

4.1 INTRODUCTION

There are no ahead markets in the NEM currently; no transactions taking place in the PD day-ahead timeframe. But, of course, this does not mean that nothing is happening. As discussed in the PD chapter, market participants are busy during the PD period preparing their plant to be deployed effectively into the spot market.

At the time of early deregulation in some overseas electricity markets, IT systems were not fast enough to run a dispatch-based spot market. Instead, a market was run ahead of time: typically day-ahead. Because forecasts of demand and supply are *reasonably* accurate in this timescale, the cleared schedules and prices in the day-ahead market were a reasonable approximation of the real-time conditions. Residual balancing of forecast errors in real-time were settled through more primitive market arrangements: eg pay at-cost or at-bid. These original day-ahead markets have now largely been complemented or replaced by spot markets based on real-time dispatch.

The timing of the original development of the NEM, and also the far-sightedness of its designers, allowed the NEM to skip this “immature” market design stage and so it has never needed or used an ahead market⁴³.

The early day-ahead overseas markets largely mimicked the pre-existing day-ahead processes of regulated utilities: of scheduling generation as well as forecasting its level of dispatch. That is to say, scheduling and market clearing occurred at the same time and within a common process. But, because the NEM has a spot market, there has been no need for market clearing during this scheduling stage. The PD process involves scheduling and coordination but *not* trading.

For those familiar with day-ahead market designs, there is a danger of seeing the NEM’s PD process as a kind of half-baked ahead market: PD has all the “bids”, schedules and clearing engines but somehow fails to take the final step of clearing these transactions and creating actual ahead trades. There is also a danger of seeing the day-ahead as the real market and the spot market as just a balancing exercise: which of course it was in those early overseas designs which began with no substantive spot market⁴⁴. But in the NEM, the spot market is the real thing, and all the more effective because it reflects the most up-to-date information on supply and demand conditions.

Given this existing NEM design, it is hard to see ahead markets placing any more than an ancillary role in the NEM design; to perhaps provide some opportunities to hedge spot price risks in the way that forward contract markets do today. Yet the ESB seems to be looking for ahead market designs that can play a much more substantive role, akin to those early day-ahead market designs.

⁴³ a financial day-ahead market was initially contemplated for the NEM design but this idea was dropped prior to NEM implementation.

⁴⁴ and, in some cases, have still not introduced one.

4.2 CONCEPTS

4.2.1 Introduction

There is a variety of possible ahead market designs that might be able to be introduced into the NEM. Rather than try to consider all the possibilities, this section attempts to identify the main design choices. Because the current NEM design has no such markets, the descriptions are necessarily quite abstract and conceptual, without even the example of a familiar market design to illustrate them. The closest cousins are the current forward markets in electricity derivatives⁴⁵.

4.2.2 Ahead Market

Generically, an *ahead market* is a platform for transactions that take place ahead of physical delivery. As already discussed, all physical delivery occurs in the dispatch market, so ahead markets would trade in these same spot market products: ie energy and market AS. For example, energy for delivery on 17-may-2020 might be traded on 16-may-2020 on a day-ahead market.

Whilst trading might potentially occur months or years ahead of real-time, for the purposes of this report, given the focus of the ESB project, only markets transacting in PD timescales (day-ahead and within-day) will be considered.

Furthermore, since there is nothing to prevent such transactions occurring currently, the focus will be narrowed further to consider only platforms which are provided and operated by AEMO, since AEMO is the ahead-market operator contemplated by the ESB options paper.

4.2.3 Financial Settlement

Spot market settlement is currently *gross*, with the entirety of metered flows – of generation or load – being settled in the spot market. If the spot market settlement design were to remain unchanged, then the settlement of the ahead contract must be based on the *difference* between the ahead and spot market prices, to prevent the seller being paid *twice* (and the buyer paying twice) at the spot price. In forward markets currently, this is done by using financial derivatives, with contract settlement based on the *difference* between the spot price and ahead price. A similar structure would be used for an ahead contract. So, in effect, an ahead product would simply be the familiar *swap* traded in forward markets today.

⁴⁵ Although there are some other existing or proposed arrangements in the NEM that have some superficial similarity to the ahead markets proposed by the ESB, in that they – unlike forward markets – involve market institutions. Firstly, the Settlement Residue Auction is a forward auction run by AEMO to sell off the inter-regional settlement residue. However, this takes place months ahead of real-time, rather than the day-ahead period. Secondly, some proposals under the COGATI MDI incorporate similar AEMO-managed auctions of financial transmission rights. These are, similarly, months-ahead rather than day-ahead markets. Finally, the Retailer Reliability Obligation (RRO) places obligations on retailers to maintain certain forward positions under certain circumstances. However, this is a regulatory obligation around *existing* forward markets rather than a new market and, again, does not operate in the ahead timeframe.

Alternatively, spot market settlement could be changed so that only the *residual* between the ahead traded quantity and the metered quantity is traded in the spot market. In this case, the ahead settlement should be at the full ahead price.

In fact, these two different settlement approaches give financially identical outcomes as an example will show. Assume that the ahead and spot transactions, for a common product, are:

- *Ahead*: 100MW at \$80/MWh
- *Spot*: 110MW at \$50/MWh

The settlement amounts under the two approaches are shown in table 2 below:

<i>Settlement Stage</i>	<i>Derivative ahead product and gross spot market settlement</i>	<i>Gross ahead product and residual spot settlement</i>
<i>Ahead market settlement</i>	$100 \times (80 - 50) = \$3000$	$100 \times 80 = \$8000$
<i>Spot market settlement</i>	$110 \times 50 = \$5500$	$(110 - 100) \times 50 = \500
Total	\$8500	\$8500

Table 2 settlement of ahead and spot markets

For the purposes of discussion, it will be assumed – without loss of generality – that spot settlement is unchanged and ahead products are simply the swap derivatives that we are already familiar with.

4.2.4 Physical Obligations

By itself, an ahead market sale does not oblige the seller to physically generate at the quantity sold ahead. It just changes the financial outcome of doing so or not doing so. This is no different to current forward markets. This would still be the case even if the ahead contract references a particular generating unit, say. An ahead sale of “100MW from unit X” would give the same financial outcome whether unit X actually generated the 100MW in the spot market or a different unit Y (owned by the same firm) generated 100MW instead⁴⁶. This is because electricity, and all spot market products, are *fungible* in the spot market: all units produce the same product and are financially interchangeable. And, of course, it is why physical units are never referred to currently in forward contracts.

This is different to how the spot market operates. A unit that sells into the dispatch auction must then *physically operate* at that level. There is not the option of running a different unit at that level instead. This is why dispatch trades are *physical*. Ahead trades, as defined so far, are purely financial.

So for an ahead trade to be *physical*, there must be some *additional* physical obligations placed on the seller; to either operate at the ahead level in dispatch or face some additional non-conformance

⁴⁶ assuming that both units sells the same product in the spot market: ie are located in the same NEM region

penalties, similar to those that apply to dispatch currently. So, a physical ahead sale of 100MW by unit X would mean:

- A 100MW sale of a financial swap; *and*
- An obligation on unit X to operate at 100MW in dispatch or face non-conformance penalties.

If we have the former without the latter, it is simply a financial derivative. If we have the latter without the former, it is equivalent to a *direction* by AEMO under the current design. So a physical ahead transaction is analogous to a combination of a derivative trade and a direction.

4.2.5 Ahead Location

Since a financial ahead product is like a financial swap against the spot price, the only relevant locational information is to establish *which* spot price is being referenced. So, under the current NEM regional design, there would simply be a reference to the region⁴⁷. If the NEM were to become a nodal market, the node would need to be specified. So, whilst “100MW from unit X” is not meaningful, “100MW from region A” is and would be the form of bids or offers in such an ahead market⁴⁸

On the other hand, since the physical ahead product references physical delivery, this would need to specify a particular meter, or group of meters, that would measure the physical flow and compare it to the ahead amount. So, a generator selling ahead could reference the particular unit, or possibly a set of units, that it is selling from. A retailer buying ahead might reference a customer meter or group of customer meters, say.

Non-physical traders, that don't own generating plant or serve customers, would therefore not be able to participate in a physical ahead market, analogous to how they are unable to participate in the (physical) dispatch market. In contrast, anybody could potentially participate in a financial ahead market.

4.2.6 Clearing Times

The spot market must clear shortly ahead of real-time, to ensure that the information it is based on (including 5-minute-ahead demand forecasts) is as accurate as possible and security is ensured⁴⁹. So it must take the form of an auction, where all offers and bids are cleared in one go.

A similar auction-based approach *could* be chosen for an ahead market: eg an auction takes place at 12pm each day, covering products relating to the following calendar day. However, since time is no longer of the essence, many alternative timings and mechanisms are possible. There could be repeated auctions: for example, auctions could take place at 6am, 12pm and 6pm, all relating to the following calendar day. Or there could be a sliding window, where auctions take place each half-hour covering

⁴⁷ or perhaps, more specifically, to the regional reference node

⁴⁸ Ahead markets in the US operate where there are nodal spot markets, so bids in these markets are in the form: 100MW from *node A*

⁴⁹ of course, there will still be some level of forecasting errors and these are handled in the NEM through the actions of FCAS providers. So, generally, the further ahead of real-time that the spot market clears, the greater reliance on FCAS or similar services

the following 24 hours. For each auction there needs to be a *gate closure* a short time before the auction takes place, with all bids and offers submitted prior to that closure being cleared in the auction.

Alternatively, trading could be *continuous*, similar to a stock market. Bids and offers could be submitted at any time and a transaction occur whenever a bid and offer (or sets of bids and offer) can be matched: for example a bid for 100MW of electricity for the period 6:00-6:30 the next day at \$50 could be cleared with an offer for 30MW of electricity over the same period at \$49. The transaction might clear 30MW at \$49.50, say. This type of platform is commonly called a *bulletin board*.

Under repeated auctions, or continuous trading, participants could accumulate financial and physical obligations. So, a generator might sell 100MW ahead at the 6am auction, say, and then sell a further 50MW ahead at the 12pm auction⁵⁰. It would therefore have sold 150MW of swaps in total and any physical obligation would reference that 150MW aggregate. Of course, a generator might sell in the first auction and then buy back some or all of this position in subsequent auctions. Indeed, if there are no physical obligations, a generator might be a buyer even in a single ahead auction.

4.2.7 Gross or Net Market

As discussed in section 2.2.3, the NEM spot market is considered a *gross* market. This is because it is a physical market and generators must conform to the dispatch instructions which come out of the cleared dispatch market. Strictly speaking, participation is not mandatory, but if you have a large power station, there is no way to produce and sell services other than by actively submitting offers into the dispatch market.

As discussed above, there is not 100% active participation in the spot market, given that neither retailers nor small generators are required to submit bids and offers, although they still must trade at the spot market price. However, the active participation of all large generators is sufficient⁵¹ to give AEMO the visibility and control it needs to manage and maintain system security by using NEMDE to clear the auction. So that is a useful and meaningful distinction between a gross market and the alternative of a *net* market where such visibility and control is not afforded to AEMO.

Assuming that the spot market is not to be dismantled – and this has never been suggested – market participants will always have the option of trading through this market rather than an ahead market. If the ahead market is financial, participants could have the additional option of trading derivatives over a different day-ahead platform: eg the ASX⁵².

So an ahead market can only be made gross if ahead participation is commercially advantageous relative to the spot market or derivatives markets. This might be, for example, by:

- Providing priority access to congested transmission capacity in the ahead market; or

⁵⁰ These would likely be at different ahead prices

⁵¹ at least for the time being. As noted previously, as load becomes increasingly responsive and as smaller generation such as rooftop PV grows in market share, there may be a need to change to a *two-sided* market with much wider active participation mandated

⁵² If it decided to set one up.

- Levying additional fees on the residual spot market volumes: eg to recover ancillary services costs

Alternatively, the ahead market might be naturally advantageous: eg if prices are less volatile than in the spot market.

4.2.8 Demand Side

As with the dispatch and PD processes, a question arises as to who would be responsible for submitting demand-side bids into an ahead market. For energy, it would be natural for retailers to bid, giving them direct responsibility for any bids that are cleared. If, on the other hand, AEMO were to bid on behalf of retailers, several difficulties emerge:

- AEMO would need to decide on the level of participation: eg would it bid 100% or 50% of the forecast demand into the ahead market, say.
- An algorithm would then be needed to allocate any cleared bids (and their associated settlement payments) between retailers.
- AEMO would be unable to submit price-dependent bids because it has no knowledge of the demand response that end customers might provide⁵³. It would therefore need to price the entire bid at the market price cap, potentially adding substantially to retailer costs if the forecast is too high and the ahead price is set at the MPC, or if the ahead price is constantly set higher than the spot price

In the light of these difficulties, it would seem essential that retailers were responsible for energy bids in ahead markets.

On the other hand, it would be natural for AEMO to retain responsibility for bidding for market AS, since retailers do not have good knowledge of how much is required for security and what share of this they are required to fund. However, similar issues arise as to how much AEMO should bid for (and how much should instead be left to purchase in the spot market) and at what price. As with existing AEMO interventions, AEMO would need to have twin objectives of ensuring system security whilst obtaining value-for-money.

4.2.9 Financial Market Clearing Engine

The design of the clearing engine depends on whether the ahead market is physical or financial. This section considers a financial ahead market and the next section a physical one.

As discussed above, a financial ahead product only references a region (or node). Furthermore, because it does not relate to a particular generating unit or customer, there is no requirement to have operational constraints on generators like there are in a physical market⁵⁴.

⁵³ AEMO has a similar issue in “bidding” into the dispatch market with its five-minute-ahead demand forecasts

⁵⁴ to again draw an analogy with forward markets, these trade in simple generic products which reference only the region and do not attempt to mimic generator operating constraints

So the clearing design could be anything from quite simple to highly complex. The simplest design would simply match bids and offers of the same product: ie the same service, trading interval and region. A more complex structure could trade strip products, say, covering multiple consecutive trading intervals. Or the clearing engine could co-optimize different products: a generator could offer to sell energy or AS in alternative feasible combinations, similar to how energy and FCAS are optimized in dispatch. The generator could even put in structured bids (eg with energy constraints) whose constraints or costs the clearing engine would need to reflect in its clearing logic. Generally, any clearing algorithm that could be used to reflect actual physical generating unit constraints (as discussed for PD) could potentially be used to clear a financial ahead market.

Since the bids and offers are regional, a question arises as to whether some clearing *between* regions could be permitted in the ahead market, reflecting the ability for interconnectors to transfer power between regions. However, the effective level of interconnector capacity would be unclear at the ahead stage, since it is also affected by intra-regional constraints which would not be able to be modelled in the financial ahead market clearing engine⁵⁵.

With the more complex clearing engines, a financial ahead market begins to resemble the existing PD process. However, there are several key differences which should not be lost sight of:

- The market is *net*: unlike with PD there is no obligation to participate
- There is no good faith obligation: generators can bid however they want; their bids are not forecasts of how they expect to bid into the spot market
- There is no representation of intra-regional transmission constraints: since all bids relate to RRP, unless nodal pricing is introduced
- And, of course, no transactions occur in the PD process: it is “for information only”.

For those who are looking to the US as examples of ahead markets⁵⁶, the significance of the NEM’s regional pricing model cannot be overemphasized. If COGATI proposals to introduce nodal pricing were agreed, that would at least address this major flaw in the ESB ahead market proposals. Of course, that is no reason to introduce nodal pricing; or to introduce ahead markets, either. But if one is looking to transplant overseas market designs to the NEM, it is important that this is done in a systematic and comprehensive way. Market designs operate as a coherent whole, and picking and choosing individual elements to introduce can lead to a flawed design. Hopefully, this will be recognized and addressed in the ESB work program, when the elements developed separately by each MDI are brought together into complete design packages.

⁵⁵ there is also the question of who is actually *offering* this inter-regional capacity into the ahead market. Effectively, some of the inter-regional settlement residue is being sold, analogous to how this is sold currently through settlement residue auctions. So, in the sense that this revenue belongs to TNSPs or their customers, they should logically have a say in how this would be offered.

⁵⁶ And some of the terminology and architectures being proposed by the ESB suggest that it is drawing heavily on these examples.

In summary, AEMO would not be able to rely on ahead market outcomes as indicating that dispatch is likely to be secure (or insecure) as it does currently with PD. In fact, it is possible that it would get no better insight into system security from seeing ahead market outcomes as they would if they looked at, say, forward market transactions on the ASX happening currently.

4.2.10 Physical Market Clearing Engine

If the ahead products were physical, bids and offers would need to reference the associated load or generation meters, respectively. Each seller would need to ensure that it could physically supply the cleared amount from the specified generating unit and so would need to incorporate any operational constraints implicitly or explicitly into its offer.

In a physical market, since bids and offers will be at different physical locations, there will be transmission flows associated with any clearing. If these flows are inconsistent with system security – eg because it would mean that some transmission lines are overloaded - it might not be possible for all cleared parties to fulfil their delivery obligations, even if nothing changes between ahead and spot. Thus it may be appropriate to include security constraints within the ahead market clearing engine to avoid this.

Alternatively, the ahead market could be *unconstrained* – where clearing is *not* constrained by these transmission constraints - with provisions that permit generators to deviate from their ahead amounts if dispatched to do so by AEMO in order to manage congestion⁵⁷.

So the options for the clearing engine design will be similar to those for the PD engine, with the additional option of including or not including system security constraints, as discussed.

The physical penalties for non-conformance to ahead schedules will affect how generators participate in this market, and could introduce biases. For example, these penalties might be *asymmetric*, with a generator penalized for operating *lower* than the ahead schedule, but not for operating *above* it. Under such a regime, a generator would probably be conservative in how much it sells ahead. A windfarm operator, say, might sell an amount ahead which it forecasts a 90% probability of exceeding in dispatch. Even a firm generator might decide to offer only some of its expected dispatch output into the ahead market, to give it some *optionality* to change its dispatch intentions should forecast conditions change.

4.2.11 Impact on Spot Market

The impact on spot market design and operation will depend upon whether the ahead market is physical or financial. If it is financial, the products traded are simply derivatives, similar to those that are already traded in forward markets. Since the latter are already accommodated by the existing spot market design, there would appear to be no need to alter the spot market were a financial ahead market introduced.

Of course, whilst the design is unchanged, trading behaviour may well change. Derivatives change a market participants exposure to the spot price: for example, a 100MW generator that has sold 70MW

⁵⁷ Overseas, ahead markets in the US are typically constrained, whereas in Europe they are typically unconstrained. This is discussed further in the Appendix.

ahead (or forward) is exposed to only 30MW at the spot price if it generates at full output, and is exposed to being *short* 70MW if it doesn't generate. This change will affect the relative profitability and risks of different trading or bidding strategies and so a trader will submit bids that reflect the ahead position.

This is commonly referred to as *defending* the forward position, but this should not be taken to imply that the generator that has sold 70MW ahead *must* ensure that it delivers at least 70MW in the spot market. However, it will typically mean that it will deliver closer to the forward amount than would be the case if it bids its generation at cost. This is illustrated in figure 16.




Figure 16: Offers to "defend" a forward position

Relative to the output it would provide if bidding at cost, the generator in figure 16 bid so its output is closer to its forward position. This has the effect of pulling down the price somewhat when the generator is *short* (forward position exceeds output) and raising the price somewhat when the generator is *long* (output exceeds forward position). The extent to which it does this depends upon several factors: in particular, the market share of the generator.

On the other hand, if the ahead market is physical, each generator must deliver in accordance with its ahead position or face possible penalties for non-conformance. There are many reasons for variations in energy dispatch: generator technical issues, congestion management, balancing of demand or supply variations etc. It would be complex to define which of these were permissible (and so did not attract non-conformance penalties) and which did not. A generator would endeavour to bid in the energy market in a way that reflected those penalties.

On the other hand, dispatch of system services might be less variable. For example, if inertia became a market AS, AEMO could buy inertia from a synchronous generator in the ahead market. That generator would simply need to ensure that the unit specified in the ahead contract was on-line as required: variations in its energy output would not matter.

4.2.12 Impact on Forward Markets

The potential impact of ahead markets on the existing forward markets will similarly depend upon whether the ahead product is financial or physical.

If ahead trading is financial, and so voluntary, it may well just complement the existing forward markets. Generators and retailers may already have taken satisfactory – from a risk management perspective – positions in the forward market and the ahead market might create an opportunity to fine-tune these positions in the light of better knowledge (eg around plant availability and weather forecasts) than was available at the time of the forward trades. This is discussed further in section 4.3.2, below.

US day-ahead markets are quite liquid, leading to forward derivative contracts typically referencing the day-ahead price rather than the spot price. In this situation, generators will then offer into the day-ahead market to “defend” these forward positions, in an analogous way to how NEM generators bid into the spot market to defend their forward (spot-price-referencing) positions. US markets have arrived at this point largely for historical reasons: day-ahead markets have always been the main traded market since wholesale energy markets were first established. Given that the NEM is starting from a different point, it seems unlikely that it would migrate to a similar position, although this is not impossible.

If ahead trading is physical, market participation essentially becomes mandatory. So, for example, a generator intending to generate 500MW may have to sell 500MW in the ahead market to avoid non-conformance penalties. It cannot then also have a 500MW position in the (spot-referencing) forward market, otherwise it will have effectively sold forward (through the combination of the forward and ahead markets) *double* its output. So, in this case, forward contracts would *have* to reference the ahead price rather than the spot price. This might imply the need for a transition process to amend or renegotiate existing forward contracts to reference the ahead price.

4.2.13 Transmission Access

In the current market, generators may be constrained-off⁵⁸ by AEMO within a region in order to maintain system security. There is no compensation for this. Instead, the pain is shared through a *tie-breaker* rationing mechanism in NEMDE⁵⁹. If the ahead market is physical and used a constrained clearing approach, a similar issue could arise there and tie-breaking might similarly be adopted.

A question arises as to how transmission capacity would be shared between a physical ahead market and the spot market. For example, capacity might be fully allocated in the ahead market, meaning that there is nothing to be allocated to spot market participants. This would obviously encourage participation in the ahead market.

⁵⁸ dispatched lower than a clearing of their offer at the spot price would imply

⁵⁹ As will be well known, generators can attempt to avoid being constrained off by rebidding to a lower price, potentially down to the market price floor of -\$1000/MWh. When several generators behind a radial constraint do this, and so all bid at the same price, NEMDE is financially indifferent to which it constrains off, and uses an administrative algorithm to decide, constraining each generator in proportion to its availability.

4.2.14 Conclusions

There is nothing wrong with having ahead markets, *per se*. Many overseas markets use them. But it is important to recognize that ahead markets and spot markets form a package, and this package must be consistent and coherent overall. The NEM has been designed without ahead markets and without the need for an ahead market. Overseas markets that do use ahead markets either have quite different spot market designs or, effectively, no spot market at all.

So, there are fundamentally two ways to go with the NEM design. Leave the spot market essentially unchanged and add in some ahead-market features to complement this. These ahead markets will then need to be voluntary and financial: essentially a form of the short-term forward markets contemplated in a recent Rule change proposal.

The other way is to introduce ahead markets which largely replace or displace the spot market. These new ahead markets will be gross and physical. But this is essentially to travel back in time to adopt earlier (overseas) market designs. This seems a backward rather than progressive step. It would take some fundamental failure in the current spot market design to even contemplate such a radical redesign.

This is not to say that the ESB is explicitly or consciously proposing such radical change. Rather, its philosophy seems to be to “nudge” PD into an ahead market. To take the architecture that already exists – PD bids and offers – and take the extra step of clearing the implied transactions. But this option does not exist; it is a mirage. PD “bids” are not bids at all but forecasts of spot market participation. The ESB has been misled by apparent similarities between PD and day-ahead markets that are no more than superficial. It is hoped that this chapter, by carefully and systematically combing through ahead markets concepts and options, has demonstrated this fundamental impossibility. So that when we come to look at the ESB’s options for attempting to “square the circle”, in the next chapter, it should be no surprise that the ESB fails to achieve the impossible.

4.3 ISSUES

4.3.1 Introduction

This section discusses the issues that are discussed in the ESB paper and which the ESB considers could be addressed by the introduction of ahead markets. Some of these issues have already been considered in section 3.3 in the context of PD-based solutions.

4.3.2 Firmness of pre-dispatch bids

As already discussed in section 3.3.1, a concern expressed by ESB in its paper is that PD bids are “non-firm”, meaning that a rebid can unexpectedly lead to forecast insecurity, which AEMO might then not have time to address through intervention. So, could the introduction of some form of ahead market improve firmness?

It is assumed that PD will continue to exist alongside any ahead market, with the existing good faith obligations. Remember PD “bids” are not actually bids, but forecasts, so the issue is really whether these forecasts would become more *accurate* if an ahead market would be introduced.

It is also important to distinguish between the bids that are submitted to the ahead market and the PD “bids”. The former are firm, but only for the short period between the gate closure⁶⁰ to the auction clearing, after which they cease to exist, as illustrated in figure 17, below. So the firmness of these bids is irrelevant to AEMO’s issue of having visibility and predictability of system security in dispatch. The question is, rather, whether ahead market outcomes will cause the resulting PD bids to be firmer.




Figure 17: Firmness of Ahead Market Bids

Consider first a physical ahead market. Since non-delivery penalties can be high, market participants will ensure that they are able to deliver in accordance with the ahead market clearing. They will schedule their plant, and bid into PD, to achieve this. The benefit of this to AEMO, in terms of visibility, will depend upon the level of participation in the ahead market. If it is net, plant that has no ahead market position will have similar firmness as now⁶¹.

As noted above, market participants will only sell into the ahead market if it advantageous relative to the spot market. It would be a clear *disadvantage* of the ahead market if it locks the generator into its ahead position, with no opportunity to amend its scheduling closer to real-time if forecasts change. Unfortunately for AEMO, this means that it is those marginal plant that are most likely to want to rebid their availability in PD that would see the most disadvantage from participating in the ahead market. So, a physical ahead market might not improve visibility and firmness unless there are regulatory penalties placed on spot market participation.

Now consider instead a financial ahead market. Obviously, there are already forward contracts and market participants bid into PD in the light of these. So, the question of whether a financial ahead

⁶⁰ bids could not be required to be firm *prior* to gate closure, because participants could then simply leave it to the last second to bid, in order to effectively avoid this obligation

⁶¹ unless, of course, such plant is precluded from participating in the spot market. That is conceptually possible, but essentially involves dismantling the spot market, which I don’t think is being contemplated

market would improve visibility amounts to considering how forward positions might change if the ahead market were introduced and how this would improve PD firmness.

If the ahead market is sufficiently liquid, it would give market participants opportunity to fine tune their forward positions. For example, forward positions currently need to, in effect, be set months ahead of real-time, based on typical weather conditions etc. At the ahead stage, weather forecasts will be available, so market participants might take the opportunity to, say, increase cover when the weather forecast is extreme and demand is expected to be high; and, correspondingly, reduce cover when the weather forecast is mild.

These changes in cover might then flow into PD, through generators scheduling more or less plant to defend these higher or lower positions, respectively. If one equates “higher cover” with “more firmness” (itself arguable) this leads to *less* firmness than currently, under low demand conditions: which, increasingly, are becoming just as problematic for system security as high demand conditions.

The other distinction between forward and ahead markets is that AEMO is the operator of the latter, so has visibility of trades taking place. However, a financial market is *net* by definition. In fact, it is quite plausible that retailers might be selling, or generators buying, in this market. So, whilst the clearing *prices* might give AEMO some indication of the conditions that the market expects the following day, the cleared *quantities* are unlikely to provide any useful information.

Indeed, if AEMO’s objective were simply to have more visibility of market participants’ forward positions, this could be done more easily through introducing rules that require this disclosure (on a confidential basis) rather than AEMO going to the trouble of setting up its own market⁶².

4.3.3 Firmness of AS PD Bids

Currently, there is limited forward trading in market AS. The markets are fragmented and relatively low value, and market participants exposures to the spot market are uncertain, given complexities in how AS costs are allocated. Nevertheless, market participants still have positions to defend. Generators are charged for some AS and gentailers would also bear a share of the costs of the services levied on load. Generators will factor this exposure into their AS bidding strategies and may seek to defend these exposures, similar to bidding in the energy market.

As noted above, AEMO might act as proxy bidder for the demand side in an ahead AS market, seeking to procure AS ahead so as to ensure system security. But it is only a proxy: the cost of these purchases would be passed on to market participants in some way, similar to how the cost of spot AS purchases are passed on. So, a gentailer selling AS to AEMO in the ahead market would, in part, be selling to itself. It is unclear how this would affect ahead market participation and the AS position that a generator or gentailer would have going into dispatch. So, again, it is uncertain whether an ahead market would make AS PD bids firmer than they are currently.

⁶² This option is discussed in section 4.4.2.

4.3.4 Scheduling and Coordination Effectiveness

As discussed in the pre-dispatch chapter, scheduling and coordination is undertaken in the PD process by a combination of the individual traders trying to optimize their own position and the PD engine then bringing these schedules together to see what this means for forecast prices. The process iterates, allowing these individual schedules to progressively coordinate towards an optimal schedule⁶³ overall. This iteration also allows new information – eg updated forecasts – to be introduced and processed, right up until real-time. Whilst the current PD process – with a PD clearing engine that mimics dispatch – has worked well to date, some possible refinements were explored that might be needed to maintain its fitness through the energy transition.

The possible role of an ahead market in scheduling and coordination depends, once again, on whether it is physical or financial. As discussed above, a financial ahead market is going to be voluntary and net, so if there is any scheduling and coordination happening, it only relates to a part of the market. Furthermore, because there are no security constraints included in the clearing process, it cannot represent or reflect the complexities of dispatch in the way that PD does.

On the other hand, a physical ahead market might be gross and could incorporate security constraints, depending upon the design details. So, potentially, the ahead market outcomes could be reasonably reflective of dispatch conditions and constraints.

However, the iterations that are key to coordination in the PD process do not exist, or are substantially reduced, in an ahead market. There can be no coordination of bids prior to an auction, and once cleared the bids no longer exist. There may be subsequent auctions and so some sort of coordination could perhaps emerge in this sequence. However, because the auction bids are real (not just forecast bids) the dynamics will be quite different, and it is not clear whether or why coordination would be improved.

To consider this most starkly, suppose hypothetically that the PD process was removed, so that scheduling and coordination process relied entirely on ahead markets. AEMO and the market would see a snapshot of forecast conditions each time the ahead auction cleared, but then would be in the dark about any subsequent changes until the next clearing. Furthermore, since each sequential bid is simply intended for the next ahead auction, there can be no concept of good faith (or, at least, the concept would need to be substantially amended) and there might be little continuity between progressive bids and outcomes.

Thus, it is hard to see how ahead markets would assist with the existing PD-facilitated scheduling and coordination process.

⁶³ or, at least, a Nash equilibrium

4.3.5 Coordination between the market and AEMO is poor

As discussed in section 3.2.8, when it intervenes in the market (through scheduling contracts or directions), AEMO is essentially just another scheduler/trader in the market. Coordination is best served by ensuring that AEMO participates fully in PD, in accordance with similar good faith obligations as market participants.

Since the scheduling and coordination requirements are essentially the same, it again seems unlikely that an ahead market, of any form, would improve the existing process.

4.3.6 Risks in the Spot market

The ESB paper notes that an ahead market might be a facility for a generator to hedge some risks associated with scheduling decisions that rely on PD forecasts that may turn out to be inaccurate. For example, a generator might commit an additional marginal unit on the basis that its costs would be covered by the PD prices, but may end up losing money if spot prices turn out lower. Similar risks might exist for a retailer calling on a customer to manage its demand on the basis of high PD prices.

An ahead market could hedge such risks. A generator could offer the output from the marginal unit into the ahead market at a price that covered its costs. If this offer clears, it can commit the unit and be confident that its costs will be covered; if the spot price is then lower than expected, settlement of the ahead contract would make up the shortfall. Of course, this requires a willing “counterparty” to the trade: perhaps a retailer bidding into the ahead market to hedge its higher than normal customer demand.

This is plausible. However, the magnitude of the risks that are being hedged seem likely to be quite modest in both relative and absolute terms. Whilst spot prices are volatile, much of this volatility is due to variations in factors (eg weather) that are already known with some degree of certainty at the day-ahead stage and reflected in the ahead price. So variations between ahead and spot prices will be relatively small, reflecting only the residual uncertainty at the ahead stage. Furthermore, the exposure is only on a small part of the generator’s overall portfolio (ie the plant which is so marginal that a trader is uncertain whether to commit it or not) and for a short period of time (eg a few hours each day). For a generator, such risk is likely to be in the noise level. Similarly, for a retailer looking to hedge the risks associated with calling demand response, the risks will be relatively small.

There is also an implicit assumption here that ahead trades can be undertaken at close to *fair value*: that is to say, the seller or buyer is not giving up too much expected profit for the sake of reducing its risk. In a liquid market (eg involving the participation of non-physical speculators), trading at fair value is plausible, due to the opportunity to arbitrage away any substantial and consistent differences between ahead price and fair value. However, if the market is thin, any significant offer is liable to pull the market price below fair value and the hedging benefits of the trade are more than offset by the cost of selling at a discount. Liquidity is, unfortunately, self-fulfilling. If traders don’t expect to get value in the market, there will be less trading and liquidity and so value will fall further.

Therefore, it seems unlikely that a financial ahead market could offer significant hedging opportunities and value for market participants⁶⁴. A physical market would be even worse as a hedging mechanism, because an ahead trade would introduce new risks of being penalized for not conforming to the ahead schedule in dispatch.

4.3.7 Market Design for a New System Service

As previously discussed, there are some missing markets in system services currently that the ESB design project is considering how to address. Alternative approaches of new market AS (ie traded in the spot market) or non-market AS (traded through a contract market) have been considered. A third possibility would be to trade in a physical ahead market *only*: ie with no associated spot market trading at all. Essentially, this is a particular form of contract market, where the tendering process for the contracts takes place at the ahead stage through some form of auction. A usual non-market AS contract would typically provide for AEMO to be able to call upon the service to be delivered in an ahead timeframe: whether one day or one hour before real-time, say. With the ahead-market AS, AEMO would know how much it needed to procure and so the obligation for physical delivery would be implied. This is illustrated in Table 3, below.

Ancillary Service Category	Forward Period (months ahead)	Ahead Period	Real-time
<i>Market Ancillary Service</i>	n/a	AEMO determines requirement and formulates NEMDE constraint	NEMDE procures and prices the service
<i>Non-market Ancillary Service</i>	AEMO and provider negotiate contracts	AEMO calls on contract and provider commits unit	Provider delivers service and paid based on contract
<i>Ahead market Ancillary Service</i>	AEMO specifies service and contract template	AEMO holds contract tender to select providers and successful tenderers commit their units	Provider delivers service and paid based on contract

Because there is no spot market, settlement of these ahead contracts would be at the full ahead price, rather than the difference between the ahead and spot prices. It would necessarily be a gross market, because there would be no alternative channel for a supplier to sell these services.

To introduce such an ahead market, the new system service would need to be incorporated into the ahead market clearing engine using constraints similar to those required by NEMDE in dispatch: eg the “controllable variables” of the cleared ahead offer amounts would need to be expressed in a linear and continuous form on the LHS of the relevant clearing constraints, with the RHS expressing AEMO’s bid

⁶⁴ and this is perhaps the reason why the market has not set up its own ahead market platforms as it has done for forward markets

amount for these services. There would also need to be a reasonable level of competition in supply of this service to ensure value-for-money for those who would bear the eventual cost of these services. These are similar to the requirements required for introducing the new system service into the spot market. So any new service that could be introduced into an ahead market could also be introduced into a spot market.

A spot market would give the additional advantage of being able to adjust the amounts procured in the light of new information arising since the ahead market cleared. For example, if a supplier in the ahead market had a technical failure that meant it would be unable to supply as promised, a spot market would provide a channel to purchase additional service amounts to cover the shortfall. In the absence of a spot market, AEMO would be forced to direct.

In summary, whilst it is possible that an ahead market in a new system service might be a *complement* to a spot market in that service, it is implausible that it could be an *alternative* to a spot market.

4.4 POTENTIAL SOLUTIONS

4.4.1 Introduction

This section considers possible alternative approaches to addressing the issues that the ESB paper raises, which would not require the introduction of ahead markets of the form envisaged by the ESB.

4.4.2 Disclosure of Forward Positions

As noted above, if one of the objectives of introducing an ahead market were to give AEMO greater visibility of MP forward positions, this could be done more simply and more comprehensively by creating a new rule to that effect⁶⁵. This could require that participants disclose all contract positions at the day-ahead (say). Only aggregate MW positions would be relevant which, for option contracts, would also need to include the strike price at which the contracts come into the money. For PPAs, such as with semi-scheduled plant, the relevant plant could be named, or alternative the forecast output submitted.

Other confidential data – such as the counterparty, the forward price or other terms and conditions – would not be strictly necessary, although counterparty information might be used as a cross-check to ensure all contracts had been submitted.

It would need to be decided whether and how to submit contracts that are not in spot market products but that that might indirectly impact on MP scheduling decisions: for example financial contracts with insurers or weather derivatives.

The information would remain confidential to AEMO and not available to competitors or other institutions such as the AER.

This potential solution is only relevant if providing this information to AEMO would assist it in its own scheduling decisions. But for the reasons discussed above – that the forward market is voluntary,

⁶⁵ this might have some similarities to the existing Retailer Reliability Obligations rules where, under certain circumstances, retailers are obliged to disclose their forward contracts to the AER

financial and regional – this is considered implausible. So this potential solution is only offered to address a putative, inferred concern of the ESB's, not a real one.

4.4.3 Forward Purchasing of Market Ancillary Services

Rather than introducing an ahead market in market AS, an alternative would be for AEMO to procure these services in *forward* markets: ie months or years ahead of real time, rather than hours. These would be financial forward products (ie simple financial derivatives), since generators would not wish to commit to running specified plant so far ahead of real time. To the extent that the forward position must be defended by bids in PD and so improves the firmness of these bids, a forward purchase should be just as effective as an ahead purchase.

In effect, AEMO would be procuring these market AS through a contract market, similar to how it would need to procure non-market AS. AEMO would need to be guided by similar objectives of ensuring system security whilst obtaining value for money. The only difference between the two is that the non-market AS contracts are *physical* (they must be, because there is no spot market for non-market AS), whilst the market AS contracts are *financial*.

For example, AEMO might buy 1000MWsecs of synchronous inertia in SA from generator X. This is a financial derivative, so X must make a contract payment to AEMO based on the difference between the contract price and the spot price for this service. X will likely offer at least 1000MW of inertia into the spot market to ensure that it defends the contract, rather than take the risk of being short when the inertia spot price goes to the MPC, say. It can choose to commit *any* combination of synchronous units that provide this inertia amount, so this gives it more flexibility than a physical contract which would require specific units to run. Where it was short of units to run (eg if some were on outage), it could buy some inertia from another generator Y in the forward market in order to reduce its exposure. Y would then need to offer inertia into the spot market to defend *its* acquired forward position⁶⁶.

A forward contract has the additional advantage for providers of hedging risk around investment or disinvestment decisions. For example, a generator considering whether to close a synchronous unit might first see whether it can sell a one-year (say) inertia contract to AEMO⁶⁷ to cover the cost of keeping the unit operational for another year. This could be important where the spot price was volatile.

4.4.4 Clearing of Multi-step Dispatch

In section 2.4.4, the possible option of a multi-step dispatch was discussed. Only the first step would be used to set dispatch instructions, but the dispatch would be co-optimized over the multiple steps to make dispatch less myopic and so better coordinate generation to meet, say, fast run-up rates in demand. However, this model places a risk on generators that the forecasts used in this multi-step

⁶⁶ This implicitly assumes that inertia from generator Y can be physically substituted for inertia from generator X in dispatch, without any security concerns. But that assumption is anyway necessary for inertia to become a market ancillary service. Conversely, if inertia cannot be treated as a commodity – since each unit or generator delivers a somewhat different service – there can be no spot market in the service, and it would have to be managed as a *non-market* ancillary service.

⁶⁷ Or it might be sold to the TNSP where the TNSP is responsible for maintaining inertia, say.

dispatch don't eventuate. For example, a FSG with a start-time of 15 minutes, say, might be committed in the multi-step dispatch on the expectation of high prices in 15 minutes' time. But if these don't eventuate, the FSG has been started unnecessarily and is unlikely to recover its start-up costs.

This risk could be hedged by having a multi-step *settlement* based on the prices and quantities cleared in the multi-step dispatch outcome. So if the multi-step dispatch runs over a one-hour study period, there is in effect an up-to-one-hour-ahead market operating that clears in accordance with the dispatch outcomes. So this is a special kind of ahead market, tied directly into the dispatch process, rather than operating as a separate clearing process in PD timescales.

As dispatch is repeated the clearing becomes multi-stage. So consider the FSP again with a 15-minute start-up time. Suppose that it appears in consecutive multi-step dispatch runs as follows:

- In a 17:00 dispatch run: the FSP is committed and is dispatched for 100MW in the 17:15-17:20 interval at a clearing price of \$1000/MWh. It is paid $100\text{MW} \times \$1000/\text{MWh} / 12$ for this.
- In the 17:05 dispatch run: the FSP is no longer needed and it has zero dispatch in the 17:15-17:20 dispatch interval, with the clearing price now set at \$100/MWh. Settlement is for the *difference* between the prior cleared position and its new position: so it pays *back* to AEMO $100\text{MW} \times \$100/\text{MWh}/12$
- in the 17:10 and 17:15 dispatch runs it is again not dispatched, so there is no further settlement

So the generator is paid \$1000/MWh, then must pay back \$100/MWh, and so receives *net*, \$900/MWh for the inconvenience of having its start cancelled.

The bids used in this multi-step dispatch are the good faith PD bids. The clearing prices set by the multi-step dispatch ensure that the generator is paid at least its bid in each dispatch clearing. So these payments should always be satisfactory to the generator. Also, since dispatch instructions only apply to the first timestep, the ahead trades cleared in subsequent steps are *financial* in nature: there are no obligations to conform to the cleared ahead schedule in dispatch, since this is liable to change anyway prior to real-time, as discussed above.

A possible concern with this design option is that each generator will be paid a combination of spot and ahead prices, with the particular combination unique to each generator. Retailers will similarly face different price combinations, based on how their cleared bid amounts vary. Therefore, given each generator and retailer faces somewhat different prices in spot, it will be difficult to design forward contracts that hedge both sides, as they do today. There will inevitably be some basis risks remaining for one or both sides. However similar risks exist currently, given that 30-minute settlement can make it difficult for, say, a FSP to defend a forward contract, although these will be mitigated somewhat by the introduction of 5MS.

This option is a form of ahead market, in that transactions occur ahead of real-time; how far ahead depends upon the length of the multi-step dispatch study period. However, many of the concerns previously expressed around the ESB-style ahead markets do not arise. Because it is tied into dispatch, it will be a gross market, so it essentially provides AEMO the same visibility as PD in terms of potential

system insecurity. However, the ahead trades are financial rather than physical, in the sense that dispatch conformance obligations still only apply to the dispatch targets for the *first* timestep. On the other hand, as noted above, there remain possible concerns about impacts on forward markets.

The ahead market transactions in this option occur very close to real-time: eg up to one-hour ahead. This potentially is the time period in which any PD weaknesses are most apparent and material. For example, if existing PD is slow to converge or track (as discussed in section 3.2), it may struggle with the high ramp rates or weather forecast errors that might arise in this timescale in a high-renewables grid⁶⁸.

However, the same notes of caution should be made as were discussed in relation to possibly changing the PD clearing engine⁶⁹ to similarly take a multi-step approach. Clearing engines are just one component of the PD-dispatch architecture and changing them will cause generators to change their scheduling processes and bidding strategies. The *overall* impact needs to be assessed: again, by operating paper trials or similar. The systems and transactions costs of introducing this complex settlement architecture would also need to be considered.

There is also a question of the materiality of the risks that this option allows generators to avoid – compared to the alternative of having a multi-step dispatch but not multi-step settlement. The considerations here are similar to those for an ESB-style ahead market, as discussed in section 4.3.6. Multi-step clearing might introduce a lot of complexity for limited benefit in terms of risk mitigation.

Given these complexities, this option is being put forward very tentatively: as something to consider and explore further.

4.5 CONCLUSIONS

There are three hurdles that the ESB needs to clear if it is to justify the introduction of ahead markets. Firstly, it needs to identify and define the particular issues that will emerge in the future if the NEM design is not changed. But the ESB paper only talks vaguely about “firmness” of PD bids. It is not clear how this could be an issue, since PD “bids” are really forecasts, and “firmness” is not a property of forecasts.

Secondly, it needs to consider whether its issues could be better addressed through design changes elsewhere. The particular issue that the ESB seems to be concerned about, or extrapolating from, is the difficult AEMO faces when synchronous units drop out of PD, leaving AEMO with a system strength insecurity. As discussed, that arises primarily because the relevant system strength services are not priced or valued in the NEM, and so the market ignores this issue. So this could and should be addressed through introducing new market AS or non-market AS. To be fair, the ESB is considering this currently, albeit through a separate MDI.

⁶⁸ AEMO’s recent Renewables Integration Studies report emphasises the challenges that are likely to arise in the future in these timescales.

⁶⁹ Section 3.4.2.

Thirdly, even if the putative issue can't be addressed elsewhere in the NEM design, the ESB needs to demonstrate how it might be addressed by ahead markets. The ESB describes two possible nexuses between PD "firmness" and ahead markets.

- that a *financial* ahead market would lead to suppliers having a higher forward position to defend than currently. But it is unclear why that would be the case, given that ahead markets would give participants the opportunity to reduce, as well as to increase, forward positions.
- That a *physical* ahead market could place obligations on generators to physically supply in conformance with an ahead schedule. But, of course, AEMO is already able to place such obligations on generators: using powers provided under AS contracts or directions.

So, to summarise, it is not clear what the ESB issues are, or whether and how ahead markets could address them.

5 ESB OPTIONS

5.1 INTRODUCTION

The ESB has set out four options for introducing ahead markets into the NEM. These were described initially in the March 2020 paper to COAG, with some details subsequently being revised or clarified in meetings with stakeholders. These options are illustrative only and intended to represent a spectrum of possible design options.

The options are described at a high level and many important aspects of each option are unspecified or unclear. Furthermore, the precise design of an option – and its effectiveness and impacts – will depend upon design decisions made in other ESB work programmes: in particular in relation to two-sided markets and new system services markets. Thus, the eventual form and impact of any of these options is inevitably uncertain.

In this chapter, the proposed design for each option is described in terms of the market design concepts defined and explained in the preceding chapters of this report. The option is then assessed in terms of whether it will be effective in addressing current and emerging issues and also whether it might create new issues. Finally, the potential impacts on market participants are considered. Given the uncertainty around the designs, this analysis is necessarily at a high level.

5.2 ESB OPTION 1

5.2.1 Description

Option 1 of the ESB introduces a process called *unit commitment for security* (UCS) that AEMO would run in PD timescales. The March ESB paper describes this process and a recent ESB webinar has added some further detail. This is the only new element in Option 1: there are no ahead markets.

The UCS is essentially a decision support tool that AEMO would use when scheduling its intervention tools and resources to ensure system security, as described in section 3.2.8⁷⁰. The scope of the UCS process is quite similar to what AEMO does currently. However, there are two important changes:

- The UCS is envisaged to be a sophisticated scheduling algorithm designed for this purpose; the current approach relies on *ad hoc* manual analysis;
- AEMO will keep the market informed of its scheduling intentions and expectations through inputs to the PD process; it does not do this currently.

The first initiative relates to AEMO's internal operations and could, and really should, be undertaken today rather than waiting until post-2025: acknowledging that other post-2025 design elements – particularly the treatment of system services – will affect the design of the scheduler. AEMO anticipates

⁷⁰ Recall that these fall into three categories: (a) calling on delivery of non-market AS pursuant to existing AS contracts; (b) calling on reserve capacity in accordance with RERT contracts; (c) issuing directions.

the need for more comprehensive information on generator operating costs and constraints as inputs into the scheduler: eg in relation to directions. Again, AEMO can and does ask for such information today, so the only change proposed is in making this information provision systematic rather than relying on the existing *ad hoc* requests.

The second initiative is similar to the “AEMO good faith” solution option discussed in section 3.4.4. Again, it is not clear why AEMO needs to wait until 2025 to do this.

One aspect of the UCS remains vague: the objective function that the scheduler is solving against. This issue was discussed in section 3.4.5. As noted in that section, the preferred approach would be for AEMO to minimize intervention in relation to spot market services and aim for least-cost intervention in other services.

5.2.2 Assessment

The UCS concept is very much aligned with the proposed solutions for improving AEMO scheduling and coordination of this with MP scheduling, discussed in the pre-dispatch chapter⁷¹, and to that extent is to be welcomed. However, the uncertainty over the objective function remains the “devil in the detail”. A key concern is that the AEMO scheduling might unnecessarily interfere with – and even over-ride – scheduling decisions made by the market.

As discussed in the pre-dispatch chapter, plant operating costs and constraints are private information held by each generator and not disclosed to the market. Each MP will take these into account in its scheduling decisions, which are communicated to the market through PD bids. However, these bids themselves are not necessarily cost-reflective: indeed, the simple PD bid structure makes it difficult to accurately reflect complex cost structures. In any case, each generator is entitled to seek a margin above its costs and this is effectively managed by the competitive process rather than by regulation.

Under the Option 1 proposal, this private cost information is provided routinely to AEMO, who could potentially use it to generate its own cost-based schedule: not just for the plant that it needs to schedule for system security, but covering *all* plant. The two schedules – the AEMO schedule and the PD schedule – may be somewhat different, for several reasons:

- AEMO is not privy to all relevant private information: perhaps because it has not thought to request it;
- market participants incorporate other factors into their scheduling objectives: eg risk management
- It is not possible for any complex scheduler to always find an optimal schedule, so two schedulers might give different sub-optimal outputs, even with identical inputs;
- market participants may seek a margin over and above cost

⁷¹ Sections 3.4.4 (“AEMO good faith obligations”) and 3.4.5 (“clarity of AEMO scheduling objectives”)

The risk is that AEMO prefers its own schedule and uses its intervention powers to over-ride it. Of course, its ability to do this will depend upon how these powers are described and delineated.

Thus, it is important that AEMO's scheduling objective is appropriately formulated that the scheduler would not advise AEMO to over-ride the market in this way.

5.2.3 Impacts on Market Participants

Aside from this concern around AEMO's scheduling objectives, the new processes proposed by AEMO in option 1 should generally be helpful to market participants, in particular by providing more transparency and predictability around AEMO's intervention actions. Possible concerns are the practical costs and impacts of developing and operating the systems needed to provide routinely the additional cost information. There is also some concern that this information might be used for other purposes: eg by the AER in examining bidding strategies.

5.3 ESB OPTION 2

5.3.1 Description

This option adds a *voluntary forward market* (VFM) to the UCS proposed under option 1. There are few details provided. The VFM would cover energy and market AS and would operate over PD timescales. It could be a bulletin board or a periodic auction. The energy VFM would be two-sided. However, the AS VFM would be one-sided, with AEMO submitting bids. The ESB paper contemplates that the quantity of AS bids might be designed to "procure some base or minimum level before real-time" or "based on forecast system security gaps".

Penalties for non-conformance with ahead schedules are not specified, so it is unclear whether the VFM is intended to be financial or physical. However, there are some suggestions that it would be physical, at least for system services. For example, the ESB paper notes:

- it is required that offers are submitted "on a unit instead of a portfolio basis";
- "trades on the VFM could be restricted to those that are physically feasible", based on comparing any trade to unit status or capacity using PD information
- Participation of non-physical participants is not permitted
- "transactions on the VFM related to system services might need to be physically binding"
- The outcome of the VFM will provide "a firm schedule for delivery of systems services before UCS and real-time operation"

If the intention is for the market to be physical, then some further specification of non-conformance penalties is needed to assess the option.

5.3.2 Assessment

The ESB hopes that the VFM would firm up PD bids and lead to less need for AEMO intervention. However, as discussed in section 4.3, this largely depends on whether the market is physical and whether it is gross. A net, financial market would do little to firm up bids or to give AEMO greater visibility in PD.

As previously discussed, problems of insecurity emerging in PD usually relate to missing markets. It is implausible that any generator would submit a rebid that led to a shortfall of energy or FCAS⁷², because it would be then missing out on earning the MPC on the quantity it had withdrawn. It would always be preferable for it to maintain its offered availability and simply rebid this into an offer band priced at the MPC. That would increase prices, of course, but would not endanger system security.

Nor is the ESB suggesting that shortfalls in energy or market AS is a problem currently that needs to be addressed. Rather, the issue is usually that a PD rebid could lead to a shortfall in a *non*-market system service, such as system strength.

The VFM by definition can only trade in spot market services: energy, FCAS and whatever new system services (such as inertia) are introduced into the spot market in the post-2025 design. Again, there is no plausible reason for PD rebids to deliberately create shortages of these new market services. The problem remains confined to non-market services. The ESB must be hoping, then, that through the VFM placing physical obligations on delivery of market services, there is a lower likelihood of a gap emerging in *non*-market system service. For example, that a VFM in a new inertia market AS will lead to a firmer commitment of synchronous units in PD, which will also help provide certainty and visibility around system strength provision.

The problem with this expectation is that different plant might be providing the different services. For example, inertia might be provided by non-synchronous technologies (eg as virtual inertia) that do not contribute to system strength. So purchasing inertia through the VFM won't necessarily address a system strength shortfall. And having more certainty that the inertia providers will be physically present in dispatch does not provide any surety that a system strength gap won't emerge in PD.

In any case, it is anticipated that a contract market will be developed for the system services that are not able to be formulated as new market AS. AEMO will then schedule the necessary units, in accordance with the contracts, as needed to provide the quantity of services needed for security. Since AEMO is *itself* the scheduler for these services, submitting its own bids into PD, the firmness of PD does not arise. So if, say, AEMO considers that 3 synchronous units are needed on-line for system strength, it will schedule these units in accordance with the contracts it holds. The shortfall issue does not arise and there is no reason to set up an additional VFM for these services for AEMO to top-up its supply.

5.3.3 Impacts on Market Participants

If the VFM is genuinely voluntary, market participants don't necessarily need to trade in it at all and therefore might not need to build the necessary trading and settlement systems. However, even if the

⁷² unless, of course, this reflected a physical failure

market is *strictly* voluntary, an MP choosing not to trade in it might put themselves at a commercial disadvantage.

Suppose that ahead prices were higher (on average) than spot prices. In a two-sided market, a higher ahead price would discourage retailers from participating and so this might serve to remove the price difference on average. However, in a one-sided market, AEMO is the proxy bidder and might not have the same financial incentive to obtain best value-for-money. At worst, AEMO could bid *all* of its market AS needs into the ahead market, leaving the AS spot market as just a residual shell.

If there were commercial imperatives to participate, market participants would need to develop the necessary trading and settlement systems to do so. To the extent that the market was physical, they would need to also ensure that plant operated in accordance with the ahead schedule, which differs of course from the current situation where *any* plant (or even no plant at all if prices are expected to stay low) can be run to defend forward positions.

On the financial side, contract positions would now be the aggregate of the forward market and ahead market positions. Generators might need to adopt a lower level of forward cover to provide headroom to sell into the ahead market, particularly where AS and energy markets were linked, and AS ahead sales could necessitate energy ahead sales too. For example, a generator selling inertia from a synchronous unit might be obliged to sell the minimum energy output of that unit at the same time.

Alternatively, new forward products might be developed which reference *ahead* prices rather than spot prices. This is typically the case in overseas markets with well-established ahead markets. Generators would then need to consider how to offer into the ahead market to defend these ahead-linked contracts, in the same way that they do currently to defend spot-linked contracts.

In summary, depending upon the detailed design of this option, there might be significant commercial disadvantages of not participating in the VFM, but also substantial operational complexities of trading in this new market.

5.4 ESB OPTION 3

5.4.1 Description

In this option, the PD process is converted into an ahead market, for one or more PD runs. Thus, rather than the cleared prices and quantities from the PD engine simply being forecasts of spot market clearing, they can now represent actual cleared trades.

So, for example, consider a 12pm PD run that shows generator A selling 10MW of an FCAS service at \$30/MWh for the period 6:00-6:30am the following day. Currently, this is simply a *forecast* that this will occur in the spot market the next day. Under the ESB option, this might now become a *transaction*: generator A has actually *sold* that product at those prices and amounts. It is intended that the ahead products are *physical*, so generator A must now ensure that the relevant plant runs in dispatch so as to provide the sold service.

It is not proposed that *every* forecast transaction from the PD run becomes an actual ahead transaction. This would only be the case for:

- “critical system services”: a term which the ESB has yet to define
- other services where the seller *volunteers* to sell on the terms established in the PD run

It is not clear how and when the seller would decide whether to agree to the trade. Possibly, this might be expressed before clearing: eg as a condition in the PD bid. Or the choice may be offered *after* PD clears: the seller first sees the cleared price and quantity from the PD run and then decide whether it wants to make an actual sale on those terms.

It is also not clear how this voluntary process would operate in a two-sided market. Presumably, there would need to be agreement from both the supply side and the demand side that the PD trade should become a binding contract. Indeed, the role of the demand side is generally unclear. Would retailer be submitting good faith forecasts into PD or would AEMO perform this role? Where AEMO is the proxy bidder – eg for market AS – what prices and volumes should it bid, given it has the option of procuring some amount of these services in dispatch instead?

PD iterates currently, being re-run every half-hour. The ESB has not defined which of these PD runs would be used to clear the ahead market. Potentially, there could be repeated clearing. For example, PD runs at 6am, 12pm and 6pm might all provide ahead market clearing for the following day, with the other PD runs being “for information only” as now. In this case presumably the sequential trades would not be additive but incremental. For example if a generator’s offer for a particular product cleared 10MW in the 6am run and then 15MW in the 12pm run, the latter trade would only be for 5MW, adding to the 10MW already sold at 6am.

The possibility of changes to the PD engine are raised. For example, the PD engine could co-optimize across multiple timesteps or even make commitment decisions. These concepts were discussed in section 3.4.2, and could be introduced in PD anyway, without the need for an associated ahead market.

5.4.2 Assessment

A critical conceptual flaw in this option as described is that it conflates two entirely different processes: the PD forecasting/coordination process and the ahead market clearing process. It requires that market participants submit a single set of bids in the ahead timeframe, which is then used for dual purposes:

- To drive pre-dispatch, as now; and
- To bid into an ahead market.

This is illustrated in figure 18, below.

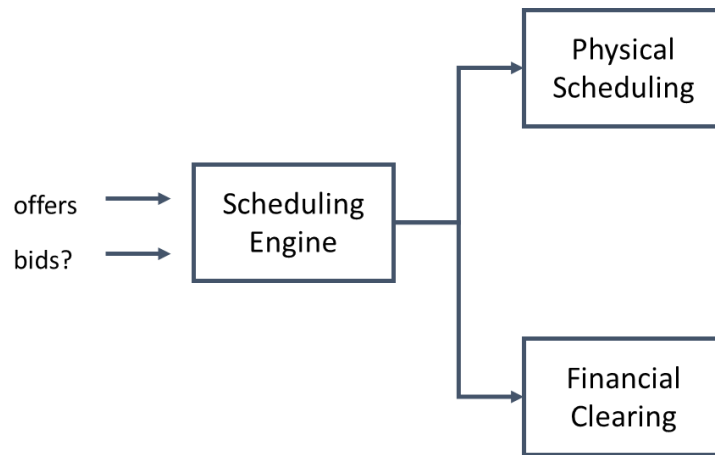


Figure 18: Conflating the Scheduling and Ahead Market Processes

Asking market participants to submit a bid that can apply to both processes puts them in an impossible position. For example, suppose a generator intends to offer a critical system service in dispatch at \$30/MWh, but wishes to offer into the ahead market at \$50/MWh⁷³. Current good faith obligations would require that it bids at \$30/MWh, since this is how it intends to bid into dispatch. But when submitted into the ahead market, this bid might then be cleared at \$30/MWh, a price that is unacceptable to the generator, who is seeking at least \$50/MWh. So the generator has a choice of bidding in “bad” faith or selling at a “loss”. It is a poor market design which places the generator in this invidious position. It is analogous to a real estate agent forcing every bidder to bid for two different properties at the same price, so each time you place a bid for property X this automatically becomes a bid for property Y too. This would be unfair and unreasonable and not conducive to obtaining the best sale price for either property.

Presumably, the ESB’s intention is to make the ahead market *de facto* gross⁷⁴. Dispatch is gross and, since PD bids must reflect dispatch intentions, they must be gross too. Since participants are forced to submit PD bids into the ahead market, the ahead market becomes gross too. For non-critical services, the supplier can perhaps choose not to agree to the ahead market clearance, so that might remain a kind of net market, although in a complex and unconventional way: ie gross bidding but net clearing.

Incidentally, this is *not* how ahead markets operate overseas. In those market designs, physical scheduling occurs through participants submitting current operating plans (COPs) providing details of how they intend to operate their plant in the spot market, similar to PD. Bidding into ahead markets is separate from this. So ahead markets can remain net whilst the COP schedule is gross. Of course, COPs must reflect ahead market commitments, but that is not the same thing as submitting the same information to both processes. This overseas architecture is illustrated in figure 19.

⁷³ it might seek a higher price in the ahead market, given the associated loss of the optionality that exists in a PD bid

⁷⁴ As discussed in section 4.2.7, only a gross market provides AEMO with useful information on system security.

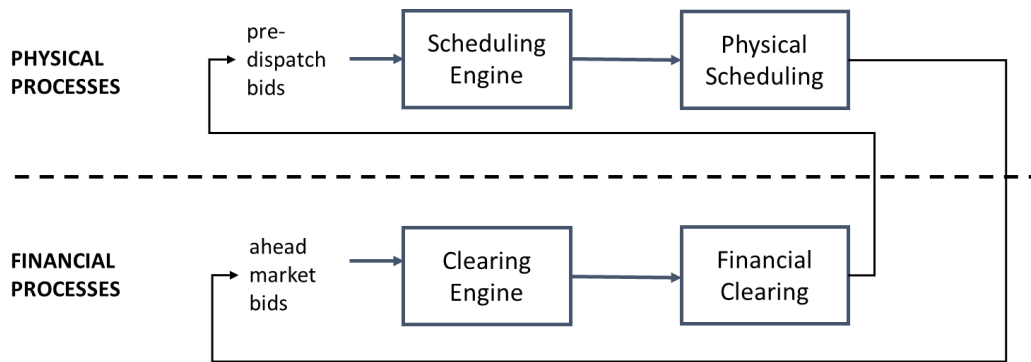


Figure 19: Separation of Scheduling and Ahead Market Processes

Presumably, the intent in making the ahead market gross is to allow for a full co-optimisation of all services against all security constraints, as occurs in PD and dispatch. This could perhaps be achieved in other ways, by making the spot market relatively unattractive, as previously discussed. However, this begs the question as to whether a gross ahead market is even desirable.

The ESB’s underlying objective is to improve scheduling and coordination of market participants and AEMO through the PD period. Since this is the role of the PD process, this is the obvious area to look at for possible reform. For example, the ESB’s suggestion of introducing multi-step clearing into the PD engine for the purposes of ahead market clearing could more simply be introduced into PD instead. There is no necessity for an ahead market here.

Achieving a gross ahead market requires the undermining of either PD (as the ESB proposes) or the spot market (to make it unattractive). Either approach is liable to adversely impact the effectiveness of the current processes.

5.4.3 Impacts on Market Participants

As discussed above, forcing generators to offer the same prices in ahead and spot markets puts them in an invidious and unnecessary position. Could “good faith” be stretched to allow for ahead market bids that are not genuine forecasts of spot market intentions? If so, there could be peculiarly discontinuities between the PD runs used for ahead market clearing and the remaining PD runs used solely for forecasting.

As with option 2, generators would need new trading and settlement systems to accommodate the new ahead market trading. Furthermore, if the ahead market were to be gross, as the ESB appears to intend, this would have implications for forward contracts, which would need to be changed and renegotiated to refer to ahead prices rather than spot prices. The ESB paper prefigures this possibility and recognizes the potential for disruption in the forward market.

Generators would also require new tools to support decisions as to whether to voluntarily accept the cleared ahead sales for non-critical services. This would be particularly complex where different services were co-produced and co-optimised. For example, a generator might offer a synchronous unit to provide inertia (a critical system service, say) into the ahead market and an offer to generate energy at

minimum load would have to accompany this. If the inertia offer were cleared, the generator might then need to decide whether to accept the accompanying energy clearance too. The unit must be physically run either way: the question is whether the energy output will be paid the ahead price or the spot price.

If the PD engine and ahead market clearing engine were designed to use structured bids and made additional scheduling decisions -eg central commitment – generators would need to also change their trading systems and bidding strategies to accommodate this.

5.5 ESB OPTION 4

5.5.1 Description

Option 4 is similar to option 3, but with one key difference: agreeing to the cleared trades would no longer be voluntary, for *any* services. Thus, the PD result for the ahead market run would be fully cleared.

There are suggestions from the ESB paper’s discussion of this option that the design implications have not been fully thought through. For example, it is also suggested that a gate closure *might* be introduced into rebidding prior to dispatch, with only rebids that do not change the maximum available capacity of a unit being permitted after the gate closure, unless for physical reasons: eg plant failure. But if the ahead market is physical – as the description strongly implies – such rebids are already in effect prohibited for plant that has been sold into the ahead market: it must run in accordance with its ahead schedule or be penalized for non-conformance. And since the ahead market is gross, this will constitute nearly all plant.

The paper also suggests that semi-scheduled generators might only bid the “firm” component of their capacity into the ahead market: eg a 90% probability-of-exceedance forecast of output rather than the usual 50% POE forecast used in PD. But since it is PD bids that are submitted into the ahead market, good faith obligations will still require the 50% forecast: a 90% exceedance forecast of output would be “bad faith”, unless the generator genuinely intended to operate the plant in this way in dispatch – which it wouldn’t of course⁷⁵. So again, there is a conflict in forcing generator to submit the same bids into PD and into the ahead market.

5.5.2 Assessment

This option suffers from the same critical flaw as the previous one: that the same set of bids must be submitted into the ahead market and into PD (reflecting intentions for the spot market). The compulsory nature of clearing worsens this anomaly: it now applies to all services rather than just critical system services.

If the description is taken at face value (and, as discussed above, there are suggestions that the ESB has not fully understood the implications of this), there will be a gross clearing of the ahead market, leaving the spot market to clear just residual differences due to forecasting errors.

⁷⁵ It wouldn’t generally voluntarily constrain the output of a zero-cost renewable plant to this artificially low level.

As noted above, this forces generators to commit all plant in accordance with the ahead market clearing. If it turns out that the associated forecasts were too high, some of this committed plant will be surplus to requirements, but would nevertheless be required to run at minimum load. If the forecast were too low, there could still need to be additional plant committed, and this would need to be scheduled through the PD process as now: or perhaps through additional ahead markets closer to real time. So it doesn't necessarily fully address the "firmness" issue.

The gate closure rule, if applied symmetrically⁷⁶, could perversely prevent commitment of *additional* generation close to real-time, potentially creating artificial supply shortages: where additional resources are operationally available to commit, but are prohibited from doing so by the gate closure rule. Alternatively, if the rule is applied asymmetrically⁷⁷, this could lead to generators committing *less* resources in PD prior to the gate closure (to the extent that this is consistent with "good faith"), to retain the option as to whether or not to run them in dispatch. This might lead PD runs to artificially signal supply shortages, which might prompt unnecessary intervention⁷⁸.

5.5.3 Impacts on Market Participants

The impacts of this option are similar to those described for option 3, although there is no longer the issue of deciding whether to voluntarily accept cleared ahead trades, since this is now mandatory. The possibility of a rebid gate closure could complicate scheduling assessments in PD and cause uncertainty over how to ensure good faith rebidding whilst prudently retaining some optionality post-gate-closure.

5.6 US MARKET DESIGNS

It is probably no coincidence that US electricity markets commonly incorporate processes similar to the ESB's proposed UCS: typically called "reliability unit commitment" (RUC). US markets also have ahead markets (usually day-ahead) and central commitment. It seems reasonable to infer that, in proposing and developing ahead market concepts, the ESB is drawing on and guided by US experience.

It is sensible and worthwhile, in a major design review such as the post-2025 project, for overseas markets to be examined and lessons and ideas to be drawn. Of course, they must be the right lessons and ideas, and this does not seem to be the case here, for several reasons.

Firstly, these design aspects of US markets are long-standing and were implemented long before the prospects of grid decarbonisation and transition came on the scene. Indeed, even today, US grids are generally not facing the associated issues – such as weather uncertainty and loss of synchronous generation – that are driving the NEM review. US spot markets typically trade the same services – energy and FCAS – as the NEM, and their ahead markets reflect this. It is not clear how the US designs

⁷⁶ ie prohibiting any rebids which either decreased availability or *increased* availability

⁷⁷ rebids to *increase* availability are permitted, whereas rebids to *decrease* availability are prohibited

⁷⁸ In fact, this is a general problem with attempting to make PD bids "firmer" by encouraging generators to be more conservative in their forecasts, making the PD outlook worse than it really is. It is vital for all parties that PD provides the best possible forecasts of dispatch conditions. If AEMO wishes to be conservative in its interventions, this should be through appropriate scheduling objectives (as discussed in section 5.2.2), *not* by distorting pre-dispatch.

would cope with new services such as system strength and synchronous inertia and whether these would be incorporated into spot and ahead markets.

Secondly, US market designs do not conflate the physical process of generator scheduling with the financial process of ahead trading in the way that the ESB's proposal do. Whilst ahead market clearing engines do allow for central commitment, this is generally voluntary and non-binding. Ahead markets outcomes do not, by themselves, guarantee security; neither do they feed into the RUC/UCS commitment processes. Instead, the physical processes are driven by information, referred to as "current operating plans" (COPs), that generators are mandated to provide to the system operator, analogous to (although less detailed than) the information that NEM generators must provide to pre-dispatch. COPs do not feed into ahead market processes in the US, and pre-dispatch bids should not be fed into ahead markets in the NEM.

Thirdly, US spot markets are cleared nodally – at least on the generator side – and use nodal spot energy prices in settlement. Ahead markets, because they reflect the spot market design, are similarly nodal. Thus, although these ahead markets are financial, each bid must nevertheless specify the transmission node at which it applies. The pattern of ahead market prices can then provide some useful physical information, in terms of the likely location and severity of transmission congestion in dispatch. Because the NEM market is regional, an ahead market would not and could not reveal similar information, and the ESB should not expect it to do so. Thus US ahead market designs would need to be adapted in order to operate effectively in the regionally-priced NEM.

Finally, the US markets have always had ahead markets; in many cases these pre-dated their spot markets. This means that ahead market participation and liquidity is high, even where participation is purely voluntary. Forward financial contracts will therefore typically reference the ahead price, leading to increased participation to "defend" these contracts. So there is a virtuous circle of liquidity begetting liquidity.

On the other hand, the NEM has never had ahead markets and already has established financial forward markets that reference the spot price. Were a voluntary, financial ahead market to be introduced in the NEM, it is likely that participation and liquidity would be low, given that potential participant could and would continue to rely on the spot market. A thin ahead market would provide limited hedging opportunities market participants and no visibility or "firmness" benefit for AEMO.

5.7 CONCLUSIONS

The ESB describes four design options. The first involves a change only to PD, by formalizing AEMO's scheduling in a new UCS process. The other three propose to introduce ahead markets in various forms.

The UCS process in option 1 seems likely to provide benefits to the market, by both formalizing the existing AEMO processes and better integrating them into PD. In effect, it ensures that AEMO provides "good faith" information into PD in the way that market participants are required to do. This will assist with coordination and should lead to more efficient and secure dispatch outcomes.

The other three options are difficult to assess, because they do not provide coherent and complete description of how the new ahead markets would operate. Any ahead market design description needs to define:

- *What products are being traded:* in particular, whether these are just financial derivatives against spot prices or whether they also include some physical obligations. If the latter, what are these obligations and what are the penalties for not conforming with them?
- *Who participates and how:* what form do bids take; who is required to participate; if participation is "mandatory" what does this mean and what are the penalties for not participating in this way?
- *Who represents the demand side:* are bids for energy submitted by retailers; are bids for market AS submitted by AEMO; if so, what are AEMO's bidding objectives?
- *How and when do the ahead markets clear:* what is the clearing objective and algorithm; when does clearing it occur; are there multiple stages; can participants choose not to execute a trade that has been cleared?
- *How is the spot market design affected:* what trading remains on the spot market? What are the benefits and costs of spot vs ahead trading? If there are conflicts between ahead and spot trades (eg in use of scarce transmission capacity) how are these resolved?

None of these questions are answered clearly or satisfactorily in the descriptions in the ESB paper.

Notwithstanding this, it is difficult conceptually to envisage how or why *any* form of ahead market could improve scheduling and PD. Certainly, all of the ideas being considered by the ESB seem more likely to damage PD or the spot market than enhance them.

6 OVERALL CONCLUSIONS

6.1 OVERVIEW

The NEM does not have ahead markets and seems to have operated successfully – in terms of efficiency, security and reliability – to date. So why are ahead markets being proposed for inclusion in a post-2025 NEM design? The ESB describes three areas of concern that it considers might be addressed by ahead markets: a lack of “firmness” in pre-dispatch bids; a potential inability of the existing pre-dispatch process to schedule generation and demand response effectively; and the difficulty that AEMO may have in managing security effectively through market intervention. These areas are causing the ESB some concerns currently, and these concerns are expected to grow as new challenges emerge in the energy transition.

There is another possible reason for considering ahead markets that, although not explicitly stated by the ESB, might be inferred: that other electricity markets – in particular in the US – have ahead markets, so why don’t we?

Conclusions on these areas are presented in turn, below.

6.2 FIRMNESS OF PRE-DISPATCH BIDS

The ESB sees problems arising from the fact that bids into pre-dispatch are not “firm”. By this it means that a unit initially bid into pre-dispatch can potentially be withdrawn, through a rebid, prior to dispatch. This could cause pre-dispatch to be insecure, requiring AEMO to intervene at short notice to restore security.

This issue arises in the provision of a services (such as system strength) that are *not* priced and paid for currently, so generators are not incentivised to take account of a possible system strength shortfall when they bid. In contrast, if the insecurity had instead come about due to a shortage of a spot market service – energy or FCAS – the forecast price would then rise to the market price cap and supply adequacy would be quickly restored (or, more likely, the rebid that caused the scarcity would never have been made in first place).

So the obvious solution to this particular issue is to incorporate system strength (or whatever service it is whose shortfall is causing this insecurity) as a spot market service. Of course, this may be easier said than done. But if the relevant service cannot be incorporated into the spot market, it cannot be incorporated into an ahead market either. The ahead market has to reference the spot market financially and/or physically, so the same services must be traded in each market.

Alternatively, for the sake of argument, suppose that non-firmness remains a concern even when the relevant services *are* priced. How might ahead markets then improve firmness? The ESB posits two possible mechanisms. Firstly, that a financial ahead product might encourage generators to bid more firmly into pre-dispatch to defend that sale, just as generators currently are incentivised to defend their forward contract position. But given that these forward positions already exist – and yet firmness is still a concern for the ESB – an ahead market can logically only make bids firmer if it causes generators to

take on a *higher* forward position than they would in its absence. That is possible, but so is a *lower* forward position, making “firmness” worse, according to the ESB’s argument.

If the ahead market were, instead, physical, generators might be required to physically generate in line with their ahead schedules: ie the quantities sold into the ahead market. But this begs the question as to why a generator would voluntarily enter into an arrangement where it is locked into a potentially unprofitable dispatch. Indeed, locking in ahead schedules might be bad for the market as a whole and could even endanger system security: eg if too *much* generation is now committed.

In summary, pre-dispatch “non-firmness” is a strength, not a weakness, of today’s NEM. But, even if it were a problem, ahead markets would not be the way to address it.

6.3 SCHEDULING EFFECTIVENESS OF PRE-DISPATCH

Scheduling is the process of preparing generation to be available for dispatch. With the energy transition, the scheduling problem could become more complex. So there is a plausible argument that the current pre-dispatch process used for scheduling, whilst effective currently and historically, might not remain fit-for-purpose over the long term.

But if the issue lies with pre-dispatch, the obvious solution lies in reforms to this process, rather than by introducing a new, ahead-market process. The ESB should have, but has not, analysed the pre-dispatch process and explored options to address potential future weaknesses.

The decentralised scheduling architecture used in the pre-dispatch process is powerful and effective. It appears that the ESB is worried that the lack of control and visibility that self-commitment provides to AEMO could make it hard for it to manage system security. In particular, with the provision of new system services, such as system strength and synchronous inertia, being highly dependent on commitment decisions, a more centralised commitment process might better ensure the security and efficiency of the scheduling process.

This could potentially be done through changes in pre-dispatch: ie by changing the PD clearing engine so that it makes commitment decisions. This is a plausible – if likely undesirable – market design and would address the ESB’s firmness concern. And, again, this does not involve or require ahead markets.

But the ESB’s proposed vehicle for central commitment is ahead markets, by clearing one or more pre-dispatch outcomes – with actual ahead *transactions*. Generators might then be financially or physically locked into this ahead schedule.

In combining pre-dispatch and ahead markets in this way, the ESB appears to be conflating two different processes: the pre-dispatch process which is a de-centralised scheduling algorithm to help prepare generation for dispatch; and an ahead market auction which clears actual ahead transactions and places associated financial and physical obligations on sellers. In trying to combine the two processes, the outcome will be that neither is done well and both are done badly. The scheduling process will be critically undermined.

The ESB needs to ensure that, in its design proposals, the two processes – ahead market and pre-dispatch are kept separate.

6.4 AEMO INTERVENTION

AEMO has tools available to it that can be deployed in the pre-dispatch period to ensure security or to rectify any emerging insecurity revealed in the pre-dispatch process. Based on pre-dispatch results, it must identify any anticipated insecurity in dispatch and then select and deploy the appropriate response. Where it has a choice of options, it needs clear intervention objectives – relating to both security and cost – and be able to select the option which best meets these.

In many respects, these challenges are similar or analogous to those facing generators when making scheduling decisions. Whilst their objectives will be different, the options – committing units etc – and timescales are similar. Both rely on pre-dispatch outcomes to inform and guide their decisions. Both are affected by complexity and uncertainty in pre-dispatch forecasts. Given this, an obvious area to look for a solution is for AEMO to develop scheduling tools analogous to those that generators have already developed.

The ESB is proposing a new “unit commitment for security” (UCS) process that does this. This is a useful initiative that is to be supported. This scheduling tool should make AEMO better able to deal with uncertainty over the pre-dispatch period (the “lack of firmness” discussed above) and removes any need to attempt to lock in ahead schedules with strong central commitment: whether using ahead markets or by other means.

The UCS proposal also addresses another concern over AEMO’s intervention process: that it lacks transparency. By flagging intervention intentions ahead of time, the pre-dispatch process will allow the market to adapt to these more rapidly and effectively.

A critical aspect of the UCS design is the “objective function” that it uses to select the preferred intervention option: which, of course, might in some instances be “do nothing”. This will have to trade off the cost of intervention (both the direct cost and the indirect cost of distorting market outcomes) against the likelihood of dispatch insecurity if intervention does not occur. If the objective is too strict, AEMO might be constantly intervening to head off the remotest prospect of dispatch insecurity. Too lax, and insecurity might lead to major grid and customer impacts. Objectives need to be balanced and transparent.

6.5 US MARKETS

In proposing and developing ahead market concepts, the ESB seems to be drawing on and guided by US experience. It is sensible and worthwhile, in a major design review such as the post-2025 project, for overseas markets to be examined and lessons and ideas to be drawn. But US markets are different to the NEM in their history, design and issues. The US has always had ahead markets, whilst the NEM has not. The US markets price spot energy nodally, whereas the NEM has a regional design. And US markets are not yet facing the energy transition issues to the degree the NEM is facing; and to the extent they are, they are not proposing solutions based on ahead markets.

But perhaps most significantly, the ahead market designs proposed by the ESB do not follow the US examples. Whilst the ESB conflates the physical scheduling process in pre-dispatch with ahead market bidding and clearing, US markets manage these things quite separately, although of course there are commercial and physical interactions between the two

6.6 FINAL CONCLUSION

When considering the introduction of new design elements into the NEM, three hurdles need to be cleared: the issues to be addressed need to be real and material; the new elements must be able to address those issues; and alternative, simpler approaches to addressing them must be explored and ruled out as less effective. These hurdles have not been cleared by the ESB, in any of their expressed areas of concern.

APPENDIX: OVERSEAS MARKETS

ERCOT

Introduction

Electricity markets in the US have some broad similarities. ERCOT has been chosen for scrutiny as it is probably closest in nature to the NEM. It is an energy-only market, with a relatively high market price cap. It is also an AC island, with a significant penetration of non-synchronous renewable generation. Like other US markets, ERCOT has both an ahead market and a real-time market.

The ERCOT architecture is shown schematically in figure A1, below.

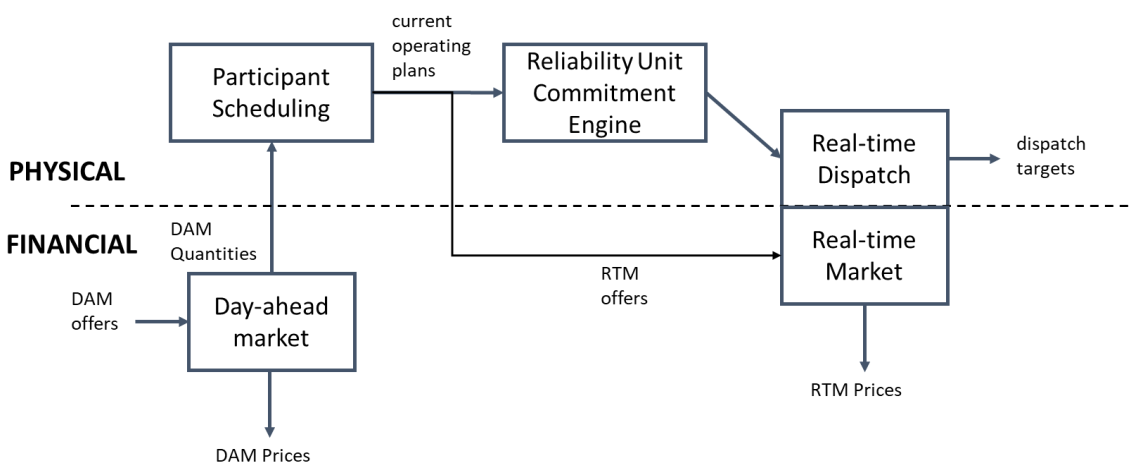


Figure A1: ERCOT Market Architecture

Real-time Market

The Real-time Market (RTM) is broadly analogous to the NEM dispatch market, although pricing is nodal (LMP) rather than regional. Energy-only bids are submitted from every physical plant and are cleared using 5-minute dispatch, with the five-minute dispatch prices then averaged to create 15-minute settlement prices.

As well as energy, ancillary services are cleared: these are just FCAS, again analogous to the NEM. The RTM also has a reserve price, known as the Operating Reserve Demand Curve (ORDC) which is added to the energy price and reflects the value of on-line or fast-start capacity. It is based on a LOLP x VOLL formulation where:

- LOLP is loss of load probability: based on the likelihood that there will be load shedding due to insufficient capacity to meet in the next hour.
- VOLL is the value of lost load: in this case the ERCOT market price cap of \$9000/MWh

All qualifying capacity – whether it is generating energy or is in reserve – is paid this reserve price.

Day-ahead Market

The day-ahead market (DAM) is a financial market that is cleared the day ahead of real-time. There are two types of energy bid structure in the DAM:

- an energy-only bid: offers for incremental energy, similar to bid structures in the NEM
- a three-part bid: with start-up cost and minimum generation costs, together with incremental energy offers for output above minimum.

The DAM is cleared with a scheduling engine which makes commitment decisions based on the three-part bids. However, because it is a financial market, there is no obligation on generators to then commit their plant in accordance with the cleared DAM schedule.

Energy-only bids are financial only and do not need to refer to physical plant. So non-physical speculators can and do trade in the DAM. This makes it likely that the DAM schedule will be a reasonable reflection of the expected RTM outcome for the following day, because to the extent there are any consistent differences, arbitrageurs would be able to make money by trading in the DAM.

Three-part bids, on the other hand, must refer to physical plant, for a few reasons:

- the extra two parts are regulated: prices must be no higher than double the true underlying costs, which must be declared;
- a three-part-bid that is cleared in the DAM is entitled to “make whole” payments

The make-whole payment ensures that a cleared bidder at least covers its start-up and minimum-load costs in the DAM. This implies (to me) that the clearing prices in the DAM are based on the incremental energy offers only. In any case, the RTM prices must be: and, as noted above, DAM prices are likely to reflect RTM prices because of arbitrageurs. The make-whole payment is only payable if the plant is physically committed. So, this may be an incentive to commit in accordance with the DAM schedule (and makes these cleared trades kind of physical).

FCAS offers have to reference the plant that is to provide the FCAS, so it appears that non-physical speculators cannot bid directly into this market. However, because the DAM is financial, there does not appear to be any obligation on FCAS providers to physically provide the FCAS in accordance with the cleared DAM schedule.

Nodal Pricing

ERCOT sets energy prices nodally, based on locational marginal prices (LMPs). These LMPs factor in congestion prices but not losses. Both the day-ahead market (DAM) and real-time market (RTM), run a security-constrained dispatch which incorporates transmission constraints. This means that bids into the DAM, even bids from non-physical speculators, have to reference the node that the bid relates to.

ERCOT holds auctions, months ahead of time, to issue so-called congestion revenue rights (CRRs), which are basically financial transmission rights (FTRs): hedges against nodal price differences. Importantly, these CRRs are struck against the day-ahead LMPs from the DAM. Thus, a holder of FTRs is hedged against nodal price differences in the DAM, not the RTM. To convert the FTRs into real-time hedges,

holders must bid into the DAM in the form of point-to-point (PTP) bids. These are essentially bids to buy FTRs that are struck against RTM nodal prices. Again, these are purely financial bids, so speculators can participate in this market too.

So, for example, a market participant who has 100MW of CRRs from node A to node B would likely then bid for 100MW of PTP from A to B in the DAM. The CRR hedges it against the DAM price outcome and so a DAM purchase converts the CRR into a real-time FTR at no additional cost. So, this provides an important incentive to bid into the DAM.

Net Settlement

In settlement of the RTM, the DAM trades are netted off the meter readings. There are other “trades” referenced in the settlement algebra, which it is understood refer to bilateral physical forward trades. For example, suppose that Gen A sells 100MW forward to retailer X. In RTM settlement, 100MW is netted off the metered output of A and also off the metered load of X. This has some similarities to the “reallocation” provided for in the NEM.

Presumably, there are also financial forward contracts that are settled by the counterparties directly rather than by ERCOT. These could potentially be struck against the DAM or the RTM prices. Since CRRs are struck against the DAM, it seems likely that forward contracts are also. However, this is unclear.

Pre-dispatch

ERCOT does not use the term pre-dispatch, but two processes in ERCOT perform a similar role to the NEM pre-dispatch process:

- current operating plans (COPs)
- reliability unit commitment (RUC)

All generators must submit their COPs up to 7 days ahead, similar to pre-dispatch. The COP contains the expected commitment status of each unit, together with other key operating parameters, such as maximum and minimum generation. COPs can be adjusted right up until real-time.

Unlike in NEM pre-dispatch, there seem to be no offer prices included in these COPs and so no associated forecasts of RTM prices. Also, there are no obvious “good faith” obligations referred to, but bids in US markets generally are much more highly regulated than in the NEM, and it would be expected that any “bad faith” COP behaviour would be prohibited.

COPs feed into ERCOT’s RUC process, which has similarities to the NEM’s existing directions process and, more closely, to the UCS process proposed by the ESB in its ahead market options. In the RUC, ERCOT runs the plant that is made available in the COPs through a security-constrained dispatch to check for any unreliability or insecurity. Where this is identified, some additional unit commitments (or possibly decommitments) are made and generators must comply with these.

Three-part bids are fed into the RUC and the clearing engine finds the cheapest commitment options that address the unreliability or insecurity. These bids may be leftover, uncleared bids from the DAM (since the DAM clears prior to the RUC run) or they might be separate bids. The regulation, noted

above, that bids must not be more than twice cost, still applies. RUC-committed units are paid their start-up and minimum generation costs, and any associated payment in the RTM is “clawed back” so that the generator is not paid twice. The “make-whole” costs (the difference between the bid cost and the RTM energy payments) are funded by retailers that have a “capacity shortfall” day-ahead. This is based on a complex formula, but might provide an additional incentive to participate in the DAM.

Discussion

There are some important differences between ahead processes in ERCOT and the options that the ESB is proposing for the NEM. Firstly, the RTM and DAM are nodal, so that the financial DAM physically represents transmission constraints and congestion prices. A financial ahead market in the regionally-priced NEM would, at most, reflect inter-regional transmission capacity and congestion: unless of course nodal pricing were to be introduced into the post-2025 NEM design.

Secondly, the ERCOT design does not conflate the operational (physical) pre-dispatch process and the financial ahead market, as the ESB options appear to do. The COP/RUC and DAM processes operate in parallel (although the DAM clears before the RUC is undertaken) and participation is quite different. There is an obligation on all generators to submit COPs, but no corresponding obligation to participate in the DAM. Conversely, non-physical speculators can participate in the DAM but clearly have no role to play in COP/RUC.

There is no obligation on generators to submit three-part bids – and to be “scheduled” – in the DAM, although many generators do this. As noted above, there may be some benefits to this, in terms of receiving additional “make-whole” payments. But it may also simply be helpful for generators in deciding which plant to physically commit. Note again, that the DAM schedule is not binding, so whilst the DAM clearing engine is in the form of a scheduler, there is no mandatory commitment of units by ERCOT, except through the RUC process.

As discussed in the main section of this report, the NEM pre-dispatch process performs this role of helping generators to schedule slow-start plant, so it is not obvious what complementary role a DAM similar to that in ERCOT might play here.

Although the RUC is referred to as a reliability unit commitment, it is not clear how much of a reliability role it plays (in the sense of committing additional capacity to cover unexpectedly high demand in the RTM) now that the ODRC adder has been introduced. Mostly, the RUC is used for congestion management in problematic zones. It is not clear why the LMP signal (at the market price cap or floor) is insufficient by itself. Perhaps, analogous to the NEM’s RERT, generators committed in the RUC might be paid higher than the ERCOT market price cap.

EUROPEAN MARKETS

Introduction

Electricity markets in the countries belonging to the EU comes in a variety of shapes and sizes, but there is some underlying commonality, encouraged by EU regulations. The EU continues to push for more convergence, to allow efficient trading of electricity between EU nations. This section outlines the common features of these markets, without going into the detailed differences which are, in any case, under constant review and reform.

A typical design architecture for a European market is shown in figure A2, below.

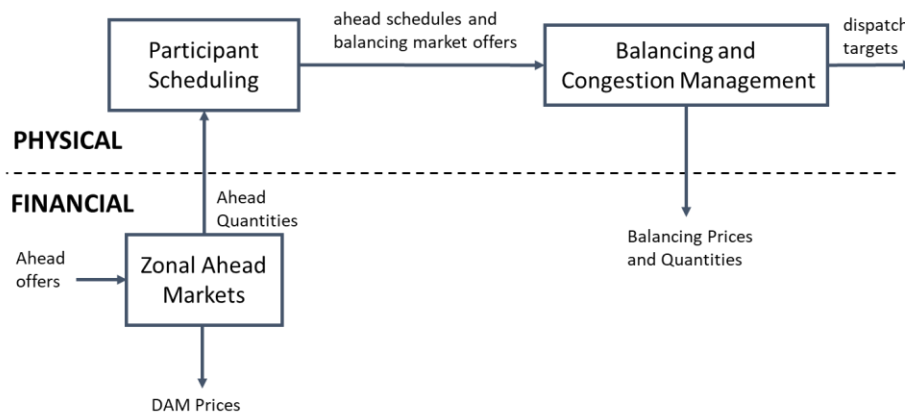


Figure A2: European Market Architecture

Ahead Markets

Ahead markets are the main trading platform for electricity. Typically, there will be a day-ahead auction, followed by staged or continuous within-day trading, up to a “gate closure” a few hours before real-time.

Since spot markets are thin or non-existent, the ahead products are essentially physical, rather than simply financial derivatives against the spot price as in the NEM. However, there is no explicit physical modelling of generating plant or transmission within the ahead market architecture. Electricity is simply considered as a commodity, with a single price in each geographic “zone”: which basically means a country. So bids and clearing engines are simple. There is no explicit or implicit scheduling of generation done by the clearing engine.

Ahead trading between zones is permitted, using a simplified representation of the actual transmission topology. Prices between zones will separate when there is congestion on these inter-zonal links, similar to regional pricing in the NEM. Unlike the NEM, though, interconnections between zones are looped, although looping may not be fully represented in the ahead clearing engine.

Balancing Market

Each generating company is responsible for scheduling its own plant within a zone to deliver its ahead sales in that zone. Of course, the aggregate of these schedules may be insecure, and transmission system operators (TSOs) in each zone are responsible for constraining generation on or off to manage congestion and ensure security. Rules around compensation differ, but typically a generator would be compensated at cost for being constrained. Generators generally pay transmission charges in return for obtaining this “firm” grid access.

The TSO is also responsible for managing imbalances between ahead purchases and physical load. To do this they call on balancing service providers (generation and demand response) to provide flexibility in various timescales, from FCAS upwards.

Balancing costs may be recovered through a combination of balancing prices, balancing spreads, and uplift on load. A balancing price is similar to a spot price, but payable only on differences between ahead trades and physical flows. It might be set by a weighted-average of balancing service providers costs (if there is “pay-at-bid” compensation for these) or a clearing price like in a spot market. The balancing spread is a difference between the balancing price paid to the “sell” side and the price paid by the “buy” side: of course the “sell” side may be generation or load, depending upon whether physical load is higher or lower than day-ahead volumes, respectively.

An “uplift” would be an additional levy payable by load but not by generation.

Discussion

The EU market architecture has some superficial similarities with the NEM in that there is zonal pricing. However, the EU approach is more coherent, in that the zonal price is established through a zonal market, whereas the NEM – as we know – has a nodal market clearing engine, but then selects the clearing price at a specified node (the RRN) to apply across the zone. So the NEM combines a US-style clearing engine with a EU-style pricing architecture.

EU ahead markets do not support generator scheduling in the sense that US markets do and that the ESB aspires to. They do, however, provide obligations on generators to generate the quantities that they have sold ahead, but with no incentives to do this in a way consistent with system security.

There is probably not much for Australia to learn from the EU markets in terms of ahead markets. Indeed, it is more likely that the EU could learn more from Australia: about how the latter has managed to successfully stitch together its separate State markets into a single, fully-integrated market.