

The storage imperative: Powering Australia's clean energy transition

Implications for the NEM

October 2024



Executive Summary

Australia's clean energy future hinges on a critical yet often misunderstood element: large-scale electricity storage. Without it, our ambitious national emissions reduction targets—43% by 2030, 82% renewable energy by 2030 and net zero emissions by 2050—remain out of reach. Electricity storage on a large scale is the perfect, and very timely, complement to intermittently available renewable energy generation. Storage allows the intertemporal shifting of energy from when available to when needed. Despite its pivotal role, the economics and operational dynamics of storage are poorly understood, creating a significant barrier to progress. Not understanding the economics of storage leads not only to poor operational performance of assets, it also induces insufficient incentives to invest.

This white paper addresses this crucial knowledge gap, first recognising the imperative for storage in Australia's energy transition and by studying the economics of storage. We present groundbreaking research that challenges conventional wisdom and offers critical insights for policymakers, the market operator, and industry stakeholders. Our work uncovers novel findings, the details of which one must pay attention to in order to best harness storage, induce more investment and, if needed, determine what intervention may be required.

IN THIS HIGH-STAKES TRANSFORMATION, STORAGE EMERGES AS BOTH THE KEY TO UNLOCKING OUR RENEWABLE POTENTIAL AND A COMPLEX HURDLE WE MUST OVERCOME.

1. NEW MARKET DYNAMICS:

Storage introduces dynamic trading strategies into the National Electricity Market (NEM), necessitating a rethink of current market rules and designs.

2. ENHANCED ISSUES IN COMPETITION POLICY:

We find storage faces a sharper price-quantity trade off than regular producers, with two consequences: (i) less intertemporal energy shifting than the capacity suggests, and (ii) greater incentives for collusion.

3. INVESTMENT CHALLENGES:

We point to the potential limits of arbitrage revenue, which raises questions about the sufficiency of the arbitrage trade to sustain the massive investment required in storage.

Based on these insights, we suggest a combination of policy recommendations and a research agenda that needs to be prosecuted to understand how to organise trade with storage:

- **No regret-policies** such as Locational Marginal Pricing (LMP), a day-ahead market (DAM) and the promotion of forward contracting. LMP delivers more precise price signals on the grid, more opportunities for storage to take advantage of price arbitrage opportunities, assists in solving grid congestion and in location choices;
- **Appropriate investment incentives:** Reassess and revise schemes like the Capacity Investment Scheme (CIS) to better align with the incentives, the operational realities of storage assets and investment needs. While the CIS aims to “de-risk” investments, it may unintentionally induce poor incentives for optimal asset use. Instead one wants to encourage full utilisation of storage assets to minimising financial burdens on taxpayers.
- **Financing Reliability:** Develop new approaches to ensure system reliability in a storagedominated grid, potentially relying on long-term contracting or new procurement mechanisms;
- **Research investment:** Fund further research to address future challenges such as the limits of arbitrage as a source of revenue to sustain storage investment, and understand market design without thermal generation.

Australia can lead the way in creating a robust, efficient, and sustainable energy future, but understanding storage is the key to its success. This work provides crucial insights for drafting effective energy policies, market design, and regulatory frameworks, ensuring we navigate the enormous challenge of our energy transition successfully.

The paper begins by setting the scene. From this, it is clear that more research is required to the benefit of market participants, the agencies overseeing the National Electricity Market (NEM) and consumers alike. We should not be surprised by this need for more, and better, information in this period of great upheaval. Why should NEM participants be expected to fully grasp the complexities of this rapidly evolving markets on the basis of back-of-the-envelope calculations?

The paper then reports results obtained by the Monash University team on the operation of one or more storage unit(s) in a market. From this work one can draw implications for market design or competition policy, and assess any need for intervention.

The last part offers some commentary on current policy initiatives and engages in some informed speculation. It comments on a key concern of the industry, namely, financing constraints, and on the Capacity Investment Scheme (CIS). Finally it asks whether the current business model based on energy arbitrage is adequate to deliver the storage capacity required for the NEM to transition away from renewables. It also offers the path to a remedy.



MONASH
University

MONASH
ENERGY
INSTITUTE



Acronyms

ACCC	Australian Competition and Consumer Commission
ARP	Advancing Renewable Program
AEMC	Australian Energy Market Commission
AEMO	Australian Energy Market Operator
BESS	Battery Energy Storage System
CIS	Capacity Investment Scheme
DAM	Day-ahead Market
GIH	Grid Innovation Hub
LMP	Locational Marginal Pricing
NEM	National Electricity Market
NEMDE	National Electricity Market Dispatch Engine
PSC	Power Systems Consultants
VRE	Variable Renewable Energy
LMP	Locational Marginal Pricing
NEM	National Electricity Market
NEMDE	National Electricity Market Dispatch Engine
PSC	Power Systems Consultants
VRE	Variable Renewable Energy

Contents

01	Background: storage and the energy transition	7
	1.1 The role of storage	8
	1.2 The quantum	9
02	Trading electricity over time	11
	2.1 Stochastic arbitrage	12
	2.2 Intra-day trading	13
	2.3 Implications	14
	2.3.1 Market design	15
	2.3.2 Competition policy and market surveillance	15
03	Policy comments and informed speculations	17
	3.1 Current policy	17
	3.1.1 The Capacity Investment Scheme	18
	3.1.2 Financing frictions	18
	3.1.3 BESS as standby reserves	18
	3.2 The limits of arbitrage	19
	3.2.1 Looming market failure?	19
	3.2.2 Storage and reliability	19
04	Summary and other works	21
	References	22

Acknowledgements

This white paper is authored by Guillaume Roger, Monash University.

This project received funding from ARENA as part of ARENA's Advancing Renewables Program under Grant number 2021/ARP016, from the Grid Innovation Hub at Monash University and from Monash Business School's Energy Lab.

This white paper is based on joint work with Sergei Balakin at Monash University.

The original works can be found on the Knowledge Sharing webpage of this project:

<https://www.monash.edu/energy-institute/grid-innovation-hub/home/integrating-energy-storage-into-the-nem>



“

IF THE COST OF INVESTMENT
IN STORAGE IS STAGGERING,
THE SPEED AT WHICH IT MUST
DEPLOYED IS EVEN MORE
DAUNTING.

”

01

Background: storage and the energy transition

In Australia the electricity industry is transitioning away from fossil fuel faster than anticipated and is outpacing the speed at which the institutions governing the market have been adjusting.

Nowhere is this disconnect between institutions and the physical reality more acute than when it comes to electricity storage. At the risk of stating the obvious, storage is critical to enable the energy transition on a large scale. Hence, at this point, storage is the bottleneck to that transition. In addition, the scope of applications of storage is widening dramatically, from operating mostly in the (small) Frequency Control Ancillary Services (FCAS) market or as a power reserve, to energy price arbitrage, assisting in managing congestion on the grid or delivering “synthetic” inertia to support the grid.

Yet we know very little of the *economics* of storage, that is, of the incentives that shape decision-making when it comes to storage. This dearth of knowledge leaves policy-making bodies like the Australian Energy Market Commission (AEMC) and the Australian Energy Market Operator (AEMO) in a conceptual and technical vacuum to develop the policies necessary to integrate storage in the NEM. For market participants, this institutional uncertainty adds to the standard business risks of a new venture. It also means that pricing the services offered by storage operators (whether buying or selling them), is still very challenging. Therefore *valuing* storage remains equally taxing, which renders investment difficult to commit to.

1.1 The role of storage

It bears recalling that storage enables the intertemporal shift of electricity production to make energy accessible when required rather than when available. The need for this intertemporal shift arises from the emergence of VRE that is not controllable, and is made evident in Figure 1.1. The left panel shows a 7-day average of total VRE supply in the NEM in 30-minute increments. The right-hand side depicts the corresponding *gross* demand information.¹ Storage is the device that can transfer the large, midday production of energy to the higher-demand period. In better matching supply and demand over time, storage is also expected to smooth prices. We delve into these details shortly.

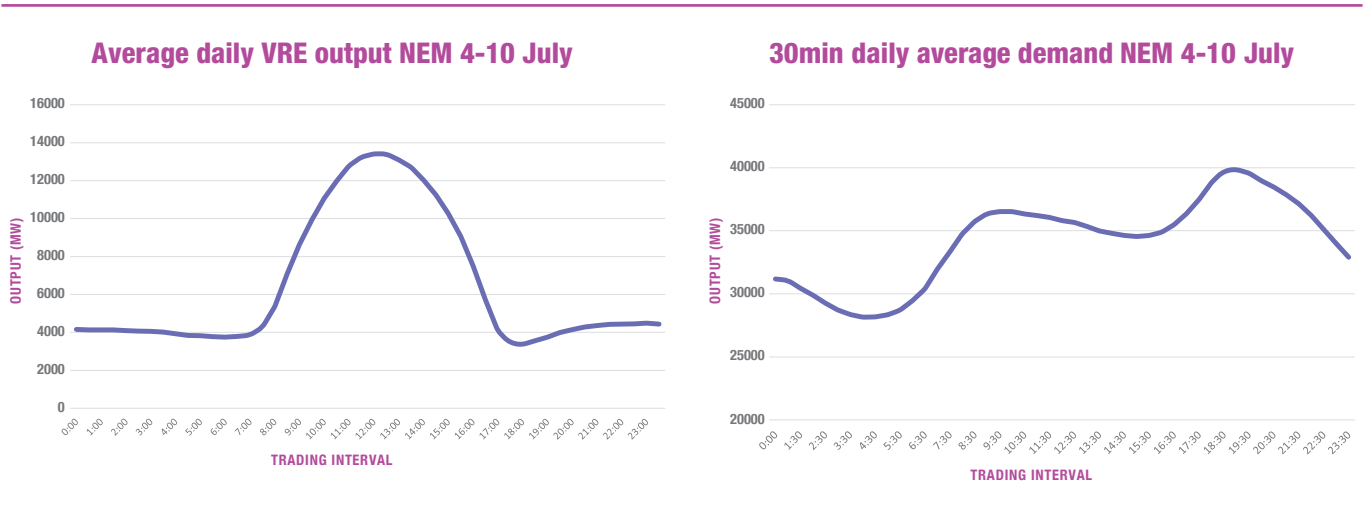


Figure 1.1: Mismatch between VRE supply and energy demand. Daily 30-min averages, 4-7 July 2024.

Source: AEMO data.

By comparison, the largest BESS to be built (not yet commissioned as of writing) offers only **650 MWh of storage today.**

Moreover, the cost of this unit is estimated to be \$400 million according to Origin Energy, which translates to **\$615,000 per MWh.**

Therefore, if relying on the same technology (lithium ion), the total cost of investment to **deliver 258 GWh is in the order of \$158 billion.**

When thinking about new transmission investment, storage ought to be included in the mix to determine the optimal transmission capacity choices, as well as the optimal location of storage capacity. These depend on demand and supply patterns on the grid, of course. Storage can also be used to supply grid support in the form of synthetic inertia or fast response services.

Unfortunately, storage can be used for less lofty purposes. It can be used to enhance existing market power, it renders tacit collusion *easier* than it is now and it can be used to manipulate the market.

¹ This Figure uses winter data and so it is immune from the larger demand variations of the summer.

1.2 The quantum

The current knowledge vacuum is all the more concerning that the necessary investment in storage to complete the energy transition is staggering. To get a sense of the quantum required, consider the following exercise.

As of writing, the dispatchable capacity of the NEM is approximately 43 GW; this power rating capacity is available at almost any time, modulo infrequent outages. To crudely demonstrate the scale of Australia's aspired clean energy transition, effecting a 50% transition requires a minimum of 21.5 GW of power capacity to be delivered by storage. Having this power available for 12 hours (overnight, roughly) therefore requires 258 GWh of *energy* capacity. By comparison, the largest BESS to be built (not yet commissioned as of writing) offers only 650 MWh of storage today. Moreover, the cost of this unit is estimated to be \$400 million according to Origin Energy, which translates to \$615,000 per MWh. Therefore, if relying on the same technology (lithium ion), the total cost of investment to deliver 258 GWh is in the order of \$158 billion. These are large, but realistic numbers: on Figure 1.1, which is typical, between 5pm and 7am, VRE produces almost nothing while demand never falls below 25 GW (and reaches about 40 GW when VRE can only deliver 10% of it). Such an expense deserves some study. It is also hard to come by, especially in such a short amount of time.

If the cost of investment in storage is staggering, the speed at which it must be deployed is even more daunting. California recently celebrated a milestone of 10 GW of storage (power) capacity in April 2024, on its way to a total installed capacity of 52 GW by 2045.² While impressive and laudable, these 52 GW are still a long way from the almost 83 GW of currently available dispatchable (power) capacity in the California market. Australia's targets are even more ambitious: it seeks to generate 82% of its electricity from VRE sources by 2035. This requires installing a total of 36.5 GW of storage capacity, or approximately 15 times what is currently available, in the next 10 years.

That is the strong sense in which storage is the true bottleneck of the energy transition. Even an ambitious target like that of California, which requires hundreds of billions of dollars, falls short of really effecting the energy transition in a reasonable time. The task is tremendous, the investment eye-watering in cost, the technology uncertain, and the policy framework to achieve all of this is not even defined.



² Source: www.energy-storage.news/california-energy-storage-revolution-is-here-says-governor-as-us-leaderstate-surpasses-10gw



“

TRADING THROUGH STORAGE IS DYNAMIC; IT TAKES PLACE OVER TIME. THEREFORE ACTIONS TAKEN NOW HAVE IMPLICATIONS FOR THE FUTURE, WITH THE COROLLARY THAT ACTIONS TAKEN IN THE PAST CONSTRAIN CURRENT ACTIONS.

”

02

Trading electricity over time

The basic premise underpinning the trading of electricity using storage is that of *intertemporal arbitrage*, which is popularly summarised as “buy low and sell high”.

For economists, the mere existence of arbitrage opportunities is regarded as an anomaly; arbitrage opportunities are symptomatic of missed trading opportunities.

But in electricity, things are a bit different and arbitrage opportunities can even be predictable. The reason of course is anticipated variations in demand and supply over the course of the day, as exemplified on Figure 1.1, which induce a quantity (and price) path that can be anticipated—that is, expected in a statistical sense. Then storage can be used to smooth these intertemporal variations: soak up energy when cheap and available, and release it when supply is scarce and prices are high. In fact, the whole energy transition rests on this intertemporal substitution, which is why storage is completely critical to transforming the power system. Throughout, this activity is referred to as “intra-day trading”. Other imbalances between supply and demand are known to arise almost systematically, but are *uncertain* and therefore not predictable. Nonetheless, they do give rise to trading opportunities – for example, when a transmission line fails or is saturated, when a generator trips or when a cloud passes over solar farm. These are “stochastic arbitrage” opportunities; they are uncertain events. Table 2.1 illustrates the point. Of course, in the real world, both stochastic arbitrage and intra-day trading may co-exist.

Time	Queensland (\$/MWh)	NSW (\$/MWh)
6.05pm	15,500	12,465
6.10pm	15,500	14,218
6.15pm	3,569	3,164
6.20pm	15,500	14,120
6.25pm	15,500	14,507
6.30pm	399	358
6.35pm	15,500	13,519
6.40pm	355	308
6.45pm	9,999	9,026
6.50pm	9,797	8,738
6.55pm	370	323
7.00pm	304	265

Figure 2.1: Prices in NSW and QLD on 16 March 2014.
Source: AEMO data.

To understand the intricacies of the seemingly simple exercise of “buying low and selling high”, it is important to bear in mind two facts that set storage apart from standard trading. First, trading through storage is dynamic; it takes place *over time*. Therefore actions taken now have implications for the future, with the corollary that actions taken in the past constrain current actions. For example, a unit that is fully charged can no longer absorb energy. Furthermore, a forward-looking, rational operator internalises this reality, which affects their precise actions in the market. Second, charging is costly while discharging generates the revenue opportunities. As a result a storage unit only ever buys what it must to maximise its revenue—modulo operating constraints such as temperature or degradation. Storage does not buy because it can, except in the face of negative prices, only because it must.

A third, important fact that is not unique to storage is that arbitrage opportunities disappear as soon as they are exercised. This is well-known to securities traders, for example, who are concerned about the so-called “price impact”. Here the price impact kicks in when selling but also when buying, so that the price difference – called the *arbitrage spread*—contracts rapidly. Furthermore, this contraction of the arbitrage spread points to the *limits of arbitrage*. That is, as the arbitrage spread contracts, the revenues to storage may decrease to the point of being insufficient to support investment in storage capacity. In economic jargon, there arise the possibility of a divergence between the *social value* of storage and its *private value*; this is the definition of a *market failure*. If this eventuality materialises, then intervention is warranted and another mechanism than arbitrage may be required.

2.1 Stochastic arbitrage

This form of arbitrage is the exercise by which a storage unit buys when the demand unexpectedly drops and sells when it unexpectedly increases. To capture this idea, we suppose demand is affected by shocks that are either negative or positive, for example with probability 1/2 each, and have value zero on average. While stylised, this set up does embody two essential features: shocks (to either demand or supply) cannot be forecast and there is no systematic trend. Again, table 2.1 illustrates such a situation. The goal is to take advantage of these variations in demand (or supply) that induce variations in prices; the difficulty is the uncertainty of these variations. In our work at Monash, we find that the optimal strategy of an operator bears two important characteristics.

First, the storage operator faces very strong incentives to withhold quantities because the arbitrage spread decreases rapidly in the quantities traded. These incentives to withhold quantities are *stronger* than those faced by a standard seller with market power. The reason is that storage must also *buy*; with market power, buying increases the purchase price in a first-order sense. The consequence then is that the arbitrage spread is eroded both ways: when selling and also when buying. To mitigate this, the storage unit limits the quantities it sells and, therefore, that it buys. This is the well-known trade-off between the extensive margin (quantity) and the intensive margin (or infra-marginal losses) that any monopolist faces, but now more acute.

The second aspect is a new phenomenon that emerges as a corollary to market power. A storage operator may find itself buying repeatedly in the face of a string of negative shocks. For a large unit, these repeated purchases can be very costly, especially because they require an equally long string of positive shocks to be sold off. Balakin and Roger (2023) [3] call this the *continuation risk*. To manage the continuation risk, a storage operator trades their capacity in multiple steps. The consequence of the continuation risk is that capacity should not be too large compared to the size of the shocks. When shocks also vary in size, as is the case in the NEM, this implies the storage unit should not be too large compared to the average shock.

Therefore, in this uncertain world, a storage unit with market power never trades its capacity in full (even if it technically can, i.e. absent any constraints). This outcome is the result of the combination of market power with the continuation risk, which induce a prudent and conservative behaviour. The consequence of this restrained trading is less price smoothing; storage does not shift as much energy from the high-production period to the high-demand period as its capacity allows.

These behavioural characteristics are not an artefact: the research team at Monash considers different kinds of random shocks—with more or less persistence and more or less symmetry to the same effect (see [3]). The best environment for a storage unit is a sequence of shocks with low persistence; that is, a sequence such as $-$, $+$, $-+$, $-$, ... which allows the unit to buy and sell in sequence almost surely. In other words, it is the variation in demand (or supply) that is the source of revenue; the larger that variation, and the more certain it is, the better off is the storage operator. The *uncertainty* around that variation is costly: if, after charging, the storage operator is uncertain whether it can discharge, it may prefer hedging its bets by charging less in the first place and keeping some spare capacity. The converse holds when selling. Making this distinction between *volatility* and *uncertainty* is important when considering whether to enter this market as participant, or contemplating any kind of intervention as a regulator or government. Any intervention should preserve volatility.

The research team studies an adjacent problem, in which a small number of storage units compete in an oligopolistic market just like the NEM—see [5]. Here too the sensitivity of the arbitrage spread to quantities makes for strong incentives to restrain quantities traded. In a competitive environment, this restraining of quantities is implemented through *collusion*, that is, the tacit understanding between two or more parties they should cooperate rather than compete head on. Collusion can take many forms: unilateral quantity restrictions, or taking turns when buying or selling, or both. Indeed, the fact that storage must both buy and sell expands the scope for collusion; collusion can now take place only when buying, when selling, or both. These incentives are so strong that when storage units are large enough, some form of collusion is the *only* equilibrium; that is, it is the only course of action guarantee positive payoffs. What is perhaps puzzling is that this collusion can be welfare improving, that is, socially desirable. This result runs completely counter to decades

of research in the economics of antitrust and of practice of antitrust law, where collusion is uniformly *undesirable* (because excessively costly to consumers). Here, collusion can be helpful to consumers because it only takes place when the aggregate storage capacity is sufficiently large. But then, the arbitrage spread is small; consumers benefit from these constrained prices and large traded quantities.

Needless to say it is still a bit unsettling to arrive at this conclusion. It is also not clear what a remedy is, nor that a remedy is necessary. At the very least, it is difficult to make a general claim: it may depend on the specifics of the case. Indeed, consider a series of independent investment decisions that result in a large capacity installed (as was seen in the late 1990s with fibre optic cable in the United States, for example). If collusion is the *only* equilibrium, then what can and should a competition authority like the AER or the ACCC do? This is not a trivial question: first, as shown by the Monash research in [5], collusion can benefit consumers; second, when a competitive equilibrium also exists, there is scope for the authorities to put in place remedies to revert to that equilibrium. It is not clear what the course of action should be when no such equilibrium exists.

In concluding this section, we attract attention to four points that are relevant to market participants and regulators:

1. Variations in demand (equivalently, supply) are the source of all income.
2. The uncertainty of these demand variations is the source of costs.
3. In the face of uncertainty, a storage operator never trades its quantity in full. Therefore there is less intertemporal shifting of energy and less price smoothing.
4. Collusion may emerge as the only, and the desirable, equilibrium, which raises the question of regulatory response.

2.2 Intra-day trading

Intra-day trading is perhaps what most practitioners have in mind when it comes to electricity storage. It is the activity whereby one or more storage unit purchases energy when it is plentiful and sells it in times of higher demand; it is the bulk shifting of energy over time. Unlike stochastic arbitrage, it rests on statistically *predictable* variations in demand, and therefore in prices, over the course of a day. This is taking advantage of the famous “duck curve” shown in Figure 2.2 for the South West Interconnected System (SWIS) of Western Australia; this could be any Australian state, or California or other markets.

Even when *predictable*, price variations (induced by variations in either demand or supply) are not always *certain*. Combined with market power, this uncertainty renders the exercise of optimal intra-day trading quite challenging as well. In particular, do shocks in a given trading day carry over into the future? Does a storage operator find it best to speculate about possible large variations in demand, or does it stick to trading the average?




Figure 2.2: The duck curve becomes increasingly pronounced.

Source: Synergy.net.au using AEMO data.

In a third paper [4], the Monash research team considers the problem of intra-day trading when the demand for energy remains uncertain at the time of charging, but more information may be available when selling. This fits a typical day for a storage operator: it charges in the middle of the day, and as the day unfolds, information about the evening peak becomes increasingly accurate.

Strategies are determined in terms of *power rates* and *duration*, which endows the operator with considerable flexibility and points to the need for accurate load and supply forecast. Indeed, *exactly when* to trade and with which *intensity* are considerably important, and completely determined by net demand.




Figure 2.3: Cyclical demand with uncertain peak.

Figure 2.3 offers a depiction of the market demand studied by the research team in [4]; its fluctuations are modelled as a periodic function of time $(\theta - (1 + \xi) \sin t)$ that is augmented by a shock ξ that ranges from -1 to $+1$, and is on average 0. Again, this is stylised; the point is that average demand can be determined using statistics, but actual demand remains uncertain. In particular here,

the demand peak is uncertain. This is a first-order question for all parties in the NEM: market participants, AEMO and traders. In the model, the magnitude of the shock is not essential; what is important is *when* this information is revealed to the storage operator.³

When the operator never learns the true demand, it must work from averages that may be estimated from historical information. Given a 0-shock (on average), the operator determines its best selling strategy $s(t)$ that is a function of time; that strategy is determined by a discharge rate that varies over time starting from some time $\pi + t_0$ to $2\pi - t_0$. Concretely, this means a storage unit need not sell (nor buy) at full tilt all the time. Indeed, the selling strategy follows the (average) demand path: it starts small and reaches its maximal discharge rate at the peak demand, and then decreases the discharge rate. The reason is that the storage operator manages its price impact. It does not simply sell at full rate when starting for the demand is still low then; nor does it delay until later when demand is higher to start selling because then it must sell too much too fast, which depresses prices. That sophisticated selling strategy is the best to mitigate the price impact. Because the shock has zero expected value, the buying path is the mirror equivalent. It is a function $b(t)$ of time starting from some threshold t_0 and finishing a time $\pi - t_0$. This is shown in Figure 2.4.

The same logic applies when the operator can avail a more accurate forecast. The information is more reliable and the shock expected by the storage operator may differ from zero, but still uncertain. Given the new information, the operator determines its best selling strategy $s(t)$, which may start at a different threshold $\pi + t_1$. Again, this determines a quantity to be purchased; the best buying strategy then minimises the procurement cost of the operator.




Figure 2.4: Selling strategy under uncertainty for different capacity levels (power rates).

If the operator learns exactly the true demand, the same approach broadly applies, but not exactly. Once the shock (ξ) is known, the operator selects the best selling strategy given ξ ; call this $s(t, \xi)$. That is, the operator can tailor its selling strategy exactly to the demand shock, not just to average demand. In particular, *when* to start (and stop) selling adjusts according to the shock ξ ; there is a threshold

time t_ξ ; for each shock realisation.⁴ In this case, more can be sold *on average*; the reason is that when selling, the storage unit does so in better conditions, so at higher prices, than when it must rely on *average* demand only.⁵ There exists a whole family of these strategies—one for each ξ . Given this, *before* the demand shock materialises, the operator can then compute how much to purchase on average. The best buying strategy is determined as the one that minimises the procurement cost of this quantity. This is shown in Figure 2.5.




Figure 2.5: Selling strategy under full information for different shocks ξ (power rates).

When the shock ξ is large, the power output peaks high and the selling window is narrow; the converse holds when the shock is small. Below a threshold shock, the stopping time $t_\xi = 0$: the unit starts selling as soon as possible. It does so with less variation in intensity over the selling interval. This is the optimal strategy to best manage its revenue, given its market power, over time.

A market participant thus should take note of three points:

1. A storage unit optimally tailors its (dis)charge rate to demand conditions; it does not systematically (dis)charge at full power capacity, nor does it wait until the demand approaches its peak. This logic applies regardless of the number of storage units participating.
2. Consequently, the (demand) peak-shaving is limited. It is not the case that the whole energy capacity is delivered at the highest demand and the highest price. Therefore prices also do not smooth out completely.
3. Forecast information is critical to adjust the best selling strategy $s(t)$. The same logic applies to the buying strategy $b(t)$.⁶

2.3 Some policy implications

In light of these works we draw some implications for market design, competition policy and the role of government intervention. Most of what is suggested below can be implemented almost immediately at very little cost, and relies on well-tried instruments.

³ Even with increasingly better forecasts, some uncertainty remains until the exact time when the storage faces actual demand.

⁴ Strictly speaking, a stopping time t_ξ .

⁵ Using this capacity in the arbitrage market would depress the arbitrage revenue; see the discussion above.

2.3.1 Market design

Competition through forward markets. The first comment to make is conceptual in nature. Whether engaging in stochastic arbitrage or intra-day trading, storage operators employ *dynamic* strategies. In other words, they unfold over time. In response, it is reasonable to speculate that a market operator should also be far-sighted and rely on its own dynamic strategy. At present, the NEM is a spot market supported by a limited amount of mostly bilateral forward contracting. It may be *time* to restore the time dimension in the NEM, which may include determining what is an acceptable bid, and how bids are ranked and cleared. Work on this specific topic is still forthcoming; it is technically very challenging.

More concretely, everything so far points to storage operators having very strong incentives to eschew competition—stronger incentives than standard generators, for example. The reason is that, to sell they must also buy. If they have market power, they are best off restraining quantities when selling and also when buying. In addition, uncertainty further contributes to restraining quantities out of a precautionary motive.

Consequently, any alteration to market design should seek to promote competition even more urgently than ever. One well-established avenue to do so is to introduce mandated forward markets; one such form of forward markets is a day-ahead market. The twin advantage of forward markets is that they are pro-competitive (Allaz and Vila, 1993) [1] and they neutralise the uncertainty that plagues the models discussed in Sections 2.1 and 2.2. That uncertainty remains in the real-time market, where it invariably must be resolved, but then the quantities involved are much smaller. Therefore the welfare consequences are correspondingly less dramatic. In Singapore, the Energy Market Authority (EMA) uses vesting contracts. Large generators must participate in vesting contracts. Currently, the vesting price is determined by the EMA based on the long-run marginal cost of the most efficient technology. *Long-term* forward contracts promote competition in that they neutralise the price effect that we discuss at length in Sections 2.1 and 2.2. Since the price of energy is fixed by contract, all that the operator can do is to manage quantities; a well-designed contract can lead the operators to operate at marginal cost. Mandating, or at least promoting, *long horizon* forward contracting also helps in financing the construction of new capacity (see Section 3.1.1).

A well-known form of forward market is a day-ahead market (DAM), which is in operation in all markets in the US. In a DAM, sellers and buyers *commit* to a schedule for delivery the next day. Introducing a DAM necessitates jumping another hurdle: a DAM requires first the introduction of Locational Marginal Pricing (LMP, or nodal pricing).

Locational Marginal Pricing. LMP is well known and has been in operation in all US markets for over a decade. Not only is it the efficient solution, it is even more so when it comes to storage: as we now know, price arbitrage relies on there being price variations to exploit. With LMP, there are more price variations at any given time because there are more prices in the market (one for each node). Recall Table 2.1, where there is only one price for each state; LMP allows for tens or hundreds of prices. There are also more price variations at any given

node as transient congestion ebbs and flows, with prices responding sharply to that congestion. Hence LMP is an essential supporting ingredient of the successful roll-out of storage capacity. Furthermore, well-located storage units at potentially congested nodes of the transmission network can alleviate that congestion, assist in better managing the network and economise on transmission investment. The introduction of LMP prior to the introduction of a DAM is required to stifle the so-called INC-DEC game. This practice is another form of intertemporal arbitrage, whereby a generator located on a constrained transmission line faces a higher price in the DAM than can be expected in the real-time market. The arbitrage is simple: sell as much as possible in the DAM, and buy back at a lower price what cannot be dispatched in real time. This problem does not arise with LMP because the nodal price in the DAM already internalises the transmission constraint.

In the NEM, almost all the ingredients required to implement LMP already exist. Today, the dispatch engine (NEMDE) computes all the relevant shadow prices as dual variables that have an immediate interpretation as prices. It also computes the appropriate Marginal Loss Factors (MLF) to settle the zonal reference price net of these MLF.

2.3.2 Competition policy and market surveillance

Storage adds to already existing concerns of competition policy. Awareness of the issues of competition policy amplified by storage may not accelerate the roll-out of storage capacity, but it does protect consumers.

Andres-Cerezo and Fabra (2023) [2] cogently make the case that a conventional generator should not own a storage unit because it can coordinate on selling, and therefore on restraining its *joint* quantities. Of course this is magnified when a firm owns multiple generators and multiple generation units—such as gentailers in Australia. To these comments one can also add that bidding and re-bidding are even easier with a storage units, which can start and stop producing at will.

An even more concerning risk is that of outright market manipulation. For a firm that owns multiple generation units, it is easy to install a small-scale storage unit that can *add* to aggregate *demand* at critical times by purchasing energy. In doing so, it pays a high price but only on small quantities, and it contributes to increasing the clearing price for *all* infra-marginal quantities. This is possible because of a convex supply function, which is typical of electricity markets. This form of market manipulation is novel, possibly not yet observed, but most definitely feasible.

In sum, storage renders existing problems of competition policy more acute, and adds new problems. It requires vigilance on the part of market surveillance authorities like the AER, which may need to acquire a new mandate and new skills, and possible intervention(s) from the ACCC. One remedy is to (a) prevent generators from owning storage units to promote new entry in this nascent market and (b) to prevent a firm from owning multiple storage units to ensure the market remains competitive. It also require greater vigilance on the part of the regulator (the AER) in its market surveillance operations. Concretely, it means greater oversight of the market and investment in skilled personnel to detect these anti-competitive practices from the data.

⁶ Here it is simple and symmetric to selling. But it need not be so in the NEM, and it is subject to uncertainty too.



“

WHAT IS SOCIALLY OPTIMAL
MAY NOT BE DELIVERED BY THE
MARKET BECAUSE THERE IS
A DIFFERENCE BETWEEN THE
SOCIAL VALUE OF STORAGE,
AND ITS PRIVATE VALUE.

03

Policy comments and informed speculations

The research we lay out in Section 2 allows us to engage in some informed comment about a couple of current policy items, and to contemplate whether the arbitrage model Australia relies on, is sufficient to fulfill the ambitions of the energy transition. A pressing issue is the rapid deployment of more storage capacity in order to increase the value of renewable energy assets, especially solar farms.

3.1 Current policy

The current policy landscape remains somewhat underdeveloped and, at times, appears fragmented; this may reflect, in part, the highly decentralised nature of the governance of the NEM. At the Federal level, the government recently introduced its Capacity Investment Scheme (CIS), which is a form of income guarantee to storage operators that is modelled after a similar scheme in NSW. It is not clear that it addresses the financing constraints that seem to plague this nascent industry. At the state level, NSW offers its own investment scheme and Victoria has entered into some bilateral contracts to support battery investment

(the Big Battery, for example). While likely well-intentioned, these schemes do not fully align with the operational realities of storage or the incentives facing market participants, requiring continuous review in this rapidly evolving sector.

3.1.1 The Capacity Investment Scheme

The details on the CIS are scarce. It is best described as a revenue guarantee to the operator(s) of a storage unit, which the federal government presents as a “collar”. If the revenue falls below a floor, the government guarantees that floor. Above a ceiling, a fraction of revenue must be rebated to the government. The levels of the floor and the ceiling, as well as the rebate rate (from 0 to 100%) are *bidding variables* (among others) in a procurement auction. The other details are not clear, including how the bid evaluation process trades off the many variables of a bid—for example, higher floor versus large rebate rate. In addition, the CIS requires a fraction of capacity to remain on standby in periods 10 of stress; in that time the operator cannot generate any revenue. Taken together, this scheme induces distorted incentives, and so makes for a very expensive guarantee that taxpayers provide. The standby requirement is also quite a deterrent to operators.

Offering a floor on revenue is an insurance against failure, but inadvertently reduces incentives for adequate management by the operator. The ceiling and share rebate in turn weaken the incentives to seek additional revenue; that is, it deters from actually using the asset to the full extent the technology allows. Indeed, if sharing the revenue, there is little incentive to chase the extra dollar, especially when that is costly (degradation) or difficult (trading strategies). Finally, setting aside a large fraction of capacity for unspecified contingencies in times of system stress may appear intuitively sound and appealing but makes little sense: it is precisely at that time that storage should be free to operate, and should be expected to do so to generate its revenue. In doing so, it sells when energy is the most valuable—and makes the most money, which supports the investment in the first place. In addition, by being active in times of stress, storage can in fact relieve that stress. So, why the CIS and what does it seek to accomplish? According to publicly available information, the CIS “derisks” investment. We presume this means easing the financing constraints. But in so doing it curtails revenue opportunities and weakens incentives so much that it runs the risk of undermining the intended outcomes. There are simpler, less onerous ways to relieve financing constraints to support investment in storage.

3.1.2 Financing frictions

Financing constraints seem to be a chief concern of developers of storage facilities. In fact, standalone developers face an enormous array of risks: financing, construction and delivery, and finally operational.

In the face of these risks, financiers—typically, project financiers—request “bankable revenue”, that is, a guaranteed revenue stream from the operating the unit. Typically, a (private) credit-rated counter-party steps in to deliver that guaranteed revenue in the form of a “tolling agreement”. This is effectively a rental contract (or a swap) whereby the counter-party rents the facility for a pre-determined price and takes over operational control. Today, the counter-party to most tolling agreements is a large retailer (a.k.a. one of three gentailers).

The motivations of all these parties, including that of financiers, are completely understandable. The construction risk looms large in any infrastructure transaction, and especially so in the NEM, where connections are subject to quite systematic delays (some of which are also understandable). The “tolling” practice resembles that

supporting the development of standard infrastructure like roads, rail, bridges or tunnels. But this presents some challenges. First, it leaves developers carrying most of the risk, namely, the construction risk. Second, already extremely concentrated gentailers get privileged access to storage assets without committing their balance sheet. This runs completely counter to the promotion of competition in the market for storage, and for the general tolerance of these large companies. Instead, the CIS may be that guarantee: a project that is supported by the CIS can deliver a guaranteed revenue.

However, a more effective approach can be considered. The CIS could potentially be bypassed by a combination of long-term forward contracts between private parties that require no, or at worst limited, support from taxpayers and have superior incentive properties. A developer of a storage facility can enter into a forward contract with a VRE generator to buy, and a retailer to deliver, energy at a prespecified date at agreed prices that are implemented by contracts for differences. From the work in Section 2.2, we know how to determine the optimal price paths when buying and selling. This can form the basis of an agreement that best takes advantage of demand fluctuations. These contracts deliver a bankable revenue that a project-financier can rely on. A developer that fails to build in time to meet its obligations owes its counter-parties penalties that are a function of realised market prices. This delivers strong incentives for delivery. The government can offer a guarantee to augment the credit-worthiness of these private parties, so that smaller generators and retailers can participate in this market. The great benefits of this solution are (a) to not interfere with the incentives of the storage operator after delivery; (b) to promote competitive behaviour in storage through adequate (long-term) contract design and (c) to not deliver expensive storage facilities at cut-rate to an already concentrated (generation) sector. As mentioned previously, concentration in storage can be very harmful. In response, one should strive to induce new entry in storage operation, rather than turning storage operations to large incumbents.

3.1.3 BESS as standby reserves

In some instances, Battery Energy Storage Systems (BESS) are used as standby reserves; it is that kind of agreement with each of the state that supports the on-going operation of the Hornsdale Power Reserve (SA) and the Big Battery (Victoria).

At present, such agreements appear to be difficult to justify, especially given their high cost to taxpayers. First, there is no shortage of the supply in the NEM, as was tested and proven in June 2022 when AEMO took control of the NEM and delivered energy to all Australians.⁷ This episode also proved there are plenty of incentive problems in the NEM. But there is no need for standby reserves. Second, who decides when these reserves should be used? Under regular market operations, the storage operators should make these decisions given prices, which typically reflect the conditions of the system and the value of energy to consumers. Third, this approach fails to take advantage of the main feature of storage, which is to charge and discharge to shift energy over time to the benefit of consumers.

In time there may be a need to build energy reserves to meet unforeseen contingencies. That time has not yet come and BESS is certainly not the right technology for this kind of event.

⁷ In June of 2022, there may have been a shortage of supply at the fixed price of \$300/MW, but not because there was no capacity—as the intervention of AEMO demonstrates.

3.2 The limits of arbitrage

Aside from ancillary markets that are small and standby reserves that are never traded, revenue generation rests on energy arbitrage as discussed in Section 2. The problem with arbitrage is that as soon as it is entered into, it tends to vanish: selling prices decrease, buying prices increase and revenue contracts. Here we draw some speculative implications of this reality.

3.2.1 Looming market failure?

As arbitrage opportunities are exercised, the arbitrage spreads shrinks—after all, that is the point of it and there should be no arbitrage opportunities in a well-functioning market. In the competitive limit, the spread contracts so much as to only match the marginal cost of supplying the service; here one can think of the round-trip efficiency losses and the degradation cost of storage unit. Such a small revenue fails to cover the fixed costs of the investment in storage. In a competitive market, this is a distinct possibility, as is the case for airlines, for example.

There is nothing wrong with arbitrage spreads decreasing; they are meant to, and are certainly not symptomatic of a market failure. The question is whether they contract so fast as to choke the investment in storage. More precisely, can the energy arbitrage revenue stream sustain enough investment in storage to support the *whole* transition (or close to)? This is a question of an empirical nature that we are unable to answer precisely at this point, but that is conceptually and technically within reach. It is a question well worth considering now, even if it cannot be answered in full. Simply put, can we get there? If not, what to do?

Such a situation *is* a market failure: what is socially optimal may not be delivered by the market because there is a difference between the *social* value of storage, and its *private* value. The private value of storage is easy to grasp: it is the revenue generated by the arbitrage spread, net of operating costs. The social value is a little more complicated: to the private value we must add the benefits to consumers—lower prices in times of high demand—and the benefits to generators—higher prices in times of low demand.

That is, the arbitrageurs (the storage units here) generate *positive externalities* for which they are not compensated. If they are not compensated, there is no investment in more capacity. Given how critical storage is, such a situation would stall the energy transition.

Then a new approach altogether is required, for that social value is real and can be shared. But trading energy alone may not allocate this social value well to all interested parties. There are ways to radically improve on this (so far, speculative) poor outcome. Broadly speaking, a typical scheme consists in taxing and redistributing. A market operator can tax transactions and offer income support to potential entrants. How much to tax can be determined using a Vickery-Clarke-Groves (VCG) mechanism, which has been a staple of incentive design since the 1980s.⁸ Here the VCG mechanism can be amended to account for the revenue the storage units can still generate from their own operations.

This is a question the Monash team is exploring in 2025. To first evaluate whether arbitrage revenue is insufficient, we must first simulate the market impact of the progressive retirement of thermal resources. If needed, we can then design a more appropriate mechanism to spur investment in storage capacity.

3.2.2 Storage and reliability

Reliability raises a similar concern. In a world of (mostly) renewable energy supported by storage, reliability as defined by AEMO today requires having storage capacity installed, and possibly on standby, to operate only a few hours of the year.⁹ This cannot be supported by energy arbitrage, as it is redundant (excess) capacity most of the time.¹⁰ Yet this capacity is deemed important and valuable in a modern power system.

What is an appropriate remedy is not clear yet. It may be that VCG subsidies (discussed in Section 3.2.1) are sufficient, or even that the price cap is sufficiently high to support that investment in capacity. If not, one can contemplate an alternative out-of-market arrangement, whereby capacity is built and remunerated at a pre-determined rate by taxpayers, and used only when all other capacity has been exhausted. The revenue it earns then contributes to relieving the taxpayers' contributions.



⁸ It consists in asking prospective taxpayers how much they are willing to pay for the service. To ensure they reveal this information truthfully, their tax rate does not depend on their information but on the information of other taxpayers. This scheme is called "incentive compatible": everyone tells the truth.

⁹ Note that the NEM is far from fitting this description today; hence power reserve are truly redundant today.

¹⁰ Using this capacity in the arbitrage market would depress the arbitrage revenue; see the discussion above.



“

MARKET POWER IS A CHIEF CONCERN OF A FORWARD-LOOKING STORAGE OPERATOR, WHETHER A LONE OPERATOR OR COMPETING UNITS.

04

Summary and other works

Storage is the critical bottleneck of any meaningful progress in the energy transition. Today a new solar farm has very little value because it produces at the wrong time of the day; the energy it generates needs to be made available later, at which time it becomes valuable. However we understand little of the economics of storage; as a result, investment stalls.

In this white paper we present some novel results on the management of storage units in environments characterised by a long horizon and uncertainty. We find that market power is a chief concern of a forward-looking storage operator, whether a lone operator or competing units. It is particularly acute because a storage operator must also buy. As a result, there is less shifting of energy from the production period to the consumption period.

From this work one can draw implications in terms of market design and competition policy. We point to the new nature of the NEM when suppliers use dynamic trading strategies, to enhanced concerns of competition policy that concern the AER and the ACCC alike, and to possible challenges in generating enough revenue from the arbitrage exercise to sustain the quantum of investment required.

We also comment on current policy initiatives, such as the CIS, and conclude that while it is a costly initiative, it may not function as effectively as intended. While the alternative to let gentailers further entrench their incumbent position in the storage space is not satisfactory, there are other means to promote investment in storage through long-term contracting. These contracts can

be mandated, or supported by some government guarantee in lieu of the expensive CIS. Likewise, as a standby facility, BESS is tremendously costly and fails to take advantage of what storage has to offer. Finally, the limits of arbitrage loom large; this potential market failure can be mitigated by other procurement mechanisms.

Some reforms can be undertaken now to spur investment in storage; these need no research and modest work to be implemented in the NEM. The first of such reforms is the introduction of LMP to (a) introduce more prices to arbitrage away, (b) give the market the right signals to correctly locate their storage assets on a network subject to transient congestion and (c) find the right mix of storage capacity and generation capacity.

This first reform is the precursor to a second one, which is the introduction of a DAM. DAMs allow for a better allocation of resources over time. There are also pro-competitive and therefore the right response to a storage market that is very sensitive to market power.

Of course many questions remain unanswered, especially when it comes to market design. More work needs to be done to evaluate whether energy arbitrage is sufficient to sustain adequate investment in storage, and if needed, to design a mechanism that exactly addresses this potential problem.

Projecting oneself into the future, to a world *without* thermal generation but dominated by variable renewable energy and storage, it is not clear how a clearing price is determined in the spot market, nor what is the best organisation of trade. Indeed, the current (approximate) double auction in the NEM is intended to elicit supply from technologies with increasing marginal cost. So, first there is a connection between clearing price and marginal cost. Second, that marginal cost is determined by the physical characteristics of the machines; it is exogenous. Absent thermal generation, the marginal cost is the opportunity cost of a storage unit willing to buy. That opportunity cost is *endogenous* to the problem, and so very difficult to determine. The same applies to demand: when there is excess supply of VRE and storage is the active buyer, demand is determined by the willingness to pay of the marginal storage unit. That willingness to buy depends on the expectations storage operators form about their ability to sell in the future; it is also endogenous to the problem, unlike currently, where demand is an exogenous object. Then it is not clear how a market is organised. This work is forthcoming.

References

- [1] Allaz, Blaise and Vila, Jean-Luc. "Cournot Competition, Forward Markets and Efficiency". In: *Journal of Economic Theory* 59.1 (1993), pp. 1–16.
- [2] Andres-Cerezo, A. and Fabra, Natalia. "Storing power: market structure matters". In: *Rand Journal of Economics* 54.1 (2023), pp. 3–53.
- [3] Balakin, Sergei and Roger, Guillaume. "Near-optimal storage strategies in electricity markets". In: *Working Paper* (2023).
- [4] Balakin, Sergei and Roger, Guillaume. "Storage cycles". In: *Working Paper* (2024).
- [5] Balakin, Sergei and Roger, Guillaume. "Storage games". In: *Working Paper* (2024).





CONTACT

Guillaume Roger
guillaume.roger@monash.edu

Monash Business School
Monash University, Caulfield Campus
900 Dandenong Road
Caulfield East VIC 3145
Australia

The information in this white paper was correct at the time of publication (October 2024).

The views expressed in *'The storage imperative: Powering Australia's clean energy transition'* are not necessarily the views of the Australian Government, nor of Monash University more broadly. The Australian Government does not accept responsibility for any information or advice contained in this paper.

CRICOS provider: Monash University 00008C.