Storage integration in the NEM

Milestone 2 Report

March 2023

Guillaume Roger

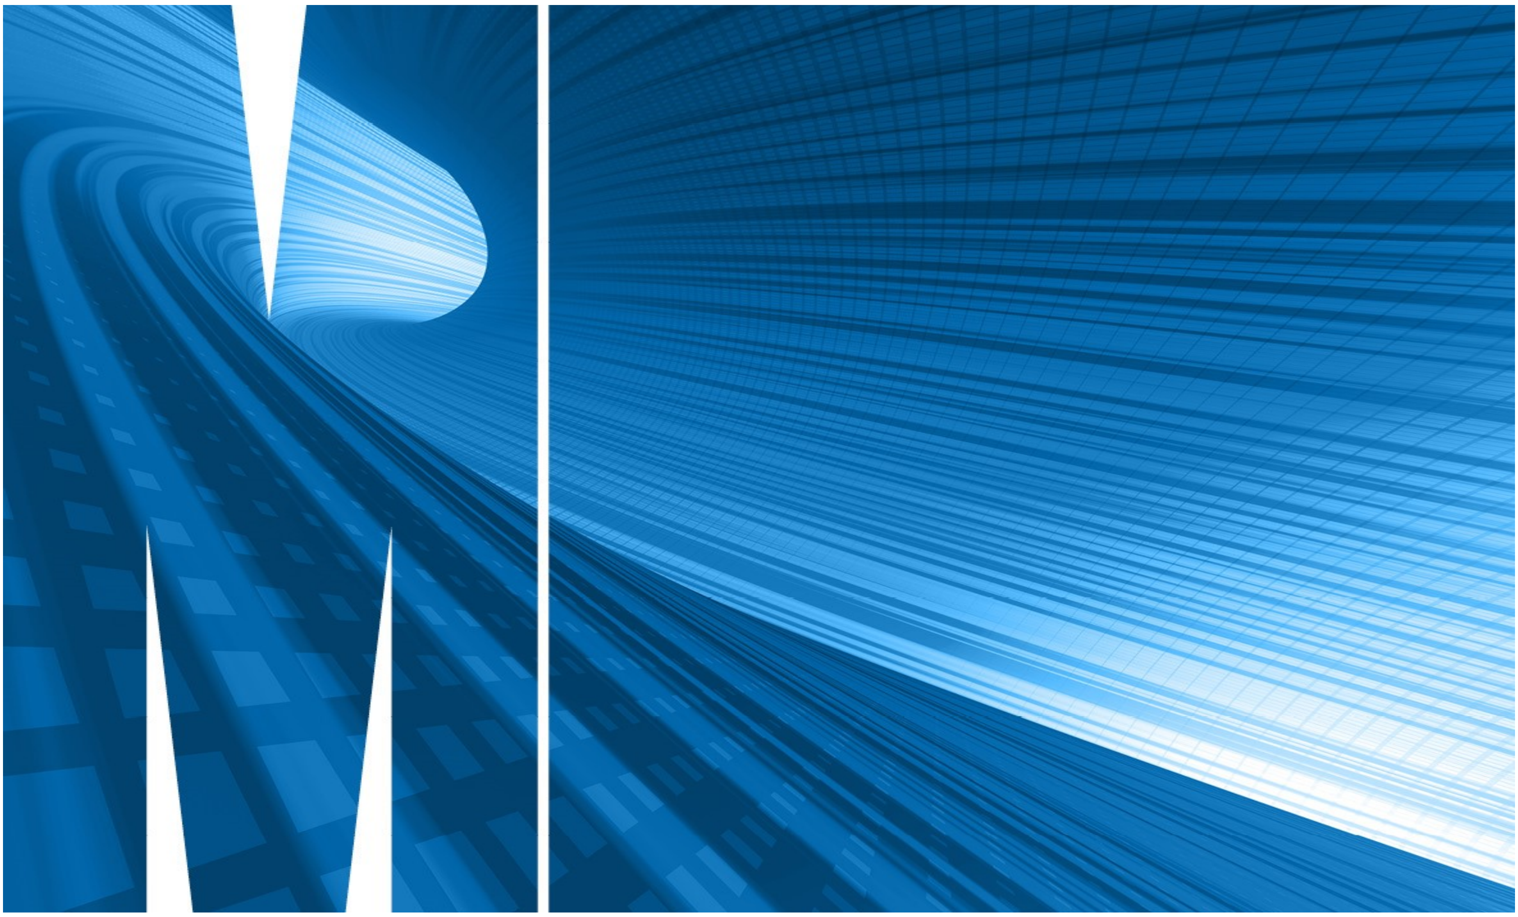

Summary of the project

Electricity storage on large scale is nothing short of revolutionary; it is the perfect, and very timely, complement to intermittently available renewable energy generation. While full of promise, it also entails plenty of challenges. Storage allows for the use of dynamic strategies in buying and selling energy, which market rules must account for; in particular, there is a distinct risk of market manipulation. Clearing and dispatching storage units in the energy market must also be amended to account for this new behaviour. Storage is likely to play a critical role in ancillary services, old (FCAS) and new (synthetic inertia). Finally, unlike conventional generators, storage presents no significant returns to scale, which may offer an opportunity for enhanced competition in the energy market.

This project intends to study the problem of storage integration into the NEM under the lens of economics. It will assist in designing the correct incentives for investment and for the use of storage. In doing so it will contribute to increasing the penetration of renewable energy, whose essential challenge remains intermittency and grid stability.

This project intends to deliver the following:-

- A manipulation-proof market design solution that specifically account for storage in the energy market;
- A dispatch rule that accounts for the dynamic strategies of storage;
- A method to value synthetic inertia, formulate a demand for inertia and a suitable market design;
- An analysis of the competitive landscape with storage, with recommendations to guide storage investments and acquisitions.

This report is the second Milestone Report. It documents early progress from the research team.

Acknowledgements

This project received funding from ARENA as part of ARENA's Advancing Renewables Program. Grant number 2021/ARP016.

Disclaimer

List of Acronyms

ARP	Advancing Renewable Program
AEMC	Australian Energy Market Commission
AEMO	Australian Energy Market Operator
GIH	Grid Innovation Hub
NEM	National Electricity Market
NEMDE	National Electricity Market Dispatch Engine
PSC	Power Systems Consultants
VRE	Variable Renewable Energy

Table of Contents

1 Introduction

Under Grant 2021/ARP016, Monash University received funding from ARENA to study "Storage integration in the NEM" with its partners AEMO, the AEMC, PSC and the GIH. This first Milestone Report is an interim report that documents preliminary findings of the project.

1.1 Background

In Australia the electricity industry is transitioning 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. Following the pioneering, government-supported commissioning of the Hornsdale Power Reserve (HPR) in 2018, grid-scale battery storage is now a realistic option with large investment being poured into it. Neoen, the largest battery storage owner in Australia, already operates 450MW (600 MWh) of capacity, with an additional 500MW coming online before the end of 2023. It is timely in that storage is critical to enable the energy transition on a large scale. In addition, the scope of applications of storage is widening dramatically, from operating mostly in the (small) FCAS market or as a power reserve, to energy price arbitrage, assisting in managing congestion or delivering synthetic inertia to support the national grid.

However, we know very little of the economics of storage, which leave policy-making bodies like AEMO and the AEMC in a conceptual and technical vacuum to develop the policies necessary to integrate storage in the NEM. As of writing there are a handful of serious academic papers written on storage (Karaduman, 2020 and Andres-Cerezo and Fabra, 2022). From these works, we do know that operating a storage unit is very different from operating a power generation unit, and that the industrial organisation of the market is an important consideration for competition policy. We also know that the incentives of private storage operators lead them to follow strategies that depart from the socially optimal allocation. Hence just letting the market decide, as *laissez-faire* would have it, is not socially efficient. Social efficiency, even in the second-best, requires market design and the formulation of corrective policies. To complete this task, policymakers need more and better information as to the economics of storage.

The knowledge vacuum the NEM is exposed to is all the more concerning that the necessary investment in storage to complete the energy transition is staggering. As of mid-2021, the dispatchable capacity of the NEM is approximately 43 GW, with an additional 14 GW of wind and solar capacity. To crudely demonstrate the scale of Australia's aspired clean energy transition that maintains 43GW of dispatchable capacity, consider effecting a 50% transition requires a minimum of 21.5 GW of storage (power). Having this power available for 12 hours (overnight, roughly) requires 258 GWh of energy capacity. By comparison, Neoen currently operates only 600 MWh of storage today. If relying on the same technology (lithium ion), the total cost of investment is in excess of \$100 Bn. Such an expense deserves some study.

The purpose of this project is to explore the economics of storage in detail to inform policymakers, such as the AEMC and AEMO. There are multiple aspects to this project, which we list below. The goal is to deliver a suite of research papers to understand the behaviour of profit-maximising storage operators, to develop an adequate market design in which storage is a significant player, to introduce markets for new services like synthetic inertia, and to understand some issues of competition policy with storage.

1.2 Outline of the project

Here we list the main workstreams that have been identified.

1.2.1 Area 1. Understanding storage and market design for storage

Storage behaves differently from standard generators and loads, and thus in many ways – starting with the fact it can be either at the same time. The first order of business is to better understand storage behaviour, especially when it comes to market manipulation. This done, one must then design market rules that explicitly account for storage and the new strategies it can employ. These are dynamic strategies, and potentially collusive strategies.

1.2.2 Area 2. Clearing and dispatching a market with storage

Dispatching storage units that employ dynamic strategies cannot be achieved using the current dispatch engine (NEMDE), which is a linear programming engine. A new dispatch algorithm must be designed for this task, and a new numerical implementation must be created. This is essentially a problem of mathematics and computer science.

1.2.3 Area 3. New services with storage

Synthetic inertia is already identified as a new service that can be supplied by batteries. However this first requires valuing inertia and then devising a market design to supply synthetic inertia services, noting that this market is obviously linked to the energy market. The principles and the methods that are used when dealing with synthetic inertia may be applied to other services.

1.2.4 Area 4. Industrial organization with storage

Because of large returns to scale, the electricity generation sector is highly concentrated. The benefit of this concentration is a low average cost, but its cost that generators routinely exercise their significant unilateral market power. Storage units do not possess significant returns to scale, so there is no social benefit to large storage units; however the risk of exercising unilateral market power remains. This area of work will study precisely this problem and suggest a limit to the size of a battery before it becomes so large that it can become harmful. It will also study the problems arising from owning multiple batteries, as well as batteries and other assets such as solar farms. Finally, because storage is also attractive to manage congestion on networks, this area will also revisit the structural separation between NSPs and energy supply.

1.3 Outputs and publications

At this point there are no publications to report and only a small number of outputs that we comment on.

2 Spot market with storage

Any market design must include a spot market to address the physical constraint of energy balance in the system. While the spot market is meant to serve immediate (or near immediate) demand, with storage the manner in which this market is organised matters a great deal because of the intertemporal links between trading periods, which storage can straddle.

Here we want to discuss exactly this point, for which we first must introduce a notion of dynamic trading in electricity. This problem bears analogy to managing a dam, but it differs greatly too in that a dam receives stochastic, exogenous inflows while a storage operator decides when and how much energy to purchase as part of an optimal strategy. In that sense, it is closer to managing a costly inventory [**HT78**] or trading securities [**Rosu09**]. We find the literature [**ACF22**], [**Karaduman20**] profoundly lacking on this point and so have to study this problem from first principles.

2.1 Optimal dynamic trading strategies

The basic principle of trading electricity using storage is that of intertemporal arbitrage, which is popularly summarised as "buy low and sell high". Implementing this simple strategy is actually very difficult, even in the simplest of environments.

With Sergei Balakin, we start with a simple Cournot model, in which there are N generators and one (possibly large) storage operator with a finite capacity k. The premise therefore is *imperfect* competition, in which marginal-cost bidding is irrelevant. In each period $t = 0, 1, 2...\infty$, a linear demand $p(q) := 1 - p$ is augmented by a stochastic shock $\varepsilon \in \{-a, +a\}$, each occurring with probability $1/2$. Generators and the storage unit decide how much energy to sell (or buy) in each of these periods in order to maximise their payoff (profit). All participants have a discount factor $\beta < 1$: they face a positive interest rate. Storage also faces efficiency losses: $\delta \leq 1$. We also note with this model, the expected payoff of a storage operator is negative in a given period.

Without going into the details, this simple game admits a very large number of equilibria, most of which are impossible to describe in general terms. That is, there exists a very large number of strategies that each of the generators and the storage unit could pursue, which are all valid in that they each individually cannot do better given what the other participants do. There is one such equilibrium that can be described, is well understood and is robust; it is the Cournot-Nash equilibrium of the stage game, in which generators set optimal quantities as:-

$$q^*(\varepsilon = -a) = \frac{1 - a + b}{N + 1} \text{ and } q^*(\varepsilon = a) = \frac{1 + a - \delta s}{N + 1},$$

where b is the quantity bought, and s the quantity sold, by the storage unit. The corresponding prices are

$$p^*(\varepsilon) = \frac{1 + \varepsilon + b}{N + 1} \text{ and } p^*(\varepsilon) = \frac{1 + \varepsilon - \delta s}{N + 1}$$

This equilibrium is the least profitable for industry participants – conversely, it is the most favourable to consumers. We'll return to this point later.

A storage unit contemplates an entire sequence of these incomes streams over the infinite horizon; moreover, the terms b and s are actually functions $b(c)$ and $s(c)$, where c is the state of charge of the storage unit. That is, how much to charge or discharge depends on how much is currently in store. This quantity is the state variable of our problem; it follows the law of motion:

$$c_t = c_{t-1} + b(c_{t-1}) - \frac{s(c_{t-1})}{\delta}, \quad c_0 = 0.$$

To determine the behaviour of a storage unit in this market, we seek the optimal sequence of purchase and sale decisions, $\{b(c_t), s(c_t)\}_{t=0}^{\infty}$ to maximise the net-present value of the cash-flows to the storage unit. After verifying that the problem admits a Bellman representation, we solve the problem:-

$$\max V_t(c) = \frac{1}{2}\left[-\frac{1-a+b(c)}{N+1}b(c) + \beta V_{t+1}(c+b(c))\right] + \frac{1}{2}\left[\frac{1+a-\delta s(c)}{N+1}\delta s(c) + \beta V_{t+1}(c-s(c))\right],$$

where the first term is the payoff from buying energy (and continuing), and the second one from selling energy (and continuing). In addition, we append the boundary conditions:

$$V(0) = \frac{1}{2-\beta}\left[-\frac{1-a+b(0)}{N+1}b(0) + \beta V_{t+1}(c+b(0))\right] \text{ and}$$

$$V(k) = \frac{1}{2-\beta}\left[\frac{1+a-\delta s(k)}{N+1}\delta s(k) + \beta V_{t+1}(k-\delta s(k))\right]$$

It turns out that this problem is completely intractable without imposing more structure to it. More precisely, one could seek a numerical solution on a computer, but it would rob us from actually understanding what goes on. Imposing more structure amounts to presuming of a trading heuristic and seeking a solution to the value function $V(c)$. While this is not fully optimal, it allows us to *a)* engage in comparative statics – that is, understand how the solution behaves when the environment (the parameters) changes and *b)* use the solution, which is almost explicit, when contemplating market design questions. We report two such heuristics.

2.1.1 Constant fraction

Under this simple heuristic, the storage operator buys or sells a constant fraction $r \in (0,1)$ of the available capacity, starting from empty. Hence the first purchase is rk; if selling the following period, the sale is $rc = r^2k$ and if buying, the purchase is $r(k-c) = rk(1-r)$ and so on. However, because the storage unit is empty at time 0, it must first purchase energy, which can only be (profitably) done with probability $1/2$. The problem to solve is deceptively complicated:

$$\max_r V(c), \; s.t. \; V(0), V(k)$$

because it looks some simple. Nonetheless, Sergei Balakin was able to uncover a recursive structure in this problem, and exploit it to compute the value function $V(c)$ of the storage operator. We leave the details of this work in the academic paper, and focus on the outcome. This function simplifies to a high-order polynomial expressed in terms of the heuristic r and the capacity k:

$$U(r;k) = rk\frac{\beta[(1+a)\delta - (1-a)] - 2(1-\beta)(1-a+rk) - rk\frac{\beta r^2(1+\delta^2)}{1-(1-r)^2\beta}}{4(N+1)(1-\beta)(1-(1-r)\beta)}.$$

This function does not lend itself to easy manipulation by hand, but it can easily and meaningfully be computed. Then the optimal strategy can also be identified. With this one can manipulate the parameters of interest (capacity, discount factor, loss factor...) to understand their impact on that optimal strategy. Below we map the function $U(r;k)$ for various values of the capacity k for $N = 2$, and with parameters $\beta = 0.95$, $\delta = 0.95$. Fig. 2.1 corresponds to high shocks $a = 0.6$ with capacity k moving from 0.15 to 1.15, and Fig. 2.2 corresponds to low shocks $a = 0.2$ where k changes from 0.05 to 0.3.

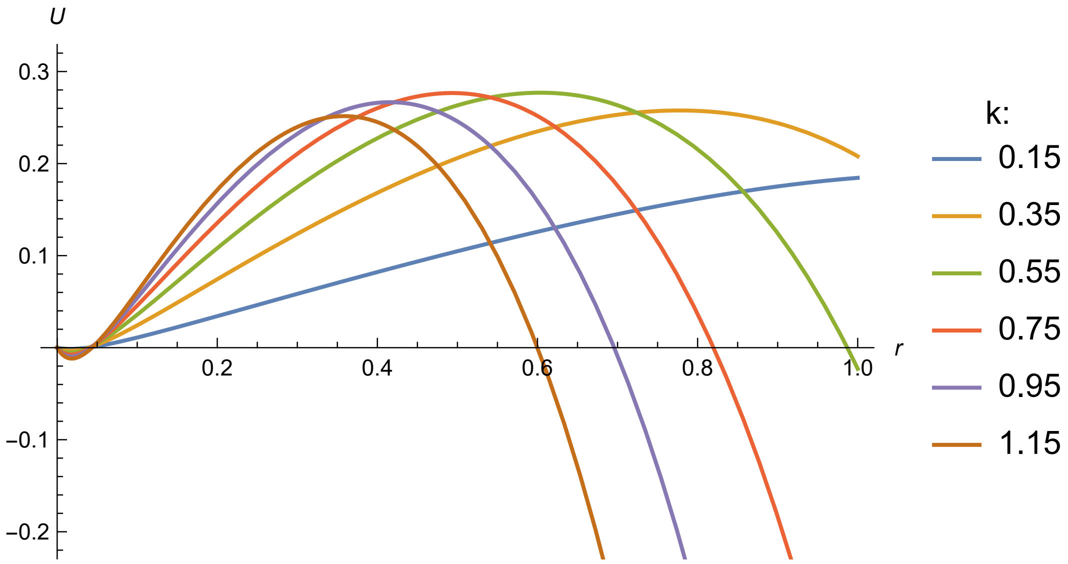

Figure 2.1: Payoff function for different capacities k when the magnitude of the shock is high enough: $a = 0.6$.

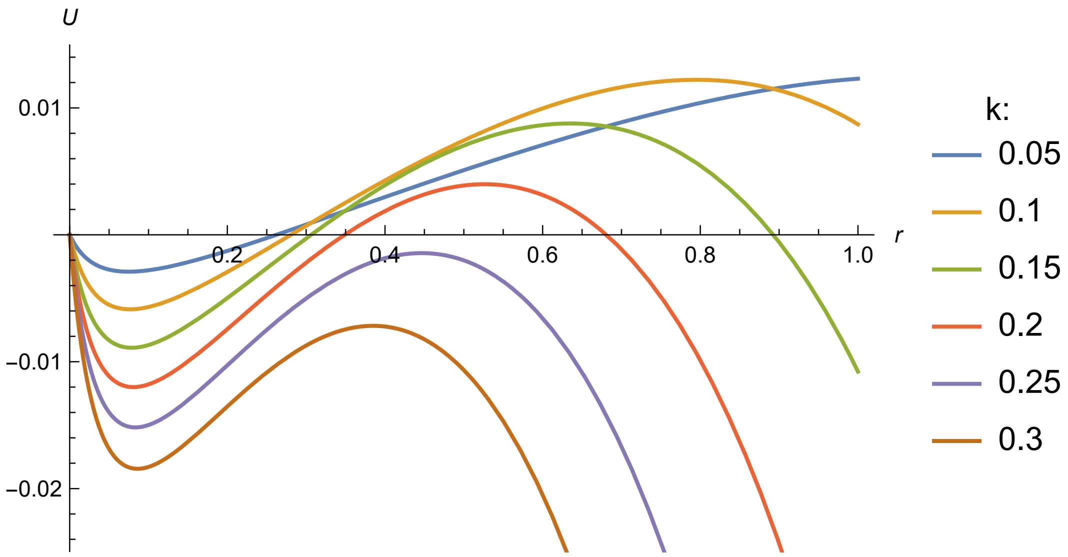

Figure 2.2: Payoff function for different capacities k when the magnitude of the shock is low: $a = 0.2$.

We see that a small unit uses all its available capacity every period. As capacity increases, the fraction decreases. Moreover, when capacity is a choice, the optimal capacity is interior: not too large, not too small. There may be more than one reason for this behaviour. One of these reasons is the price impact of storage: for a given capacity, when a unit can sell large quantities, these large quantities depress prices when selling and increase prices when buying. This unilateral market power is internalised and quantities are restrained. Second, a large capacity nullifies the arbitrage spread that storage feeds on. Therefore, for an optimal choice r, which is really a function $r(k)$, the storage operator restrains capacity investment as well. We also see that very small fractions (of capacity) induce negative payoffs, even though there is no fixed cost to operate the storage unit, nor to hold energy. We believe the reason is that storage starts empty; with small fractions forever, the future (captured by the continuation value βV) is worth nothing, the present weighs comparatively more, and the present starts negative. This third reason is the continuation motive. A successul storage operator must balance all three considerations.

Note that relative size k/a of the storage unit is the same across the two figures. A storage unit is large when that ratio is large: it can address a large fraction of, or even all, the shocks. Hence, while relative size matters a great deal for profitability and for the optimal strategy of a storage unit, the absolute size of shocks is also important. Absent large enough shocks, there is no arbitrage spread to exploit and storage cannot survive.

These results also suggest that large units are not efficient: the capital investment is underutilized, but still has to be paid for. Furthermore, any unit should be used on a sufficiently large scale. Finally, a market operator may want to have the means to increase the fraction r that the storage operator chooses.

2.1.2 Constant quantities

Here the storage operator buys or sell a constant quantity X (e.g 10MW) each period, starting from empty as well. To avoid having to deal with partial fills at the boundaries, we let $X := k/m$, so m is the number of steps to move from empty to full. In the set of admissible strategies, restricting m to be an integer may not be fully optimal, but we expect the corresponding loss to be small – if it exists.

The optimal behaviour differs from the proportional case above. Under the proportional heuristic, boundaries are never reached: the storage unit can never be completely full nor completely empty in finite time.[1] But here it is completely empty or completely full with positive probability. Reaching these boundaries induces a reflection of the probability mass, which then generates "waves" between these bounds. Sergei, once again, is able to characterise this process, and compute the value function. Let:

$$B(X) := \frac{1-a+X}{N+1}X \text{ and } S(X) := \frac{1+a-\delta X}{N+1}\delta X,$$

where B represents buying and S selling. Then

$$\begin{aligned} V(X) = \frac{-B(X)+\beta S(X)}{2(1-\beta)} + \frac{1+\beta}{\beta}\Bigg(& B(X)\sum_{j=0}^{\infty}\left(\frac{\beta}{2}\right)^{(m+1)(2j+1)}\sum_{i=0}^{\infty}\left(\frac{\beta}{2}\right)^{2i} C^{i}_{2i+(m+1)(2j+1)} \\ & - S(X)\sum_{j=1}^{\infty}\left(\frac{\beta}{2}\right)^{2(m+1)j}\sum_{i=0}^{\infty}\left(\frac{\beta}{2}\right)^{2i} C^{i}_{2i+2(m+1)j}\Bigg) - \frac{S(X)}{2\beta}\sum_{i=1}^{\infty}\left(\frac{\beta}{2}\right)^{2i}\left(2C^{i-1}_{2i-1}+\beta C^{i}_{2i}\right), \end{aligned}$$

where C^i_j denotes binomial coefficients. For the interested reader, the first term in this expression is very simple: it is the spread on energy arbitrage discounted over an infinite horizon. The remaining terms are complicated; it is the cost of these "waves". That is, this is the cost of being stuck at the boundaries with positive probability – for example, having a full storage but not being able to sell.[2] This clearly acts like a penalty on the first, simple term. This function is depicted below.

Below we reproduce the same pictures as for the proportional case.

With constant quantities, the optimal choices can only be discrete; they amount to choosing the optimal number of steps $m \in \mathbb{N} \setminus 0$. As for the proportional case, a small unit uses large steps – only one step in the extreme – and a large one uses (many) small steps. However, somewhat surprisingly, for a given capacity the storage operator can achieve higher profit under this heuristic than under the proportional approach. We understand the reason to be that the constant-quantity heuristic removes all the very small sales (and purchases) that occur under the proportional one.

[1] Boundaries may be approached, however asymptotically only.

[2] For the really keen reader, that second, complicated term is a hypergeometric function of the binomial coefficients of the problem.

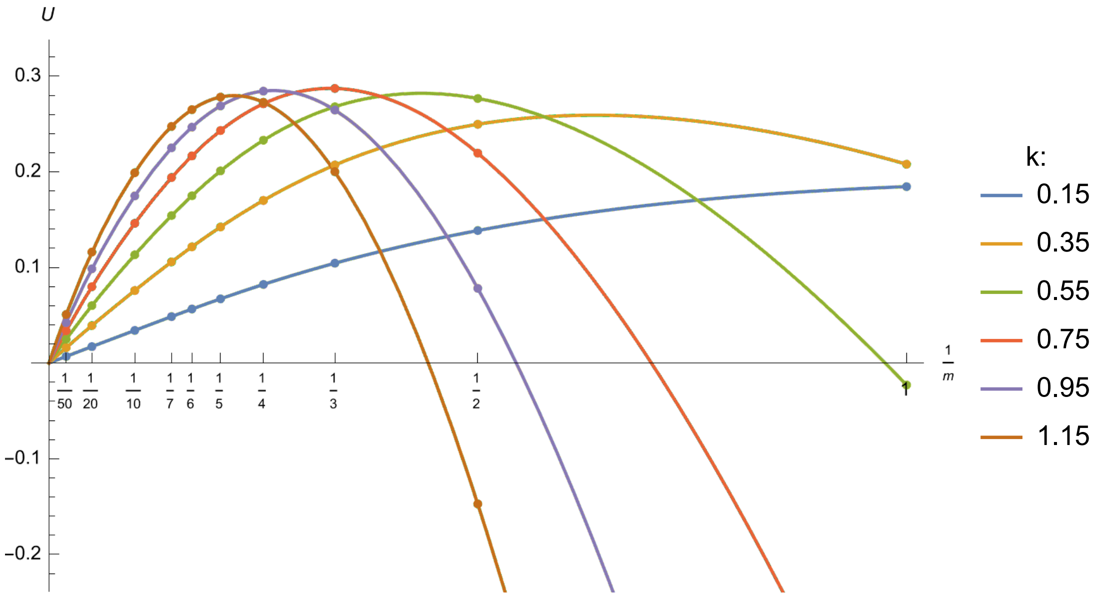

Figure 2.3: Payoff function for different capacities k when the magnitude of the shock is high enough: $a = 0.6$.

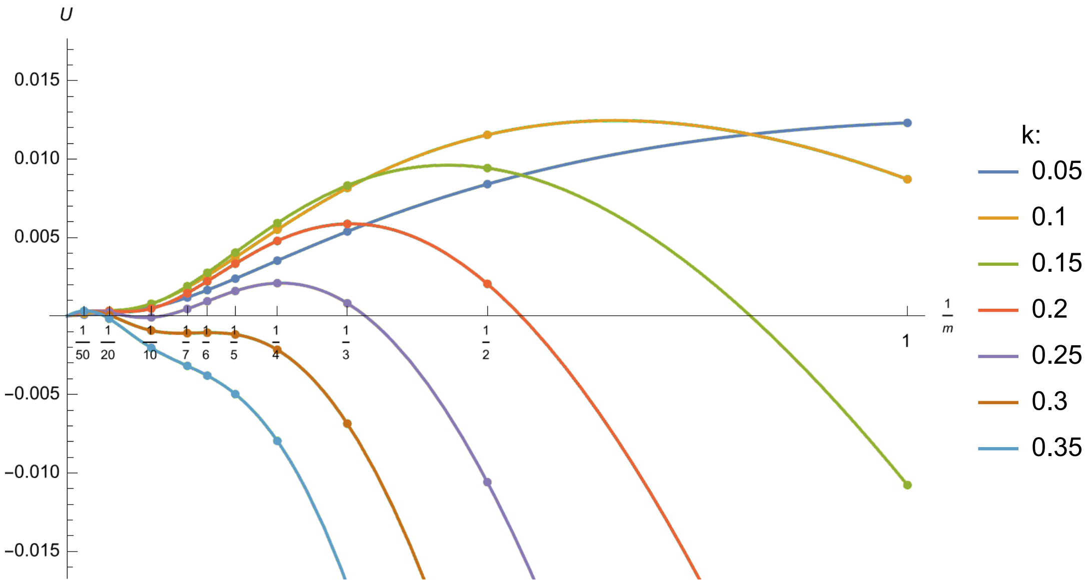

Figure 2.4: Payoff function for different capacities k when the magnitude of the shock is low: $a = 0.2$.

2.1.3 Takeaways

Much to our surprise, the optimal quantities under either heuristics are *identical*. That is, a storage unit with capacity 0.6 sets the optimal fraction to $1/3$ and optimally sells quantity 0.2 under the constant quantity heuristic (i.e. $m = 3$). However it generates more profits under the latter. The reason is that under constant quantities, the storage operator always sells or buys a sizable quantity of energy, and so avoid the pitfalls of very small transactions.

Capacity underutilization remains puzzling: that storage operators internalise their market power is one, likely partial, explanation. We believe it may be rooted in the nature of the stochastic shocks we model. It is very possible that a richer stochastic structure with either little persistence or a lot of persistence delivers higher utilisation: with little persistence, a full unit has strong incentives to discharge in full upon a positive demand shocks. Conversely with a lot of persistence, a full unit can expect multiple selling opportunities in a row, and it is less likely to get stuck full for long. Hence it may charge more in prior periods too.

Very small shocks do not serve a storage operator well, even if the ratio of k/a capacity to shock does not change. Here too, the size of transactions matters. With small shocks the quantities a storage operator buys and sells are small; but we know that leads to losses. This accords with the intuitive understanding that storage needs volatility. The details we add is "*which volatility ?*" The size of shocks matters, and we expect the nature of the stochastic process also matters.

2.2 Exclusionary equilibrium

As mentioned earlier, this equilibrium we focus on is by far not the only one. Indeed there exists an equilibrium, in which generators can construct a collusive strategy so that the storage operator *never* finds it profitable to purchase energy. Since it starts empty, it never operates.

We attract attention to this equilibrium, which we have been able to construct, for two reasons. First, it is a real risk as storage is only nascent in Australia. Storage is useful to renewable energy generation, but it is a very serious competitor for gas generators. They have every incentives to preempt entry by storage, and may be able to coordinate on such an equilibrium. This equilibrium is likely not the only exclusionary equilibrium; that is, there may be more ways than one to preempt entry. Second, the Cournot-Nash equilibrium we are able to study is the least profitable to generators, whereas the exclusionary equilibrium is more profitable, and therefore more attractive.

2.3 The path forward

The results presented so far are preliminary: they stem from a very simple model that does not capture all characteristics of the NEM. They are also a first step to understand the behaviour of storage operators, which feeds into the market design problem.

Next we tackle some extensions of these two basic models by allowing for:-

- Capacity constraints on generators;
- Markov chains with little and much persistence;
- A time-varying demand, which can accommodate deep storage;
- Cost-minimizing allocations.

With these extensions in hand we can then turn to market design. In particular, the cost-minimizing allocations inform us as to what is desirable for consumers. The exercise of market design then amounts to finding a set of rules, including a bidding space, that deliver these cost-minimizing allocations.

3 The combinatorial problem

The premise of this problem is that there exists complementarities between trading periods, which are induced by the start-up costs of large generators. These start-up costs introduce non-convex payoffs to generators. Combined with the energy-balance condition, these non-convexities induce phenomena such as negative prices and high price volatility.

In the case of multi-unit auctions, a combinatorial approach is able to harness complementarities between units for sale – for example, in the case of spectrum auctions, contiguous geographic units are more valuable together and so can attract higher bids if packaged together. The reason is that receiving $A + B$ is more attractive than A and B with some probability. We seek to build on this insight to package (time-)contiguous trading intervals.

3.1 Start-up costs and their impact

Start-up costs have long been known but their impact on the NEM has not been well documented. Using data from Western Australia (so, the WEM rather thena the NEM), Jha and Leslie (2022) show that start-up costs have a significant impact on price dynamics. First, these start-up costs are convex. Second, combined with a large VRE penetration, they induce higher price volatility and even higher average prices. Finally, the combination of VRE penetration and start-up costs of thermal generators induce *higher* CO2 emissions. The reason for all this is that a large VRE fleet displaces thermal generators in the merit order. Yet these generators are needed to meet peak demand when VRE is unable to supply, but have very little time to ramp up. The convexity of the cost function is reflected in both the price path and in the emission level.

Start-up costs have other consequences on the behaviour of the market, as suggested before. More precisely, these costs induce what economists call non-convexities; it is well known that non-convexities typically lead to non-existence of a Walrasian equilibrium. That is, the market may not clear. But in electricity, the market *must* clear: market clearing is identical to the energy-balance condition. Hence something must give; what gives here is chaotic bidding, negative prices and high price volatility, which are all a symptom of a market that does not clear well.

Combinatorial bidding has a second benefit beyond harnessing complementarities. It also renders bidding contingent in the following sense: bidding on a 1-hour duration block, for example, is equivalent to bidding on a 5-minute block *conditional* on also securing dispatch on the other 11 intervals. Contingent bidding is less risky for sellers, who in turn can bid more aggressively. This is expected to lead to lower prices.

3.2 Constructing blocks

Introducing block bidding is technically more difficult than it is to discuss conceptually. Unlike a spectrum auction, for example, the supply is elastic and the demand for blocks does not correspond to a valuation. In other words, determining the demand for blocks of energy is not easy: there is an aggregate demand for energy, not a demand for 1-hour blocks of energy, for example. However it is essential to determine the price at which this energy can be traded.

With Ningyi Sun, we introduced auxiliary markets for energy in arbitrary block durations. More precisely, we conceive of a game between the market design (or operator, as one wishes) and the suppliers. First the designer assigns an inelastic demand d_i for each such block market i, of which there may be n, that can be informed by the forecast load shape. The duration of a block is a key characteristic. Second, given each $d_i,\ i = 1, 2, ...n$, sellers j bid in supply functions $S^j(p_i)$ in each of these markets, as is currently the case in the NEM.

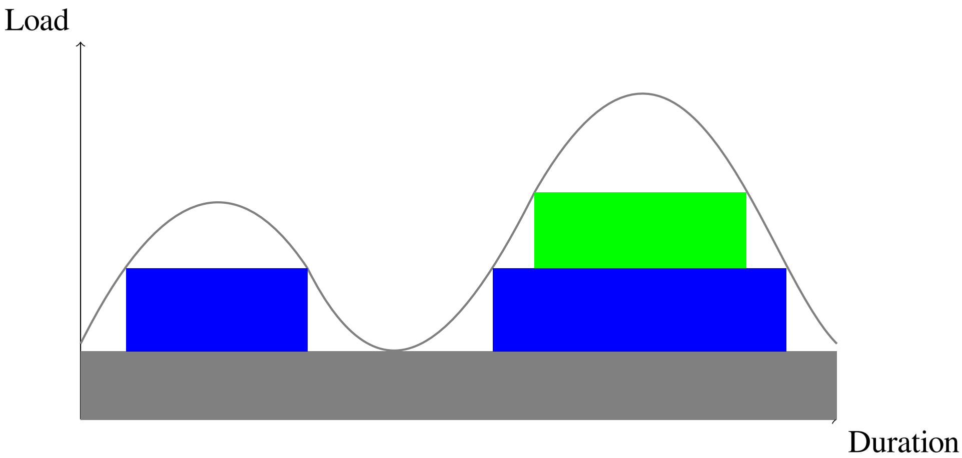

Figure 3.1: Load shape and block design – illustrative

Such a game is solved using backward induction. Given the vector **d** of auxiliary demands, each of these n markets clears according to

$$S(p_i) = d_i \rightarrow p_i := p_i(d_i) = S^{-1}(d_i).$$

where $S(p_i) = \sum_j S^j(p_i)$ is the aggregate supply function in market i. Then for all 5-minute intervals t contained in block i, the residual demand is

$$RD(p_t) = D(p_t) - d_i,$$

where $D(p_t)$ is the forecast demand at time t. This formulation makes it plain that block bidding is a substitute for bidding in the 5-minute market. Since there are n such markets, we have in fact

$$RD(p_t) = D(p_t) - \sum_i^n d_i,$$

The 5-minute market clears in the usual way:

$$S(p_t) = RD(p_t) \rightarrow p_t := p_t(RD_t) = p_t(D_t - \sum_i^n d_i)$$

Given the best-reply of suppliers (their bidding functions for a given vector **d**), the designer then wants to determine the best possible "synthetic" demand d_i for each interval i. Let m_i denote the number of 5-minute intervals contained in block i; for example, for a block of 1-hour duration, $m_i = 12$. For the designer, this amounts to solving the problem

$$\min_{\mathbf{d}} \sum_{i=1}^{n} \left[\sum_{t=1}^{m_i} p_t \cdot \left(D_t - \sum_{j \neq i}^{n} d_j - d_i \right) + p_i d_i m_i \right]. \tag{1}$$

3.3 Some results

A solution to Problem (1) is characterised by the necessary conditions:

$$\forall i,\ \sum_{t=1}^{m_i}\left[p_t + \left(D_t - \sum_{j\neq i}^{n} d_j - d_i\right)\frac{\partial p_t}{\partial d_i}\right] = \left[p_i + d_i\frac{\partial p_i}{\partial d_i}\right] m_i, \tag{2}$$

where $\frac{\partial p_x}{\partial d_i} = \frac{\partial}{\partial p_i}S^{-1} > 0,\ x = t, i$, and $S^{-1}(\cdot)$ is a convex function.[3] The term $\frac{\partial p_x}{\partial d_i}$ is the price impact of changing the quantity d_i (on the elastic function $S(\cdot)$. Since the function $S^{-1}(\cdot)$ is convex, each of these n equations has a unique solution – the LHS of (2) decreases, while its RHS increases in the term d_i.

If a solution to this problem exists, it is such that the vector **d** of optimal quantities balances the marginal impact in any market i with the sum of marginal impacts in each of the m_i intervals it corresponds to. One loose interpretation is that the clearing price p_i of block i is the average clearing price of the m_i intervals it contains. To be obvious:

$$\forall i,\ p_i + d_i\frac{\partial p_i}{\partial d_i} = \frac{1}{m_i}\sum_{t=1}^{m_i}\left[p_t + \left(D_t - \sum_{j\neq i}^{n} d_j - d_i\right)\frac{\partial p_t}{\partial d_i}\right].$$

This is a very natural condition that amounts to no-arbitrage between any of the block markets and the 5-minute market.

Solving Problem (1) first requires showing the system of n equations (2) it gives rise to, *jointly* has a solution. Indeed, we are able to show that a solution does exist. Furthermore, at the optimum **d**, we can prove that average procurement costs are lower. In fact, these multiple markets have the same effect as forward markets do on the spot market: they alleviate the infra-marginal effect in all markets, which makes for more aggressive bidding in all markets.

3.4 The path forward

This solution of introducing block-bidding is nice, and a bit of a breakthrough, but so far only a theoretical construct. It presents the great benefit of clearing markets very simply and naturally, once the demand parameters d_i are set. These parameters can be optimised over. In principle, the vector **d** can be reset every day.

Next we wish to make this solution more relevant empirically, and test it to check it does not return crazy allocation. More precisely:-

- determine whether to disclose **d** to market participant (see [**KM89**] on the role of uncertainty on the uniqueness of an equilibrium);
- use the load-shape data to determine the block durations and bounds on size;
- test the market design in the lab using experimental subjects;
- reconcile the day-ahead market with the spot market;
- scale up to the size of NEM.

[3] $S^{-1}(\cdot)$ is convex because $S(p)$ is concave; $S(p)$ is concave because marginal costs are at least weakly convex.

4 Synthetic inertia

As mentioned in the introduction, storage can deliver new services to support the grid as the natural features that are the by-product of thermal generation progressively disappear. One such service is synthetic inertia, and the project aims to bring some clarifications as to how a market for a service like synthetic inertia may be organised.

4.1 Approach

The main difficulty in designing a market for synthetic inertia is that there is no demand one can rely on – like there is a demand for energy delivery. Furthermore, constructing a demand function – as we do for example with the auxiliary markets in Section (3) – requires *valuing* synthetic inertia as a service.

This exercise itself is difficult in that AEMO is not yet able to even describe what one more unit of inertia (for example, the increment dH) does to the NEM, neither locally nor globally. This engineering problem is outside the scope of this project, but it feeds into it.

To circumvent this problem, Dr Omer Karaduman (Stanford) suggested using high-frequency data from AEMO to identify large output drops from individual generation units. We are using 4-second generation and frequency data from AEMO to identify and estimate the demand for inertia. There are several challenges to do this. Inertia is not really a product but rather a physical response that is naturally supplied by active producers with rotational power generation after generator/transmission trip; that response occurs well under 1 second. Our 4 second data doesn't provide the ideal level of detail to fully recover the curve of frequency recovery. On top of that, depending on the complementarities of other market products such as FCAS 5-second raise, for example, there are other factors that impact frequency other than the inertia.

Despite these challenges, we can estimate the potential for synthetic inertia in the system. With 4-second data, we can identify generation and transmission failures. By estimating the size of the failure, the inertia potential at the time in the system, and FCAS-market enablements (that is, actual supply of energy in response to an FCAS contingency), we can estimate the frequency recovery under different levels of inertia responses. Then, by considering the future generation fleet and probability of trip events, we can determine how much synthetic inertia needs to be procured at specific locations of the grid. With this, we can ensure that power markets can continue to operate efficiently and effectively.

4.2 Progress

We obtained a very large data set from AEMO that contains all output information in 4-second increments. The raw data stands in excess of 2 Tb, which is too large for anyone to handle. Hence we hired a data scientist to help us organise this information in a more manageable manner. This reduced the size of the dataset to 400 Gb only.

From this data we now know how to identify output drops. Next we need to cross-reference them to distinguish between transmission and power plant failures, and match these events to FCAS 5-second contingency enablements. The net variations in the frequency responses are the inertial responses we are looking for. In the last step, we can recover the distribution of these events at various critical points of the grid.

5 Summary

In this report we present some novel results on the management of a storage unit in a simple stochastic environment. This first step is essential to understand how a storage operator responds to the institutional environment, that is, to the market rules. We find that market power is a chief concern of a forward looking storage operator; in equilibrium, a comparatively large storage unit uses a small fraction of its capacity (always less than 50%, and much less the large it is). The two primary motivations for such a strategy are *a)* the unilateral market power of the unit, which adversely moves prices and *b)* prudence as positive shock may be followed by another positive shock – in which case the storage unit wants to sell again. While only preliminary, this work already informs policy: large units are problematic and wasteful.

Future work on this topic includes modification of the basic model, especially to account for periods of predictably low and high demand (i.e. a time-varying average demand), the introduction of capacity constraints and persistence in shocks. Once that is understood one can think of market design.

We also introduce combinatorial bidding in electricity markets, which can be implemented as a day-ahead market. The optimal design implements a no-arbitrage condition between markets, which equivalently states that the marginal value of electricity is identical across markets.

Finally we make some progress in the empirical identification of a demand for synthetic inertia, which is the necessary first step to contemplate setting up a market for essential system services.