MONASH
University

Storage: trading strategies under uncertainty
A white paper

September 2023
Guillaume Roger – Monash University

ARENA

Australian Government
Australian Renewable Energy Agency

PSC

Specialist Consultants to the Electricity Industry

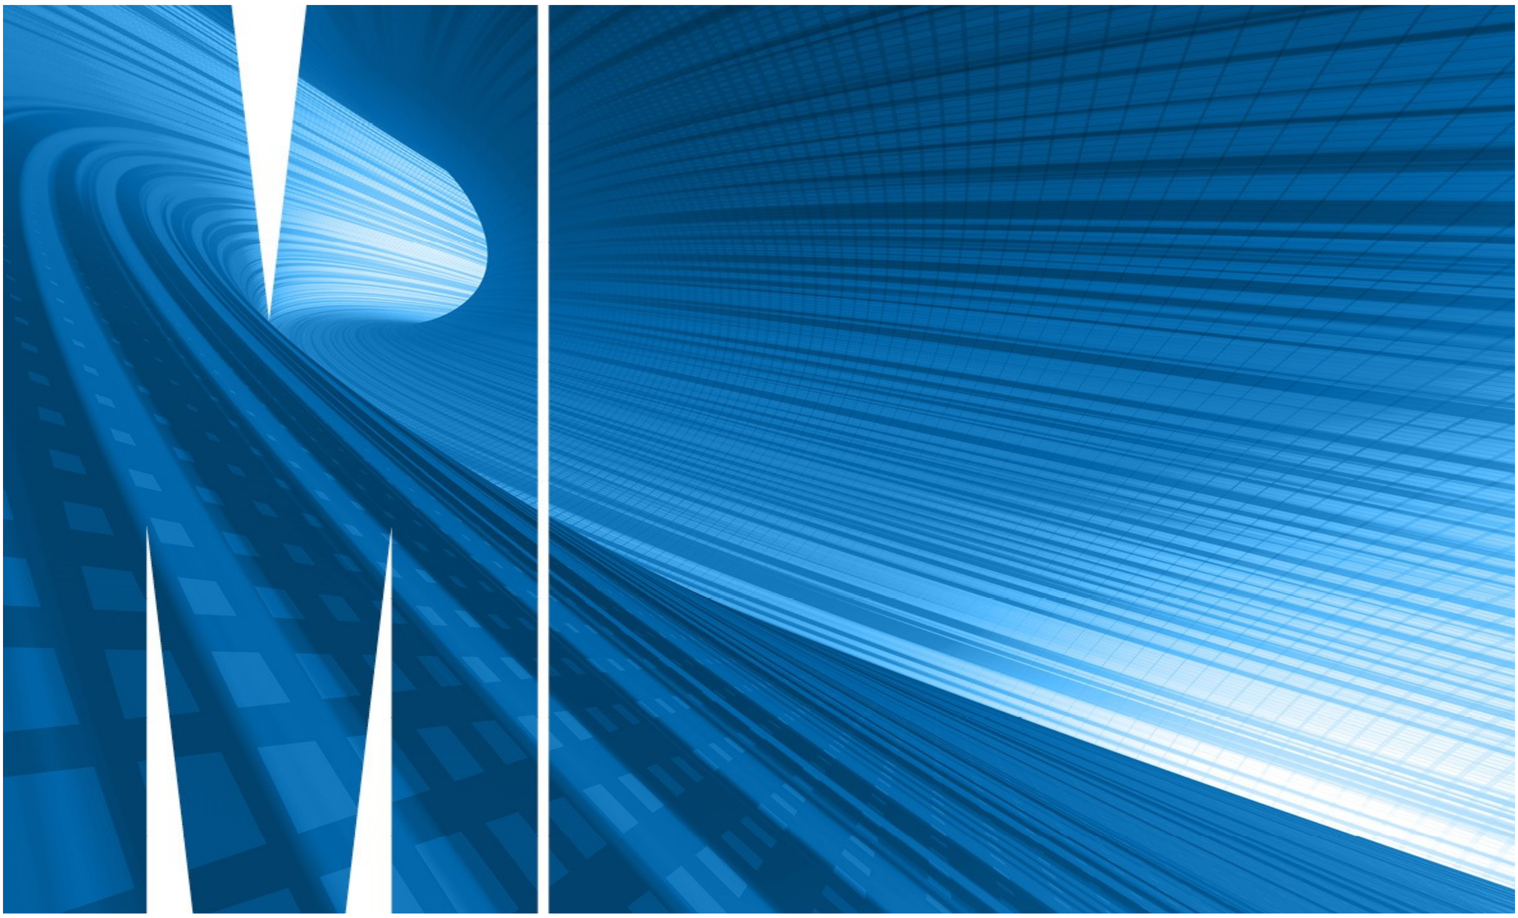

Summary so far

Electricity storage on large scale is the perfect, and very timely, complement to intermittently available renewable energy generation. While full of promise, it also presents plenty of challenges. Storage allows for the use of dynamic strategies in buying and selling energy, which is a novel dimension in most markets but especially in an energy-only spot market like the NEM.

This white paper reports on early work that informs the trading of energy based on the use of storage. The simple motto "buy low and sell high" – or, arbitrage – does hold, however with significant caveats that shed some light on how exactly energy ought to be traded over a long horizon. This may be of use to market participants, market operators and market designers alike.

Attention is limited to one (or more) storage operator(s) facing random shocks to demand, but for whom average demand remains constant; there is no intra-day cycle. Yet of course there is still an important role for storage to play in arbitraging these shocks and in smoothing quantities and prices. Storage faces a standard trade-off between current payoff and continuation, but the details of this trade-off are novel. First, a storage unit with market power affects prices and so it affects the very arbitrage spread it seeks to pocket. Hence storage confronts a sharper version of the standard price-quantity trade-off – selling a marginal unit requires decreasing the price on all infra-marginal quantities. Here, it also applies to *buying* energy. Second, uncertainty weighs heavily and induces a new *continuation risk*. A storage unit can find itself empty and unable to buy or, worse, full and unable to sell. With a large capacity, the only way to mitigate this risk is to delay its occurrence, that is, to withhold quantities. Absent market power, this consideration is of second order: the arbitrage spread remains constant and attractive enough to make the simple decision to buy or sell – the continuation risk does not vanish but it has little impact. This has consequences for optimal strategies, investment decisions, competition policy and market design.

When multiple storage units compete, managing quantities is even more important. Competing operators can do so by colluding. We identifying multiple forms of collusion; this multiplicity is rooted in the multiplicity of actions a storage operator must take: buy and sell. Hence operators may engage in *partial* collusion, whereby they take turns to buy but sell together, or in full collusion, which involves taking turns to buy and also to sell. Finally, we also find that sometimes collusion is the only equilibrium: competing generates systematic losses.

Acknowledgements

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

This "white paper" is based on joint 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.

Disclaimer

List of Acronyms

ARP	Advancing Renewable Program
AEMC	Australian Energy Market Commission
AEMO	Australian Energy Market Operator
BESS	Battery Energy Storage System
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

This paper reports the extant of academic research on the economics of storage in an approachable way; it also documents preliminary findings of a specific project ("Storage integration in the NEM") that is funded by ARENA and supported by PSC.

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. At the risk of stating the obvious, 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. This institutional uncertainty add to the standard business risks of a new venture. As of writing there are a handful of serious academic papers written on storage (Karaduman, 2020 [6], Andres-Cerezo and Fabra, 2022 [2] and Butters, Dorsey and Gowrisankaran, 2023 [3]). 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 the formulation of corrective policies by bodies like the ACCC and the AEMC, which have direct implications for the development of the industry.

This knowledge vacuum 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. Hence this paper.

1.2 Things storage can do

Storage enables the intertemporal shift of electricity production to make energy accessible when required rather than when available. This largely applies to VRE, but not exclusively. In doing so, storage is expected to smooth demand and therefore prices.

Another way of saying this is that is allows for energy arbitrage; this arbitrage can take advantage of predictable price variations – the intra-day cycle – and *unexpected* price shocks – stochastic arbitrage.

Storage can also be used to manage congestion on transmission and distributions networks, which opens the question of location choices as well as market design (locational marginal pricing). When thinking about new transmission investment, storage ought to be included in the mix to determine the optimal transmission (or distribution) capacity choices. 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 can be used to manipulate the market, and it renders tacit collusion *easier* than it is now.

1.3 Scope of this white paper

This paper cannot treat all these topics (yet), so it is confined to energy arbitrage. It reports some results from other researchers as well as the work performed by the Monash team on *stochastic arbitrage*.

The goal of our project ("Storage integration in the NEM") is broader; that is, to explore the economics of storage in detail. There are multiple aspects to this project, on which we plan to report regularly. We are developing 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.

2 Trading electricity over time

Any market design must include a spot market to address the physical constraint of energy balance in the system – hence, with or without a day-ahead market. 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. The basic principle of trading electricity using storage is that of intertemporal arbitrage, which is popularly summarised as "buy low and sell high".

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 [5] or trading securities [8], [4]. It does differ profoundly from these works, where traders seek to diversify idiosyncratic risk. To the contrary, arbitraging electricity prices is tantamount to supplying insurance against aggregate risk. Perhaps this explains why implementing this seemingly simple strategy is actually very difficult, even in the simplest of environments.

2.1 A brief review of the literature

There exist a nascent literature on storage, however only a handful of papers stand out. [6] studies the performance and the market impact of the HPR, which makes this work particularly relevant to readers of this paper. The main findings are that both the market power of the HPR, and the response of other market participants to that market power, have to be accounted for; second, the HPR, at least in 2018, could not operate profitably in the energy market only. So energy arbitrage seems to be more difficult to implement than simple stating "buy low and sell high".

Market power matters, even for a relatively small unit like the HPR; with market power, the very participation of the storage unit in the market curtails the arbitrage spread. When the HPR buys, the price increases; when it sell, the price decreases. In some sense, it is the point of arbitrage, and that is well internalised in [6]. But that part seems to left out by market participants at times. It is quite clear that accounting for market power is important to make the correct investment decision.

[3] study the Californian market; they make the simplifying, but heroic, assumption that storage units behave competitively; that is, the have no market power. Even then, as the aggregate capacity increases, the arbitrage spread (rapidly) narrows, the revenue accruing to storage stops increasing and investment in storage slows down. They even estimate a degree of cannibalisation of the VRE revenue by storage in the early stages of the ramp-up phase. That is, while storage nicely complements VRE generation in the period of low demand, it becomes a substitute as soon as demand increases – and storage (aggregate) capacity is large enough.

[2] actually makes the point that market power does matter, especially when a storage unit can be owned by a (dispatchable) generation unit. In that paper there is no uncertainty; market power induces the standard price-quantity trade-off, so quantities traded decrease. The important observation is that in periods of high demand, storage and generation are substitute (competitors). It is also exactly in those periods that this competition should be harnessed to its full extent. But if the same entity owns generation and storage, this is exactly the time when it wants to reduce competition the most. Hence a dispatchable generator should not own a storage unit.

[1] study another aspect of storage that is connected to the intra-day cycle of load. It is often argued that storage and VRE are complements; this is true if VRE and demand are negatively correlated. Solar generation is broadly negatively correlated with demand. However *not all* VRE is negatively correlated to demand; for example, in parts of the world, wind is uncorrelated, or even positively correlated with

demand. In this case storage and VRE may be *substitutes*, which *decreases* the arbitrage spread and the surplus to both storage and VRE. This carries over to the returns on each of these investments, and therefore the investment decisions.

We find the literature [2], [1], [6] and [3] short on the details necessary to truly understand the mechanics of the simple idea of "buy low, sell high". Therefore we have to study this problem from first principles.

2.2 Near-optimal trading strategies under uncertainty

In this main section we cover two papers we completed recently to better understand how to trade electricity (*a*) over time and (*b*) under uncertainty. We find that uncertainty in particular is very costly, however in a way that is novel and which we make precise.

2.2.1 Tools of the trade

The situation we wish to understand is best described as a stochastic game [7]. This is a fancy word to depict and analyse a situation in which actors (players) face an uncertain environment over time and take actions to maximise their payoffs. These actions simultaneously respond to this uncertain environment, and also affect it. In the context of storage, an action is simply to charge or discharge in response to some state of demand – equivalently, price level. In charging, for example, the unit contributes to increasing demand; having charged, it is ready to discharge, that is, ready to contribute to increase supply. That is the sense in which actions affect states. But of course, things are uncertain so a storage unit may be ready to discharge, but unable to.

Without going into the details, even a very simple version of this 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 the players 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).[1] Because of these properties we focus on this equilibrium. Even then, completely describing the *optimal dynamic* strategy of the storage operator(s) remains impossible. To make progress we use *heuristics*: we impose a type of behaviour and in that class, we identify the best (optimal) behaviour.

In the sequel we dispense with the technical details such as mathematical notation, definitions, characterisation of equilibrium and so on. Rather we focus on explaining how uncertainty unfolds and affects the decision of (a) storage unit(s), and on the results we derive. We refer keen readers to the original academic papers that can be found on the Knowledge Sharing page of the project: https://www.monash.edu/energy-institute/grid-innovation-hub/home/integrating-energy-storage-into-the-nem.

Throughout we maintain the following elements. There are $N < \infty$ generators with market power. The premise therefore is *imperfect* competition, in which marginal-cost bidding is irrelevant. Demand is stochastic, that is, subject to positive or negative shocks in each period $t = 0, 1, 2 ... \infty$. We vary the size of these shocks and the exact nature of the random processes. 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: $\delta \leq 1$, and storage has a finite capacity k. What really distinguishes a storage unit from generators is (*a*) it contemplates an entire sequence of these incomes streams over the infinite horizon and (*b*) it

[1] Small technical comment: for simplicity we model a quantity game rather than a supply function equilibrium, which is often used in electricity markets. We argue not much is lost in this because the uncertainty we model is only a two-point process.

must charge (buy) before discharging (selling). Also storage always starts empty. These characteristics are completely essential and induce very specific behaviours.

2.2.2 A single storage unit

In this first, exploratory paper we suppose there is a single (possibly large) storage operator. We study two heuristics; one is the constant-fraction heuristic, whereby the storage operator always buys or sells the same fraction $r \leq 1$ of its capacity k – for example, $1/2$. The second heuristic prescribes constant quantities – for example, 50 MWh. This is equivalent to dividing capacity k by some number m; then it takes m steps to charge or discahrge in full. The choice of either heuristic induces a (endogenous) stochastic process with rich details, which we study in the academic output.

For each of these heuristic we study a problem in which the shock is binary (positive or negative) with the same probability. With this, the otherwise intractable problem becomes just difficult and the payoff of the the storage operator can be reduced to a high-order polynomial function of the capacity k. We can also study some comparative statics with respect to both k and the size of the shocks ε. In Figures 2.1 and 2.2 we depict these payoff functions, first for the constant fraction and then for the constant quantities. In both cases, we set the discount factor β to 0.95, the round-trip efficiency to 0.9 and the shock ε to 0.6.

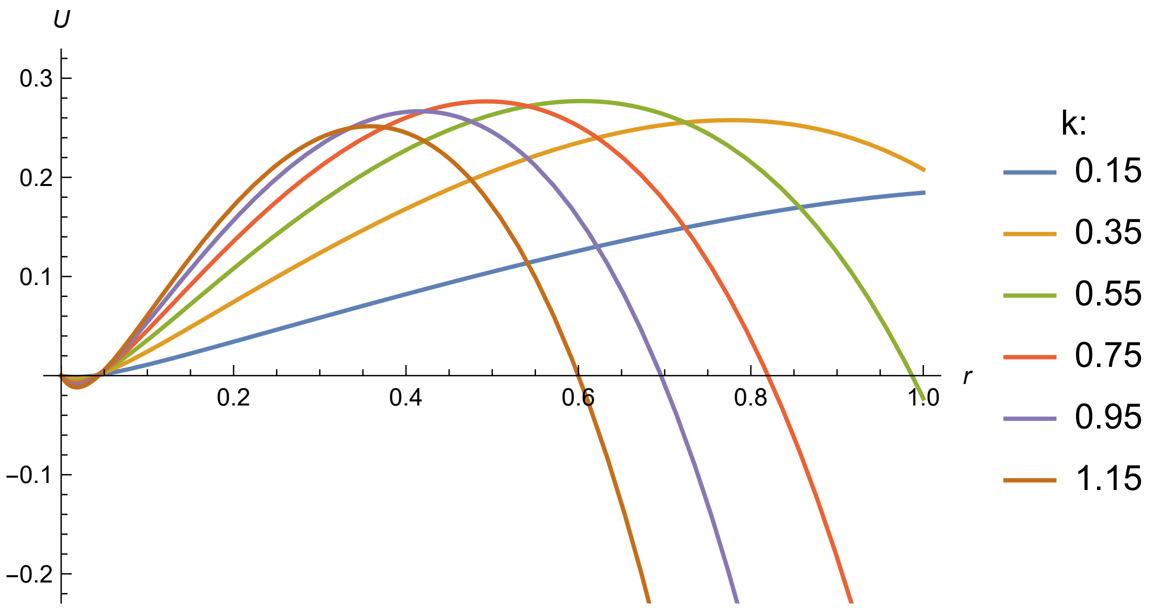

Figure 2.1: Payoff functions with constant fraction for different capacities k.

Some details do differ but both heuristics convey the same substantive messages. The storage unit must balance two considerations: its contemporaneous payoff and its continuation value. This is not new in dynamic problems, but how it materialises here is novel.

When the capacity k is small, the storage unit charges and discharges in full at every opportunity. That is, the contemporaneous payoff is the dominant consideration, and continuation is second-order. The fundamental reason is that storage has no market power (with a small capacity); its actions do not change prices (significantly), so the contemporaneous arbitrage spread remains largely intact. Furthermore, the next trading interval is (probabilistically) a simple repetition of the current cycle, which renders continuation almost trivial.

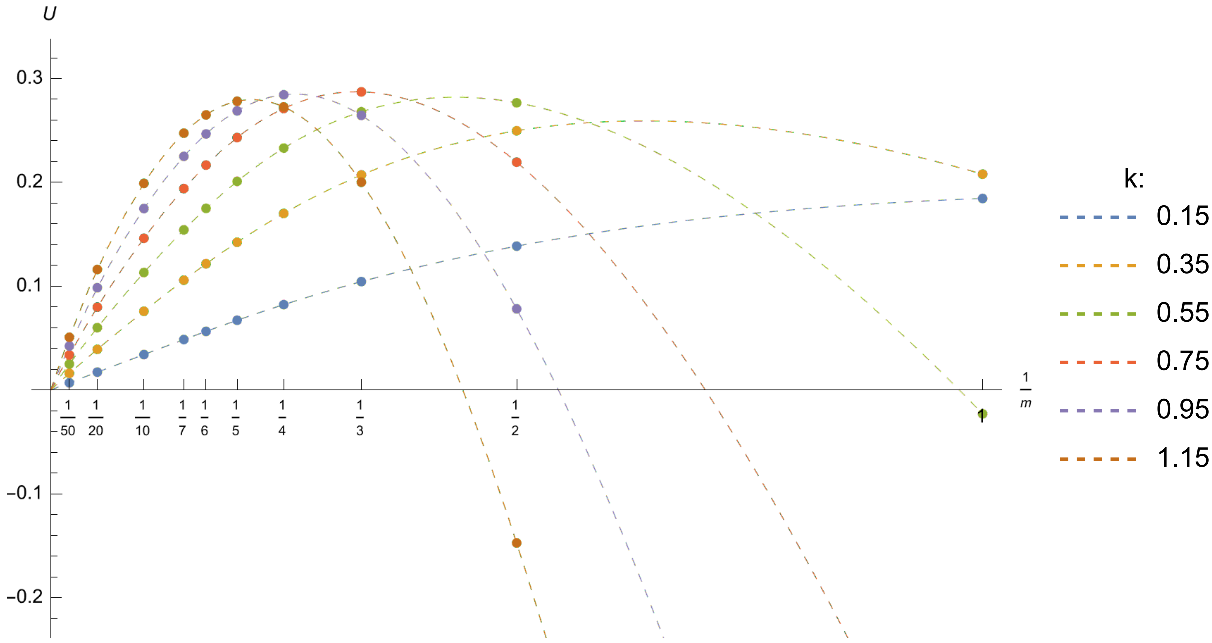

Figure 2.2: Payoff functions with constant quantities for different capacities k.

With a larger capacity, storage has market power – think of Snowy 2.0, for example; so it must adjust its behaviour and thus for two reasons. First, it must withhold quantities to preserve its contemporaneous arbitrage spread. The actions of a unit with market power can, and do, change prices – for the afficionados, this is the trade-off between marginal quantities and infra-marginal losses. For a unit that engages in arbitrage, these prices increase when buying and decrease when selling. That is, the spread can narrow rapidly. The second reason is completely specific to storage. With a large capacity, being stuck empty and unable to buy, or worse, full and unable to sell, is costly. This occurs with strictly positive probability: for example, the probability of 3 consecutive negative shocks is $1/8$ – so not trivial at all. To manage this risk, the only thing the unit can do is to delay it occurrence by reducing the *fraction of its capacity* traded each time (not necessarily the quantity). In other words, to manage the *continuation risk*, a large unit uses less of its capacity. Because the cost of this continuation risk increases with capacity, the larger the capacity, the smaller the capacity utilization. This costs becomes the dominant concern as capacity increases so that "very large" units become less profitable.

2.2.3 Extensions of the basic model

Our stark model can be enhanced to accommodate technical and institutional considerations of the NEM. We allow for more general stochastic processes to describe the random environment and let some generators be capacity-constrained – as all of them are in any electricity market.

The stochastic process: Markov chain. For the constant-quantity heuristic we relax the assumption of independent shocks and let the exogenous stochastic process follow a Markov chain with more or less persistence. In a persistent sequence, the same shock is repeated with high probability and so may give rise to a path with a string of negative chocks, for example $(-a, -a, -a, -a, ...)$. Conversely, a low persistent sequence is such that shocks revert every period: $-a, a, -a, a, ...$. We find that persistence of shocks is an important dimension too. It is intuitive that low persistence is best for the storage operator because is almost guarantees charging and discharging, and the unit is almost never stuck empty or full. So it helps in alleviating the continuation risk and allows larger quantities to be traded. Figure 2.3 depicts this constant-quantity heuristic for $m = 1, 2, 3$ – that is, for charging and discharging in 1,2 or 3 steps.

Again, for small capacity k (relative to the shock $a, -a$), charging and discharging in full dominates – and for the same reasons as exposed earlier. But as capacity increases, the precautionary motive triggered by the continuation risk takes over. Even for arbitrarily low persistence, that motive rapidly dominates. However the gain from $m = 2$ to $m = 3$ is modest – except for a very large capacity, in which case the market is flooded even in two steps ($m = 2$).

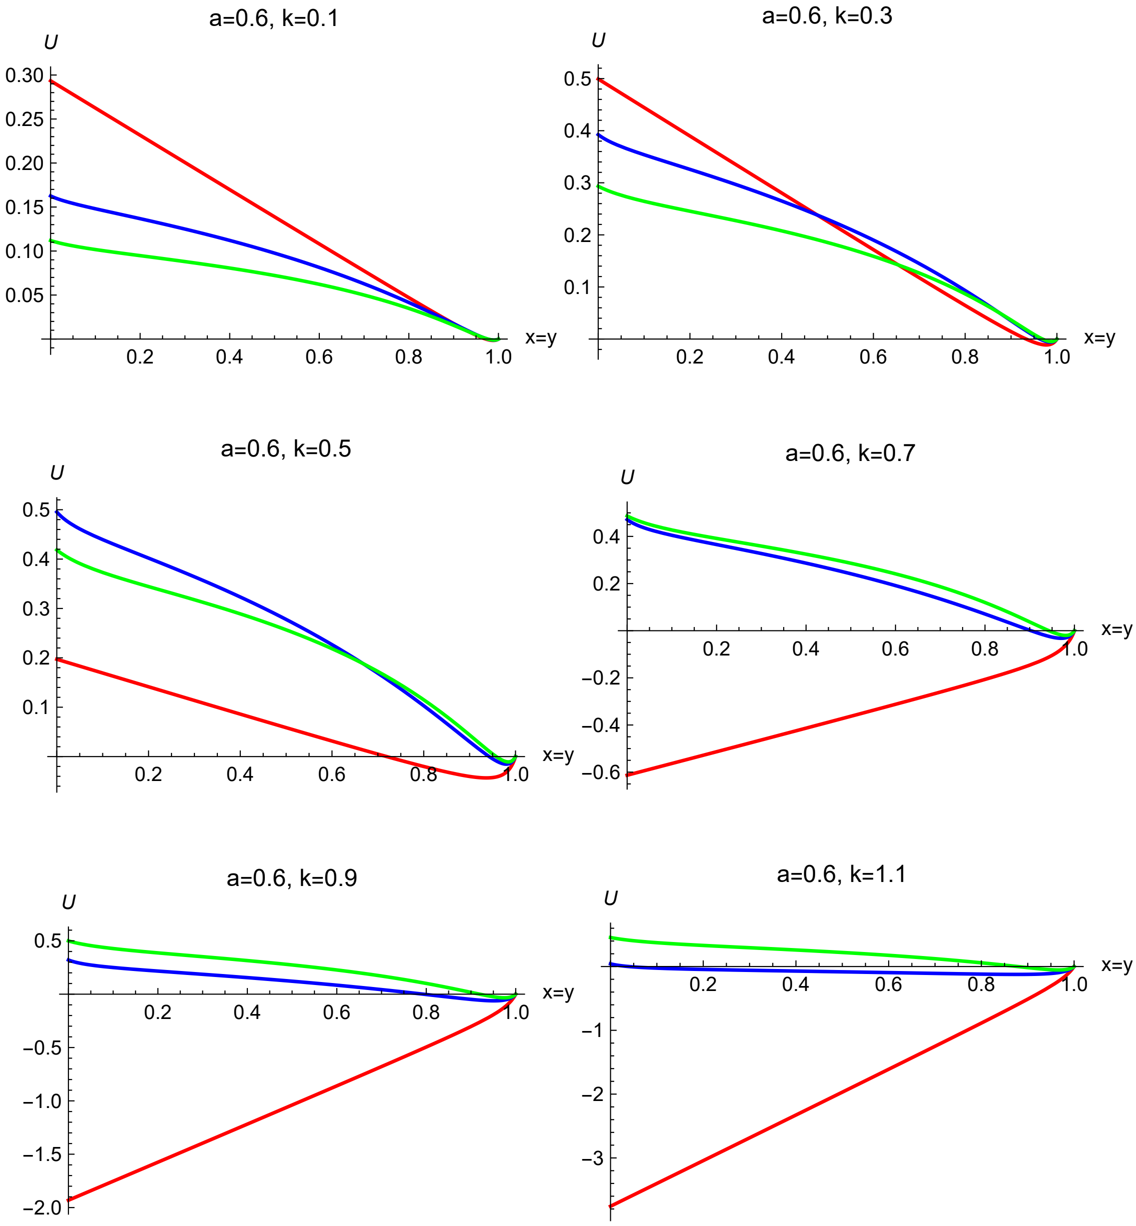

Figure 2.3: Payoffs under symmetric Markov shocks for divisible (green for $k/3$ and blue for $k/2$) and indivisible (red) capacities for different k and $a = 0.6$.

Capacity constraints. When capacity constraints start binding, prices rise and larger price differences invite (more) arbitrage. Indeed, as an increasing fraction of generators become constrained, the surplus to the storage operator increases for all values of the demand shock. This is quite intuitive.

Suppose now that the number of generators increases, however all the additional generators are capacity constrained – for example, gas plants. This adds to aggregate capacity, but in a qualified way. In this case the impact on the surplus of the storage operator depends on the interaction of the demand shock and the number of constrained generators. The reason is that constraints bind only when the demand shock is positive, that is, when storage sells. So increasing the number of constrained generators affects buying and selling energy at a different rate. More precisely, the impact of the shock a on the spread is linear, while the impact of the number of constrained generators on the spread is geometric; at some point it starts dominating.

2.2.4 Exclusionary equilibrium

As mentioned earlier, the equilibrium (the repeated Cournot-Nash 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. To do so, generators tacitly collude on the joint-maximising (monopoly) profit every period; if the storage operator still decides to buy, they revert immediately to the Cournot repetition we just exposed. This renders energy purchase prohibitively expensive: storage buys at the monopoly price and can only sell at a more competitive price.

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 (a complement), but it is a very serious competitor for gas generators (a substitute). 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 in the first place. Therefore, even if the equilibrium we construct is stark, some version of it may well unfold in the market place and slow the emergence of (independent) storage operators.

2.2.5 Competing storage units

Thus far our single storage unit has to manage the implications of its own market power and its own continuation risk. There is nothing like competition to discipline market power, and indeed, the discussion so far suggests that if market power is mitigated, then the continuation risk may be moot. We turn to this exact problem, however still letting storage operators engage in *imperfect* competition – just like generators engage in imperfect competition. There are now two storage units in the market; everything else remains as described earlier, except that we simplify our setting further and consider only the constant-quantity heuristic with $m = 1$. This is an important detail: a storage unit can only buy or sell its whole capacity k. This restriction affects the manner in which collusion can emerge.

It is well known that competing over time invites collusion. All cartels exist because of this time dimension, which allows for subsequent punishment (for example, price wars) if deviating from the agreement. This agreement need not be explicit, which is illegal. It can be sustained by equilibrium play and the threat of punishment if deviating from equilibrium. In real life, these deviations do occur, and sometimes mistakes are made; indeed, price wars are observed. In Australia, one of the best known tacit cartel is that of petrol stations, which seems to exist in all capital cities. It operates as a cycle, which any driver is surely aware of. There are many ways to sustain this kind of cooperation. We focus on equilibria that are simple in that they can be sustained by permanent reversal to a less attractive equilibrium; this is called "grim trigger". There are many others.

In a market that is essentially defined by its dynamic nature, and in which the arbitrage spread is so sensitive to quantities, collusion is a first-order – possibly the most important – consideration. We find that collusion is pervasive and can take multiple, complicated forms. There are three main reasons for this state of affairs:

- the dynamic nature of the market;

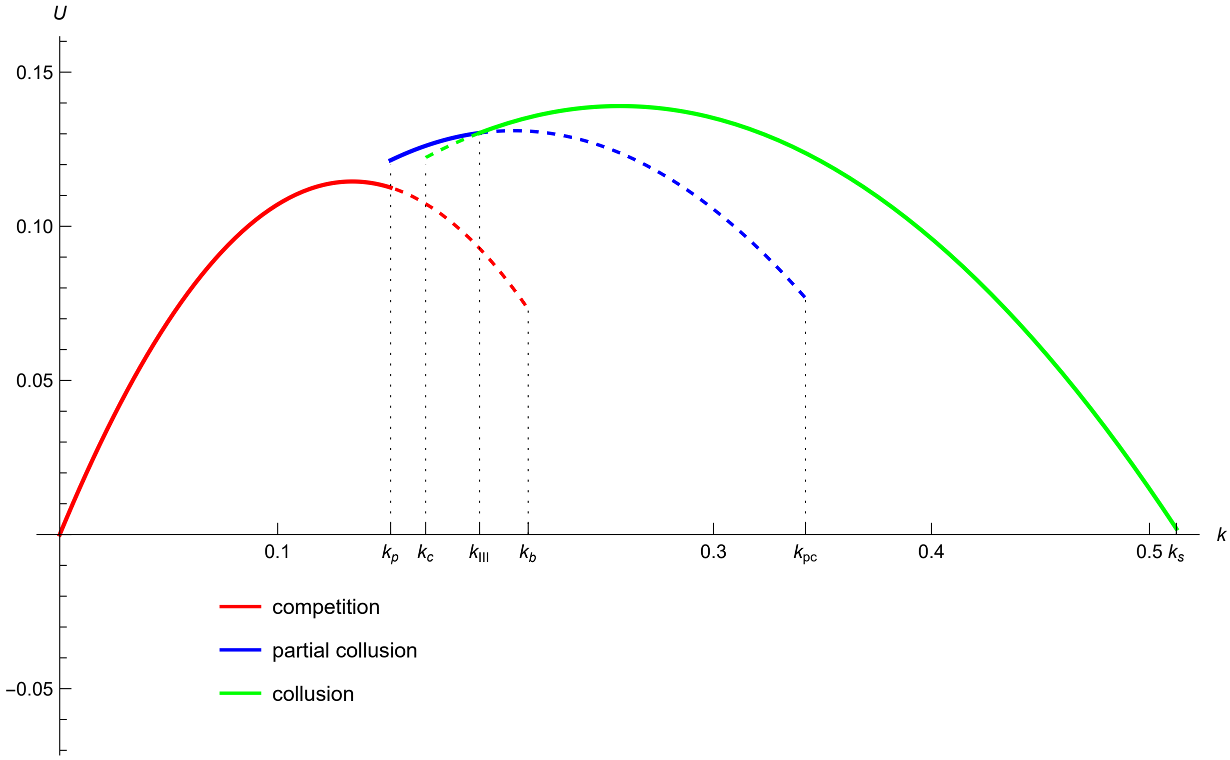

Figure 2.4: Equilibrium payoffs for different strategies for $n = 2$, $a = 0.6$, $\beta = \delta = 0.95$.

- the spread narrows sharply in quantities traded;
- storage must buy and sell, both of which are strategic. Hence collusion can exist with respect to one or the other action, or both.

In the class that we study, we find four types of equilibria, the existence of which depends on the capacity of the storage units.[2] At this point it is best to turn to Figure 2.4 that depicts the equilibrium payoffs.

The first equilibrium is quite natural: it is an equilibrium in which storage units compete alongside the other generators. The payoff from this equilibrium is depicted in red. For small enough capacities (k), this is the only equilibrium of interest. As capacities increase, this behaviour is too deleterious to the arbitrage spread: the profits of the "competitive equilibrium" decrease because the spread is progressively eroded by increasing quantities. But collusion emerges as a viable alternative – whereas it was irrelevant before. Collusion is not attractive for small capacities because the activity of the storage units has no discernible effect on the arbitrage spread.

We identify two kinds of collusive equilibria (in the class of grim-trigger equilibria). One is complete collusion: storage operators take turns to charge and discharge (so, only one at a time) in order to preserve their arbitrage surplus. This is an almost standard collusive behaviour, except for the fact that storage operators must collude on two actions. The payoffs are depicted in green. Firms flip a coin to determine who moves first, and then they take turns to charge and discharge; in this way they reduce the quantities traded and preserve the arbitrage spread. The other equilibrium is specific to storage; it is a form of *partial* collusion: they take turns to charge but sell simultaneously. This is novel and owes to the fact that two actions are required to complete a cycle (namely, charge and discharge). The converse is not an equilibrium because of the asymmetric effect of charging and discharging. Thanks to efficiency losses,

[2]There are many more equilibria; we focus on those yielding the higher payoffs to storage operators. Multiple equilibria can coexists: a collusive equilibrium is typically sustained by the threat of reversal to *another*, less pleasant equilibrium.

the quantities sold are lower than the quantities bought, so the benefit of collusion is greater when buying than when selling. If firms find it beneficial to collude on selling, they collude fully.

We see this equilibrium of partial collusion delivers higher payoffs than full collusion for intermediate capacities. When capacities are not too large, the arbitrage spread is not heavily eroded by the participation of both units. On the other hand, full collusion requires taking turns, which delays trade. Partial collusion is the compromise between these two forces; it is most profitable to collude on buying because the quantity effect is larger.

2.2.6 Takeaways

Trading energy over the long horizon in an uncertain environment is very sensitive to market power. That market power induces the usual price-quantity trade off that we know from introductory economics. It also gives rise to a new trade-off that addresses the *continuation risk*. That continuation risk is mitigated by withholding trades, that is, buying less and selling less. Hence capacity is underutilised. The larger the capacity – therefore, the more market power – the costlier is the continuation risk. For very large capacity, it becomes the dominant consideration, to the point where a greater capacity decreases surplus. Therefore, this should be internalised in the choice of capacity at the investment stage.

Similar considerations exist when storage units compete, even if they manifest differently. Large units manage the impact of their market power by engaging in collusive behaviour to reduce quantities traded. Because both buying and selling affect the total arbitrage spread, that spread narrows rapidly as total capacity increases. This renders collusion particularly attractive. As both buying and selling are required to generate surplus, collusion can occur on both actions (full collusion), or one action only (partial collusion). The asymmetry between buying and selling implies partial collusion exists in one direction only. Finally, for large enough capacities, *only* collusive equilibria can exist and only in a specific form.

Some conclusions for competition policy can be drawn from these results. Market surveillance will likely need to sharpen its focus to detect collusive behaviour and make a distinction between collusive strategies and operational mistakes. With the behaviour so sensitive to capacity, it is quite clear that large units are not desirable. A cap on the size of BESS installations is likely desirable, especially if there are no significant returns to scale. Our results also reinforce those of [2]: the AEMC and the ACCC should not allow thermal generators to also own storage units. However, storage entry may be assisted by the guaranteed access to a generation source like a solar farm – modulo the essential consideration of correlation (see [1]).

We are not yet in a position to discuss market design beyond stating that if storage operators play a dynamic strategy, then so should the market operator. This opens the question of "what is dynamic market clearing?"

2.3 The path forward

The results presented so far are preliminary: they are rooted in a simplified model that does not capture all characteristics of the NEM; for example, average demand is constant. They are also a first step to understand the behaviour of storage operators, which feeds into the problem of market design.

Currently we are working on models that complement this first pass. Specifically:-

- we model time-varying demand, which can accommodate deep storage. The goal is to understand when and how much to charge (and discharge) deal with the intra-day cycle when storage has market power and faces some uncertainty;

- we head toward understanding the cost-minimizing dispatch over the long horizon.

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.

Finally we are concerned about the ease with which a market can be manipulated with storage. Because a storage unit sells, it can engage in bidding and re-bidding just like any other generator – however at a much lower cost and much more precisely. In addition, because a storage unit *buys*, it can also contribute to increase demand in critical times, which benefits to all infra-marginal units being dispatched. This is another reason to be more than skeptical to let owners of large fleets of generators to also own storage.

3 Summary and other works

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, whether a lone operator or competing units.

In the case of a single operator, a comparatively large storage unit uses a small fraction of its capacity (always less than 50%, and much less the large it is). This strategy is rooted in the market power of the storage unit, and is motivated by two considerations: *(a)* storage acts like a monopolist on its own arbitrage spread; and *(b)* it faces a material continuation risk, which requires prudence.

Competing operators find collusion attractive as a solution to manage their own price impact. This collusion may be partial (when charging only) or complete (on charging and discharging). In some case, collusion is the only equilibrium; that is, it is the only way for storage units to generate a positive return.

While only preliminary, this work already (cautiously) 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), as well as finding the cost-minimizing allocation. Once that is understood one can think of market design: these are the rules of the market game that can implement the cost-minimizing allocation. Finally we plan to also study the scope of market manipulation with storage.

In other works, we study 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 also 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.

References

[1] Andres-Cerezo, A. and Fabra, Natalia. "Storage and Renewables Energies: Friends of Foes". In: *Working Paper* (2023).

[2] Andres-Cerezo, A. and Fabra, Natalia. "Storing power: market structure matters". In: *Rand Journal of Economics* 54.1 (2023), pp. 3–53.

[3] Butters, R. Andrew, Dorsey, Jackson, and Gowrisankaran, Gautam. "Soaking Up the Sun: Battery Investment, Renewable Energy, and Market Equilibrium." In: *R&R Econometrica* (Working Paper).

[4] Glebkin, Sergei, Malamud, Semyon, and Teguia, Alberto. "Illiquidity and higher cumulants". In: *The Review of Financial Studies* (Forthcoming).

[5] Harrison, J. M. and Taylor, Alison. "Optimal control of a Brownian storage system". In: *Stochastic processes and their applications* 6 (1978), pp. 179–194.

[6] Karaduman, Omer. "The economics of grid-scale storage". In: *mimeo* Stanford University (2020).

[7] Shapley, Lloyd. "Stochastic games". In: *Proceedings of the National Academy of Sciences* 39 (1953), pp. 1095–1100.

[8] Vayanos, Dimitri. "Strategic trading and welfare in a dynamic market". In: *The Review of Economic Studies* 66 (1999), pp. 219–254.