Addressing the risk of algorithmic collusion

Box 1: Chapter overview

AI techniques are extensively used in trading and commodity markets, taking real-time data to inform trading decisions that continually adapt to changing market conditions.

Pricing algorithms are already a pervasive feature in the NEM. Over time, these algorithms will employ AI and machine learning to a greater extent.

Algorithmic pricing and AI could provide a range of benefits for the NEM:

- It should increase market efficiency and allow decisions to be made more quickly, cheaply, and consistently
- Auto-bidding software can reduce barriers to entry for generators by reducing reliance on people to make complicated bidding decisions
- It could be used to support enhanced decision-making beyond what AEMO already achieves.

At the same time, the process of ‘self-learning’ by algorithms may ultimately lead AI to implement bidding or pricing strategies that closely resemble collusion, leading to higher prices for consumers.[1]

Such self-learning capability could also result in bidding strategies that seek to maximise profits by pushing the power system to the edges of its operating envelope. Furthermore, the lack of transparency of decisions by AI algorithms can make monitoring and enforcement action problematic.

Policymakers therefore face a difficult ‘trilemma’ when managing the risks of algorithmic collusion.

1. Waiting for a collusive outcome to occur is naturally not an acceptable policy response.
2. Detection of algorithmic collusion, even if possible, may be prohibitively expensive.
3. Restricting the use of AI or pricing algorithms in totality is also unattractive, given its potential benefits.

Key takeaways and recommendations:

The NEM is a market that is susceptible to collusion, and the risks have increased in recent years due to two inter-related trends:

- An increase in information available to participants that has the potential to reveal information about competitors’ strategies;
- The development of AI tools which, in simple terms, reduce the costs of sustaining collusion, and potentially increase the costs of detecting collusion.

While there are a range of potential regulatory responses, care needs to be taken to not unduly restrict the use of AI given its use has the potential to reduce costs for consumers.

The information that is provided to the market should be the minimum required to promote more efficient bidding behaviour. Operators and policymakers should avoid inadvertently publishing information that could be used to predict a generator’s current or future pricing strategies, or to reveal a competitor’s ability to store or generate electricity.

Promoting more responsive consumer demand also reduces the risk of collusion in the NEM.

To progress this work, we recommend the Commission establish a working group with the Australian Competition & Consumer Commission (ACCC) and Australian Energy Regulator (AER) to develop and identify appropriate responses to address the risks of collusion.

Source: [1] OECD, *Algorithms and collusion: competition policy in the digital age*, September 2017. Available here.

This chapter assesses the risk of algorithmic collusion in the NEM, discussing:

1. Why the adoption of AI has created the risk of algorithmic collusion.
2. Why the NEM is a market increasingly vulnerable to algorithmic collusion.

While it is difficult to develop a perfect policy response to the risk of algorithmic collusion in the NEM, policy should focus on understanding the optimal level of information provided to the market, reducing barriers to entry, and promoting a 'two-sided' market.[16]

1 AI has created the risk of algorithmic collusion

This section explores the following questions:

1. What is collusion?
2. What market characteristics support collusion?
3. What is algorithmic collusion?
4. How does algorithmic collusion change when considering AI?

1.1 What is collusion?

Collusion is a form of anti-competitive conduct in which parties, who ordinarily compete against each other, coordinate to increase profits. There are two types of collusion:

1. Explicit collusion
2. Tacit collusion.

Explicit collusion occurs where anti-competitive behaviour is maintained with written or oral agreements — in other words, cartel conduct. It is also possible for explicit collusion to occur where there has been a 'meeting of the minds' or an understanding between the parties.

The most direct way for firms to achieve this is to interact directly and agree to fix prices or output quantities, coordinate bid-rigging behaviour,[17]or allocate markets to particular parties. For example, consider two childcare centres operating in a country town. Explicitly collusive conduct would arise if the two firms decided to increase prices by five per cent after discussing the change in a meeting.

Conversely, tacit collusion refers to forms of "anti-competitive coordination which can be achieved without any need for explicit agreement, but which competitors are able to maintain by recognising their mutual interdependence".[18]

Tacit collusion often arises over time, as firms in the market observe and 'learn' to co-ordinate their behaviour to achieve a given goal, for example, to synchronise price cycles, avoid price cutting, or engage in non-price forms of competition.

Knittel and Stango (2003)[19] provide a practical example of tacit collusion, finding that credit card companies in the US tended to move prices towards non-binding price ceilings over time.[20] In this case, price coordination was facilitated by the non-binding price ceiling acting as a focal point.

16 A market is 'two-sided' when buyers and sellers use a platform, broker, or other intermediary to exchange offers to buy or sell goods and services'. International Center for Law and Economics, *Issue Spotlight: Two-Sided Markets,* November 2022. Available here.

17 'Bid rigging, also known as collusive tendering, happens when suppliers discuss and agree among themselves who should win a tender, and at what price'. ACCC, *Cartels*. Available here.

18 OECD, *Algorithms and collusion: competition policy in the digital age*, 2017, p. 19. Available here

19 C. Knittel and Stango, V. *Price ceilings as focal points for tacit collusion: Evidence from credit cards,* American Economic Review, 93:1703–1729, 2003. Available here.

20 A non-binding price ceiling is a price ceiling that is above the equilibrium price in a market.

1.2 What market characteristics support collusion?

For collusion to be sustainable in a market, the incentive for participants to engage in collusion must be sufficiently strong, *and* the 'reward' to a participant for deviating away from collusive prices must be sufficiently weak. The markets with the following characteristics have the highest risk of collusion:

- **Homogeneous (similar) goods:** in markets with homogeneous goods, consumer decisions are based almost exclusively on price. Therefore, incremental benefits to producers from colluding are larger.
- **Significant barriers to entry and market power:** a small number of firms and high barriers to entry increase the potential benefits, and ability, to adopt higher prices.
- **Limited buyer power and inelastic demand:** if there is limited ability for consumers to influence the price or quantity in the market — for example, through bargaining, switching, or 'playing' producers off one another — markets are less able to avoid collusive pricing.
- **Frequent interaction and competitor strategy visibility:** markets with frequent interaction and transparency of competitor strategies have a greater risk for collusive outcomes as it is easier for participants to monitor, identify, and retaliate against those who deviate from the common policy.

Importantly, these risk characteristics are common for markets with and without algorithms that set prices. The main difference is the scope for pricing algorithms to more efficiently use information to understand the current and future strategies of competitors. To that end, pricing algorithms "can amplify the risks and make tacit collusion a more frequent market outcome".[21]

1.3 What is algorithmic collusion?

Algorithmic collusion is a form of tacit collusion where the collusive outcome arises as a result of pricing algorithms.

Pricing algorithms are used extensively across a variety of industries in order to automate firm pricing decisions. Most of the pricing strategies employed are simple, such as pricing at 1.1 times the lowest competitor's price, however, a growing number are more complicated and take advantage of machine learning or other AI methods.

In recent years, several examples of algorithmic collusion have reached headlines. Among them is a simple example of how two companies priced the textbook *The Making of a Fly* on Amazon:

- One seller adopted an algorithmic pricing strategy that is 1.27 times the average price of competitors.
- Another seller adopted a strategy that set a price equal to 0.9983 times the lowest price of any competitor.
- This led to the price of *The Making of a Fly* spiralling upwards, eventually reaching a price of USD23 million in 2011.[22]

Algorithms increase the risk of tacit collusion. Uncertainty among competitors about future prices and supply choices is a key competitive uncertainty that creates competition and innovation. This

21 OECD, *Algorithms and collusion: competition policy in the digital age*, 2017, p. 34. Available here.

22 J. Chan, *Algorithmic collusion and Australian competition law: trouble ahead for the National Electricity Market,* UNSW Law Journal 44(4), 2021, p. 1380. Available here.

uncertainty is reduced as the pricing algorithm itself contains the firm's pricing strategies. In turn, this increases the risk of coordination to gain market advantage.

Algorithmic pricing also increases the predictability of a firm's actions, and, through repeated interactions with others, its strategy could be 'decoded' or otherwise better anticipated. Furthermore, algorithmic pricing acts as a commitment device[23] to these strategies, making it more likely that competitors will 'trust' the algorithm's behaviour.

There is a limited body of literature that examines the empirical effects of algorithmic collusion because of the difficulty in its detection and its recency as a field of interest. Despite this, a consensus is emerging that algorithmic pricing increases the risk of anti-competitive, collusive outcomes and, subsequently, higher prices.

- Assad et al. (2020) found empirical evidence of tacitly collusive strategies once algorithmic pricing was adopted in the German retail gasoline market. They found that the adoption of algorithmic pricing led to increases in mean petrol station margins of 9%, and where the petrol stations were in ZIP-code duopolies, of 29%.[24]
- There is also evidence that tacit collusion more broadly leads to higher prices. Byrne & de Roos (2019) found that tacit collusion among petrol retailers in Perth between 2010 and 2015 increased profit margins by up to 57%. This arose because of a leader-follower dynamic where BP led price increases to facilitate a "mutual understanding among rivals of a new, profit enhancing focal pricing structure".[25]
- The same study also show that the breadth and depth of the information that was made available to competitors (such as day-ahead pricing for all stations) facilitated tacitly collusive outcomes.

Emerging literature shows that providing more information to the market can lead to higher prices for consumers,[26] particularly when consumers' decision-making is sticky. AI increases the likelihood and worsens the impact of such an outcome, given the role of AI in utilising information to better understand consumers' willingness to pay and other competitors' likely reactions to changes in price or quantity.

1.4 How does AI reduce the barriers to algorithmic collusion?

The use of AI and learning algorithms reduce the barriers to collusion. This is achieved through:

- **Consistent pricing decisions:** as AI and learning algorithms are constructed using code (which is static), their use by market participants acts as a commitment device to a certain pricing strategy. This means that there is greater visibility over firms' pricing strategies, making it easier to sustain collusion with a greater number of competitors.
- **Information integration:** AI and learning algorithms can better respond to large volumes of dynamic information (such as new market conditions or new variables) as opposed to conventional pricing algorithms. Specifically, they are able to iterate their strategies based on available information, and are therefore able to learn to charge higher-than-normal prices without communicating with one another. In addition, they are able to quickly identify and punish any actors who deviate away from the collusive strategy.

23 'Commitment devices attempt to enforce people's voluntarily imposed restrictions until they have accomplished their goals, or their voluntarily imposed penalties for failing to accomplish their goals'. Rogers et al., *Commitment Devices*, Harvard University 'Viewpoints, 2014. Available here.

24 S. Assad, et al., *Algorithmic pricing and competition: empirical evidence from the German retail gasoline market*, CESifo Working Paper No. 8521, 2020. Available here.

25 D. Byrne and de Roos, N., *Learning to coordinate: a study in retail gasoline*, American Economic Review, 109:591–619, 2019. Available here.

26 See, for example, D. Byrne, et al. *Informed sources and the role of platforms for facilitating anti-competitive communication*, unpublished working paper, 2022. Available here.

- **Competitor strategy identification:** it is easier to test competitors' strategies (and the market more broadly) when employing AI and learning algorithms. Firms are able to rapidly bid and rebid prices in order to examine their behaviour. In turn, this makes it easier for firms to reach a supracompetitive outcome.[27]

In this way, AI increases the risks posed by conventional algorithmic pricing in realising a collusive outcome.

For example, consider again the textbook *The Making of a Fly*. Had the sellers used a machine learning algorithm, the algorithm may have tested slowly increasing prices, observed the marginal changes in consumer demand, and gradually discovered the price that jointly maximises profit with the alternative seller.

Academic literature provides some evidence of similar behaviour. Calvano et al. (2020) found in a simulated oligopoly environment that "Q-learning pricing algorithms systematically learn to collude"[28] and that the "algorithms consistently learn to charge supracompetitive prices, without communicating with one another".[29]

Algorithmic collusion via machine-learning algorithms may be more difficult to detect and punish:

- **It is not clear how algorithmic collusion would be detected**: the literature on how to effectively detect algorithmic collusion is still emerging, with limited detection methods identified.
- **Detection using algorithm inputs would not be sufficient:** a natural starting point for detection would be to consider what information an algorithm is using. However, Courthoud[30] shows that the inspection of the algorithm's inputs might not be sufficient in order to determine whether an algorithm can learn collusive strategies.
- **Detection methods may be impractical to enforce:** another proposed detection method requires that algorithms be retrained. However, this requires firms to effectively 'reset' their algorithms and their learnt behaviour. Noting that this training and learning process is often time-consuming and expensive, "one cannot ask firms to reset their algorithms in order to check for collusion".[31]
- **Detection methods may be resource intensive for regulators**: although there are some identified methods that rely on observational data only,[32] detection methods work by exploiting the fact that learning algorithms explore alternative strategies. This allows the observer to understand behaviour in counterfactual scenarios through experiments. However, the monitoring and evaluation of a significant volume of observed data would likely be resource intensive for regulators.

Whilst it is possible to detect algorithmic collusion (albeit with some difficulty), punishing it is likewise difficult. In particular, while the NER prohibits false or misleading bids, it allows subjective intent for the purposes of rebids – that is, generators may re-bid on the basis of their subjectively held expectations or beliefs (including if these are neither likely nor reasonable).

This creates a broad scope for generators to shift offered quantities in the market before the dispatch period, and for other generators to adjust their quantities in response, thus increasing the risks of collusion. In addition, monitoring all market actors for algorithmic collusion may be

27 A supracompetitive pricing outcome is one in which prices are higher than prices in normal competition.

28 Q-learning is a reinforcement learning algorithm that learns the value of an action in a state space, allowing the algorithm to iterate and maximise expected total reward.

29 E. Calvano, et al., *Artificial intelligence, algorithmic pricing, and collusion*, American Economic Review, 110: 3267–3297, 2020. Available here.

30 M. Courthoud, *Algorithmic collusion in online marketplaces*, working paper, 2021, p.21. Available here.

31 Ibid. p.18

32 Ibid.

difficult and costly, particularly considering the volume and velocity of data generated in the bid-rebid process.

2 The NEM is a market increasingly vulnerable to algorithmic collusion

The characteristics of the NEM make it particularly vulnerable to algorithmic collusion.

Box 2 serves as theoretical support to the arguments made in this chapter by providing an overview of how generators in the NEM bid (offer) electricity volumes to the market. The reader who is already familiar with this process may skip Box 2.

Box 2: Understanding generator bids in the NEM

For most Australians, the electricity that they consume is generated and sold through the NEM. The NEM determines how much electricity is produced by each competing generator using a real-time, single-price auction approach. Each competitor submits real-time bids* to the market operator (AEMO) to provide volumes of electricity at increasingly high prices.

Specifically, generators bid ten separate price and quantity pairs that represent their willingness to provide specified volumes of power at different prices. AEMO uses this information to meet the forecast level of customer demand. Whilst the 10 bid price bands are fixed for the day at 'gate close' the night before, generators can rebid quantities in each band anytime virtually up to the start of the dispatch calculation.

The dispatch process aggregates and stacks bids from cheapest to most expensive to determine the marginal generator required to meet electricity demand. It also determines opportunities to import/export electricity between regions within transmission constraints and power system security requirements. Consequently, the market operator determines dispatch instructions for every participant and a clearing price for each region every five minutes – that is, the single price that is paid to all participants whose bids are accepted and dispatched by the operator.

The NEM is designed to encourage each generator to bid the minimum price that they would be willing to accept to supply that unit of electricity. Provided a generator cannot predict – or collude to set – the price and quantity offered by other competitors, this approach operates efficiently – and in the long-term interest of consumers. This is because if they bid their lowest price and the market clearing price is higher, they still earn the higher price. But, if the generator bids above their lowest price, and the market clearing price is between their bid and their lowest price, they are worse off as they may not be dispatched at a time when they could have recovered their minimum price.

This market design reduces the risks of collusion compared to a pay-what-you-bid approach, where competitors have a stronger incentive to collude in order to assess the others' probable bids (Goldstein, 1962).

Source: AEMC.

Note: *While these are technically offers, they are described as bids in the NEM.

Note: H. Goldstein, *The Friedman proposal for auctioning Treasury Bills*, Journal of Political Economy, 70(4), 1962. Available here.

The characteristics that make the NEM a market susceptible to collusive behaviour are:

- **Homogeneous good:** electricity is a homogeneous good. A megawatt-hour from one generator is equivalent to one megawatt-hour from another.
- **Significant barriers to entry and market power:** there are large barriers to entry in the NEM in terms of cost and technical constraints.
  - First, there are large up-front fixed costs in order to construct generating units (once constructed, variable costs are near-zero for renewable generation, and are largely a function of input fuel costs for thermal generators).

    - Second, there is scarcity of suitable land for generator development, as generators typically are required to be co-located near existing transmission infrastructure (or, in the alternate, face the costs of developing this infrastructure themselves).
    - Third, congestion of transmission lines (due to capacity constraints) can also create intermittent periods of regional market power.
    - Finally, for generators, the connections process face technical and time constraints (e.g. it requires detailed modelling) that cause delays. Project timelines from inception to energisation can take upwards of three years. As a result, the largest ten generating companies accounted for nearly 80% of NEM-wide generation in 2023.[33]
- **Limited buyer power and inelastic demand:** today, the NEM is a one-sided market where there is little opportunity for wholesale consumers to exert buyer power (that is, in the wholesale market). Similarly, electricity demand is relatively inelastic – consumers will demand electricity as needed and are historically limited in their ability to be price-responsive. However, this might change with a variety of current and future reforms seeking to integrate consumer energy resources (CER) into the market.[34]
- **Frequent interaction and competitor strategy visibility:** competitors bid for electricity every five minutes, and can submit re-bids at any point in time before each settlement period. They have access to a large amount of current and future information provided by the market operator. As explained below, these factors may increase the risk of supracompetitive prices and collusion.

Algorithmic collusion in the NEM could lead to higher prices for energy consumers. Box 3 presents an illustrative example of how bidding behaviour might adjust to support a collusive outcome.

Box 3: Risks for higher prices due to algorithmic collusion

Under a competitive market outcome, each individual generator would bid in a somewhat predictable pattern, with three clusters of bids:

1. The first tends towards the market price floor, to ensure that they are dispatched in the next period (for example, if they are a baseload generator).
2. The second cluster is typically centred around the generators' short-run marginal cost and is typically aligned with their financial market volume.
3. The third cluster tends towards the market price cap, to ensure that they are able to take advantage of any extreme price events for volumes not covered by financial instruments.

Provided algorithms and auto-bidding software can utilise the information available to them from AEMO and other sources, generators may be able to co-ordinate (including tacitly) to shift the marginal bid to a higher price band, so that all generators can earn a higher price.

This may occur as a result of the group of coordinating generators shifting small amounts of quantities away from the lower bid bands at the intraday level, or by shifting prices at the inter-day level.

The former case is of particular concern as it is an extremely low-risk option for generators. Indeed, each coordinating generator is only risking a very small percentage of their total offered quantity for a much higher payoff.

In other words, the losses of the coordinating generators given a competitor deviation are small,

33 Internal calculations.

34 For example, the AEMC is currently evaluating the design of a process to integrate price-responsive resources into the NEM dispatch process in the rule change 'Integrating price-responsive resources into the NEM'. More information is available here.

meaning that a collusive outcome, and thereby higher prices, are more likely.

Figure 2.1 shows a competitive market equilibrium, whereas Figure 2.2 shows a market price increase due to the quantity in the first seven bid bands being withheld by ten per cent.

Figure 2.1: An illustrative competitive market equilibrium

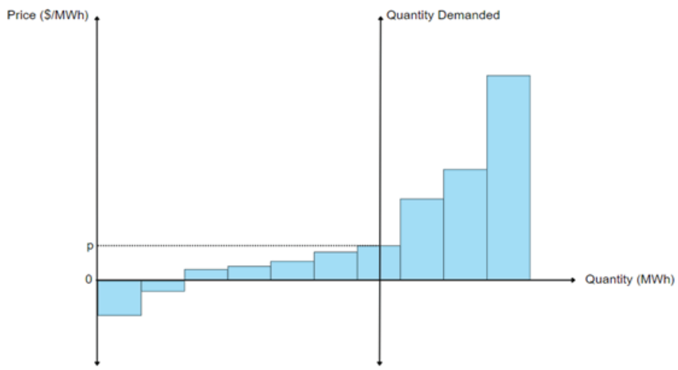

Source: AEMC analysis

Note: This figure provides an illustrative "bid stack" under a situation where each bid is set at the minimum price that a generator would be willing to supply that generation.

Figure 2.2: An illustrative equilibrium with algorithmic collusion

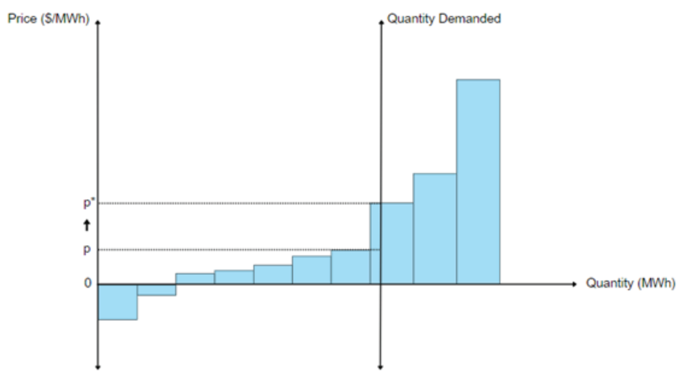

Source: AEMC analysis

Note: This figure provides an illustrative "bid stack", and the impact on the market clearing price if each generator (or algorithm) colludes to withhold a small amount of supply.

Source: AEMC.

We identified a number of recent structural changes that may have worsened the risks of algorithmic collusion in the NEM. In recent years, several changes have increased the frequency of interaction between competitors and information provided to competitors about current and future market conditions. The list below provides further details on these structural changes:

- **Increasing interaction:** in 2021, the NEM moved from a 30-minute settlement period to a 5-minute settlement period. This increased the interaction between competitors from 48 to 288 times per day.
- **Increased competitor strategy visibility:** As part of the dispatch process AEMO also publishes all network constraint data for every 5-minute pre-dispatch period.[35] Over the last 15 years an increasing number of participants and now some commercial software companies offer services that 'unpack' this information such that a participant can virtually see their

35 https://www.aemo.com.au/energy-systems/electricity/national-electricity-market-nem/system-operations/congestion-information-resource

competitors dispatch targets for the rest of the trading day. Pre-dispatch is run every 5 minutes, allowing participants to see the effect on their competitors' dispatch from each rebid and identify whether they are the marginal generator.

- **Increased competitor strategy visibility:** AEMO provides access – on a fee-for-service basis – to an offline version of the NEM Dispatch Engine (NEMDE),[36] which could provide more visibility over the potential impacts of different pricing strategies.

At the same time, there has been increased adoption of pricing algorithms, or auto-bidding software, in the NEM. An industry report found that, in 2021, 15 different bidding technology providers were providing "autonomous determination and submission of bids" into the NEM,[37]with the most prominent example AMS (now Mosaic, after being acquired by Fluence in 2020). Mosaic works by forecasting future prices across both energy and FCAS markets using machine learning across thousands of variables.[38] It then calculates the optimal bids across all markets for dispatch periods from the next period, to the next day. Finally, it also creates market-compliant bid and rebid files; although our understanding is that these files, while technically compliant, are typically relatively simplistic and contain limited information.

Although the NEM is vulnerable to collusion, the wholesale market's close relationship with the contract market should, on balance, dampen these risks.

In the short term, the contracts (or hedges) that generators and retailers enter into reduce their exposure to the wholesale market spot price. That is, the strike price of cap contracts limits exposure to the wholesale spot price above the strike price, and swap contracts limit spot price exposure entirely for the contracted quantity. This reduces the short-term incentive to collude as the returns to collusion are limited to only the uncontracted capacity and the cap contracted capacity (if the price is strictly less than the cap).

In the longer term, the higher prices in the wholesale market sustained through collusion may create a vicious cycle where price expectations increase, leading to higher contract prices and lower liquidity. This would disincentivise the purchase of additional contracts, feeding back into the cycle, and reducing the effectiveness of the contract market in moderating the returns to collusion. Over time, as contract positions expire, these effects on incentives fade. This means that the contract market may do less to dampen the risks of collusion in the long term. In addition, participants who have a sufficiently low discount rate or who are otherwise sufficiently patient will still collude in the long run.

On balance, it is likely that the contract market will still have some dampening effect on the risks of collusion, as participants value the revenue certainty that contracts offer.

3 Policy considerations for algorithmic collusion in the NEM

This section outlines broad considerations in responding to the risk of algorithmic collusion in the NEM. These considerations include:

- Measures to detect and monitor for algorithmic collusion, which could involve requiring market participants to report information to describe their current and future use of algorithms
- Measures that require participants to promote algorithmic transparency and accountability

36 https://aemo.com.au/en/energy-systems/market-it-systems/electricity-system-guides/nemde-queue-service

37 Generator Insights 2021, Appendix 22. Summarised here.

38 https://fluenceenergy.com/energy-bidding-software-australia-nem/

- Limiting particular types of behaviour in the market – which could apply to some or all participants
- Changing how much information is provided to the market, and when, to limit the ability of competitors to collude
- Reforms to competition law and enforcement frameworks to better capture (and punish) algorithmic collusion
- Reducing barriers to entry, changes to merger control rules, and more generally, to encourage more responsive consumer demand
- Measures that are adopted on a precautionary basis, or after a risk has materialised.

Enacting each of the above measures comes with trade-offs. Regulating AI is difficult because of its unique interactions with the market. It is difficult to directly regulate or legislate against algorithmic collusion principally because there are significant technological and cost barriers to monitoring or detection. In addition, regulation may reduce competition by increasing the barriers to the use of new technologies. Finally, regulation can stifle innovation and productivity by reducing the profit incentive to adopt new technologies.

3.1 Approaches for monitoring or detecting algorithmic collusion have limitations

One option to address algorithmic collusion is to use behavioural screens that flag potential collusion using available data (OECD, 2021).[39] For example, patterns in price and quantity offers (that is, bidding activity) can be analysed by regulators to identify potential collusion. This analysis could be extended with additional information about the behaviour of AI, data inputs or each algorithm's source code.[40]

Indeed, the literature has begun to propose sophisticated methods to detect algorithmic collusion.[41]However, this is an emergent area of research and, at present, can only be applied ex-post with significant data availability. To monitor algorithmic collusion, Chan (2020) proposed expanding the ACCC's information-gathering powers under the Competition and Consumer Act 2010 (Cth) (CCA).[42] Specifically, generators would be required to notify the ACCC about their future use of AI, in order to account for the fact that this use can change over time.

However, monitoring and detection have limitations when applied to AI:

- **Effectively monitoring the use of AI could be costly.** For the user of AI, it involves tracking the behaviour of a complex, changing system, and providing the required information to a regulator. This poses challenges to regulators, as they would need to keep up with the ever-evolving technical expertise (knowledge from computer science, engineering, finance, and economics) to detect and assess algorithms.
- **The market dynamics in the NEM are also complex to assess.** The sheer volume of data required for this level of analysis is potentially prohibitive because of the frequent market interactions, particularly with some renewable and storage participants rebidding multiple times in every 5-minute dispatch interval. As an illustrative example, a recent market analysis has showed that there were well over 1 million bids during October 2021 in the NEM.[43]

39 OECD, *OECD business and finance outlook 2021: AI in business and finance, chapter 4, competition and AI*, 2021. Available here.

40 UK Competition & Markets Authority, *Algorithms: How they can reduce competition and harm consumers*, 2021. Available here.

41 See, for example, M. Courthoud, *Algorithmic collusion in online marketplaces*, working paper, 2021. Available here.

42 J. Chan, *Algorithmic collusion and Australian competition law: trouble ahead for the National Electricity Market,* UNSW Law Journal 44(4), 2021. Available here.

43 Generator Insights 2021. Summarised here.

- **Monitoring and data collection can reduce innovation.** Suppose agents are worried about future punishment if a change to the use of AI unexpectedly results in an actual or perceived anti-competitive outcome. In other words, there is now a significant disincentive to innovate across the AI domain for fear of regulator punishment.

As outlined below, the NER already contains a number of requirements that protect against false or misleading bidding behaviour. As a potential next step, monitoring could be facilitated if the bids had to specify if a participant was using AI and/or an Autorebidder. This would allow for analysis of bidding, and rebidding, behaviour across the different types of market participants.

The false and misleading prohibitions under the NER require the participant to keep contemporaneous records that should record the information on which decisions were made, the basis for the strategy and who authorised its activation and use.[44]

3.2 Promoting transparency and accountability can set expectations for AI use

Beyond providing information to the regulator or operator, market participants could have obligations that limit how they use AI.

For example, the US Public Policy Council of the Association for Computing Machinery (USACM) has proposed seven principles for algorithmic transparency and accountability.[45]

More broadly, voluntary (or mandatory) codes of conduct, as well as Board and CEO attestations, have been used in a variety of contexts to reduce the risk of bad behaviour. Doing so could set guardrails and expectations for the use of AI technology in the market.

The NEM already promotes accountability over bidding:

- Clause 4.9.2(d) requires each participant to always have appropriate personnel (not machines) available to respond immediately to dispatch instructions issued by AEMO (be they electronic, written, or oral).
- Clause 4.9.8 (b) requires participants to always have bids with which they can comply.
- Clause 4.11.3 requires generators to provide contact details for those personnel. Such measures arguably reduce the risks of pricing strategies being solely the responsibility of an auto-bidder.
- Clause 3.8.22A states that a bid or rebid cannot be changed unless the market participant becomes aware of a change in the material conditions, and when a material change does occur that a rebid must be made as soon as practicable.

However, there are limits to how accountable owners of algorithms can be made for the actions of AI. Full accountability would require that someone be able to explain why a particular bidding outcome was produced, but as noted by the OECD, “that might be an impossible task when machine learning systems have made autonomous decisions that have not been instructed by anyone”.[46] This is further complicated by the fact that AI algorithms are typically ‘black box’ systems, where there is no visibility of the mechanics used to transform inputs to outputs.

Furthermore, anecdotally we have heard that while some bidding teams comprise a number of people during the day, they are often a single trader - or no traders if they are using an auto-bidder - overnight. Although this is a breach of the rules and AER’s technical compliance guidelines, the

44 NER Clause 3.89.22A.

45 USACM, *Statement on algorithmic transparency and accountability*, 2017. Available here.

46 OECD, O*ECD business and finance outlook 2021: AI in business and finance, chapter 6, Algorithms and market regulation*, p. 48, 2021. Available here.

clauses only promote accountability if they are paired with an effective monitoring regime that periodically verifies whether compliance is occurring.

There are further limits to promoting transparency in the use of AI. As noted by the ACCC, developing an algorithm can involve significant upfront investment. Therefore, "opacity on the working of these algorithms prevents potential competitors from copying and otherwise free-riding from an online marketplace's investment; allowing such free-riding and/or sellers to game the algorithm would clearly result in a poor outcome for consumers and businesses."[47]

3.3 The costs and benefits of market behaviour restrictions must remain balanced

The risk of algorithmic collusion is influenced by actions that are permitted by market participants and whether certain behaviour is restricted for all participants or solely pricing algorithms.

There are currently some implicit restrictions on market behaviour that reduce the risk of algorithmic collusion. For example, the market price cap (MPC) can dissuade collusion in the case where the collusive equilibrium price is above the cap. The MPC therefore reduces the returns from collusion, which makes collusive strategies harder to sustain.

Providing the action of generators and other resources continue to operate within the 5-minute dispatch process and targets are assigned through NEMDE's security-constrained dispatch process, there are positive guardrails to protect power system security. For example, while the action of CER is controlled through participation in NEMDE, the market is aware of the potential power system implications and can adjust their dispatch instructions accordingly.

More broadly, measures could be taken to limit bidding and re-bidding behaviour, or to limit the amount of auto-bidding that is permissible. However, the pro-competitive benefits of pricing algorithms would need to be accounted for before simply curbing the use of AI to prevent anti-competitive conduct.

Auto-bidding technology is likely to be more important in promoting emerging technologies, such as those that enable price response demand (and supply) through consumer energy resources. For example, small-scale batteries may benefit from relatively frequent bidding and rebidding activity when making decisions about charging or discharging in response to changing market conditions. Auto-bidders may reduce the costs of such activities.

3.4 Information provided to the market may be a double-edged sword

The NEM operates efficiently when each generator bids at the minimum price that they would be willing to supply that unit of electricity – their reservation price – even though there will be periods where the price may be many multiples above this bid.

The primary benefit of information provision should be to encourage consumer response (that is, switching providers or changing consumption decisions), typically by reducing search costs for consumers. However, adding more information around bids/market offers can potentially carry risks because:

- The NEM is a one-sided market where demand is for the most part extremely limited. It therefore follows that publishing information about the NEM is not particularly useful for consumer decision-making.[48]

47 ACCC, *Digital platform services inquiry, Interim report No. 4 - General online retail marketplaces*, March 2022, p. 59. Available here.

48 Note that at the margin, there remain incentives for large consumers to reduce their electricity consumption under certain circumstances.

- Byrne and DeRoos (2022) highlight a growing empirical literature that shows information sharing – often provided through government-run price information platforms – can facilitate anticompetitive conduct and higher prices for consumers.[49]

To reduce the ability of competitors to behave strategically, purely illustrative examples of policy responses could include delaying the publication of historical bidding and price information or perhaps less controversially, only publishing the constraints that are currently violating or binding in the NEM.

Increasing the scrutiny of rebidding could also be considered, for instance through questions about when rebidding is allowed, what information is needed to justify a rebid, and monitoring of rebidding itself. Policies should consider the minimum information needed to promote more efficient bidding behaviour.

Current and future changes to information provision requirements should avoid inadvertently publishing information that could be used to predict a generator's current or future pricing strategies, or to reveal a competitor's ability to store or generate electricity.

3.5 It may be difficult to fully capture algorithmic collusion through competition law

Competition law is important in addressing the incentives for collusion, as it specifies what types of behaviour are illegal and the expected punishment for participants engaging in collusion. However, previous research highlights that it is unclear whether the CCA applies to algorithmic collusion. Chan (2021)[50] argues that the CCA does not prohibit algorithmic collusion in the NEM, highlighting:

- There is ambiguity about whether algorithmic collusion is encompassed by the cartel, concerted practices, or misuse of market power prohibitions in the CCA
- There is no clear legal precedent given algorithmic pricing has not been tested under Australian competition law, nor have the recent changes to the concerted practices prohibition under the CCA.

The author proposed two amendments to the CCA to address the risk of algorithmic collusion in the NEM:

- **Expand the information-gathering power of the ACCC:** accounting for machine learning algorithms by requiring generators who wish to use AI to notify the ACCC about changes in the algorithms they use.
- **Expand the concerted practice prohibitions in the CCA:** the evidence to prove a concerted practice would be based on "ex-post economic analysis, awareness, and no reasonable steps to prevent the conduct". [51]This addresses the concern that it would be difficult to apply the current prohibitions to tacit or autonomous algorithmic collusion in the NEM.

However, interventions aimed at reducing the risk of anti-competitive behaviour by algorithms must also balance the risk of stifling the pro-competitive benefits of such technology. For example, imposing reporting obligations would increase the cost of using such technology. This could also potentially reduce competition by limiting the uptake of new technologies. These interventions may also be difficult to practically implement, because:

49 Byrne and de Roos. *Startup Search Costs,* SSNR, September 2020. Available here.

50 Chan. *Algorithmic collusion and Australian competition law: Trouble ahead for the national electricity market?* UNSW Law Journal, 44(4), 1365–1408, 2021. Available here.

51 Ibid. p. 1405.

- It may be difficult to prescribe what constitutes a change in the algorithm that should be communicated to the ACCC.
- It may be difficult for a participant to understand what in the algorithm's code has changed, and to subsequently evaluate whether this change may have an anti-competitive effect. These defining challenges create the risk of under- and over-reporting of information, which increases the costs of the reform.

3.6 Reducing barriers to entry and participation in the market can reduce risks

The most effective policies to reduce the risk of algorithmic collusion are those that promote competitive and dynamic markets. To that end, several current and future rule changes by the AEMC indirectly address the risk of algorithmic collusion.

Measures to reduce barriers to entry, and address the risks of high market concentration, reduce the ability of competitors to coordinate and reduce the profits that can be earned if they were to collude.

Policymakers should ensure that the costs of entry and exit are as low as possible for new electricity generation (e.g. by ensuring planning system constraints are the minimum needed to meet environmental and community objectives).

The recently finalised rule on expanding the transmission ringfencing framework[52], (which promotes effective competition in the market for contestable connections) will (at the margin) support new generation by ensuring the costs of connecting the network are as low and streamlined as possible.

The stochastic nature of electricity supply and demand inevitably creates system constraints. Bottlenecks and constraints increase the risks of algorithmic collusion by increasing regional market power for short periods of time. Within the electricity transmission space, the AEMC is currently considering how to address these constraints by considering options for flexibility in the allocation of interconnector costs[53] such that investment in transmission projects between different parts of the NEM proceeds where they reduce electricity prices. Increased adoption of battery storage across the network could also address these risks, as they may reduce the price impact of network constraints in periods of scarcity.

Currently, the NEM is a one-sided market. However, encouraging responsive consumer demand is important to address the risk of algorithmic collusion. With responsive demand, an increase in price leads to a reduction in demand, reducing the profit that can be earned if prices are above competitive levels. As a promising step in this direction, we note that the AEMC is currently working on reforms that will help energy consumers respond to price signals and integrate consumer energy resources (CER, e.g., batteries, pool pumps) in the wholesale electricity markets. Notable examples are 'Unlocking CER benefits through flexible trading'[54]and 'Accelerating smart meter deployment'[55]currently being considered by the Commission as part of our rule change process.

52 AEMC. Expanding the transmission ring-fencing framework. Available here.

53 AEMC. Providing flexibility in the allocation of interconnector costs. Available here.

54 AEMC. Unlocking CER benefits through flexible trading rule change. Available here.

55 AEMC. Accelerating smart meter deployment. Available here.

3.7 We recommend establishing a working group with the ACCC and AER

This chapter has provided a high-level overview of the risks of algorithmic collusion in the NEM. It suggests the risks are increasing, but also, that there are difficult trade-offs to navigate to identify an appropriate regulatory response.

To progress this work, we recommend the Commission to establish a working group with the ACCC, as the competition regulator, and AER, which is responsible for compliance within in the energy sector. Within this group, the three organisations would work together to consider or design appropriate regulatory response.