Electricity hedge market-making cost benefit assessment

12 September 2025

concept
consulting

About Concept Consulting Group

We have been providing useful, high-quality advice and analysis for more than 25 years. Our roots are in the electricity sector and our practice has grown from there. We have developed deep expertise across the wider energy sector, and in environmental and resource economics. We have also translated our skills to assignments in telecommunications and water infrastructure.
Our directors have all held senior executive roles in the energy sector, and our team has a breadth of policy, regulatory, economic analysis, strategy, modelling, forecasting and reporting expertise. Our clients include large users, suppliers, regulators and governments. Our practical experience and range of skills means we can tackle difficult problems and provide advice you can use.
For more information, please visit www.concept.co.nz.

Disclaimer

Except as expressly provided for in our engagement terms, Concept and its staff and contractors shall not, and do not, accept any liability for errors or omissions in this report or for any consequences of reliance on its content, conclusions or any material, correspondence of any form or discussions, arising out of or associated with its preparation.

The analysis and opinions set out in this report reflect Concept's best professional judgement at the time of writing. Concept shall not be liable for, and expressly excludes in advance any liability to update the analysis or information contained in this report after the date of the report, whether or not it has an effect on the findings and conclusions contained in the report.

This report remains subject to any other qualifications or limitations set out in the engagement terms.

No part of this report may be published without prior written approval of Concept.

Contents

Executive summary

Market making in New Zealand's electricity futures market ensures liquidity, transparent forward prices, and access to standardised hedging products. While baseload contracts are well established, shaped products such as peak and super-peak remain underdeveloped, limiting participants' ability to manage risk efficiently and for forward prices to differentiate key drivers of that risk. Transparent forward prices are increasingly critical for operational decisions, investment planning, and supporting retail competition.

This report evaluates the economic costs and benefits of alternative market-making arrangements across four scenarios: baseload only; baseload plus peak; baseload plus super-peak; and all three products. Impacts on hedgers, investors, and market makers are considered.

Key findings

- Price signals: Super-peak products provide the clearest and most consistent signals, including during high-risk periods. Peak products are weaker but could be more informative if extended to weekends. Baseload alone offers generic signals that provide less insight into system stress.
- Hedging benefits: Baseload contracts deliver substantial risk reduction for retailers and solar generators. Adding either a peak or super-peak product further reduces earnings-at-risk for retailers, but a third product adds minimal incremental benefit. Access, spread, and liquidity will affect practical hedging outcomes.
- Market-making risk: Earnings-at-risk is sensitive to buyer-seller imbalances more so than product characteristics. Extreme risks can be mitigated via spread and volume settings. Super-peak is not inherently riskier than baseload, though short-term seller scarcity may increase risk exposure.

Economic benefits

- Retail competition: Access to baseload and shaped products could reduce a notional independent retailer's financial exposure by ~8%. Extending this across all independent retailers could provide retail competition benefits estimated at $2.2m per year, increasing to $3m per year by 2032.
- Generation investment: Shaped products improve timing and targeting of investment. Avoiding a modest share of shortage, fuel, and carbon costs could conceptually save $6.25m per year by 2032; and a 1% reduction in earnings-at-risk for future solar and battery investments could reduce financing costs by $1.2m per year by 2032.
- Broader effects: Benefits also include improved fuel management, demand-side investment, and operational efficiency, growing as flexible demand participation expands.

Recommendations

A combination of baseload and super-peak products is likely to achieve the greatest benefits. Calibrated spreads and volumes, and eventual transition to exchange trading for super-peak is recommended to:

- Lower retailer earnings-at-risk and support competition.
- Improve price discovery and investment efficiency.
- Manage market-maker risk without compromising liquidity.

1 Introduction

1.1 Context

Market making in New Zealand's electricity futures market has been in place since the ASX futures market was first established in 2010. Since then, arrangements have evolved to support forward price discovery and transparency, and providing participants with access to standardised hedge products. Currently, four large generator-retailers act as market makers under regulatory obligation, alongside a commercial market maker contracted by the Authority.

Trading activity has concentrated on standardised baseload contracts, while other shaped products—such as peak or super-peak contracts—have seen less development. At the same time, the over-the-counter (OTC) market continues to provide bespoke but less transparent hedging options. As a result, the benefits of price transparency and liquidity have not always extended across the full range of risk management needs in the sector.

Transparent and reliable forward prices are increasingly important for guiding operational and investment decisions, including new generation investment, the efficient operation of existing assets, and pricing of power purchase agreements (PPAs). Hedge diversity also influences retail competition, as independent retailers and new entrants rely on them to manage risk and offer competitive pricing to consumers.

This report supports the Electricity Authority's current review of whether market-making arrangements remain fit for purpose.

1.2 Objective and scope

The purpose of this report is to assess the economic costs and benefits of alternative market-making arrangements, with a focus on how different product suites and contract specifications could better meet the needs of the evolving electricity system. The analysis also considers implications for the parties that provide market-making services, and how market-making settings could best enable the benefits to be realised.

The assessment is structured around four "anchor" scenarios:

- **Scenario A: Baseload only** – considers whether current baseload arrangements are sufficient to provide robust price discovery.
- **Scenario B: Baseload + peak** – considers the benefits and costs of adding peak products, and appropriate profiles, volumes, and bid-ask spreads.
- **Scenario C: Baseload + super-peak** – considers the potential value of a super-peak product, and appropriate profiles, volumes and bid-ask spreads.
- **Scenario D: Baseload + peak + super-peak** – considers the combined impact of a full suite of products.

The aim is to provide an evidence base on the value of revised market-making arrangements in the context of a market that is continuing to evolve.

Caveat: The analysis and modelling presented are intended to provide guidance on potential market-making arrangements rather than to prescribe exact bid-ask spreads or volumes. In practice, optimal settings will depend on evolving market-wide demand, liquidity, and participant behaviour. We have not sought to determine the "right" settings in detail. Accordingly, the figures and scenarios presented should be interpreted as indicative, supporting informed decision-making rather than representing definitive prescriptions.

2 Market making costs and benefits

2.1 All participants benefit from market making

A market maker is a participant that stands ready to both buy and sell a product at publicly quoted prices. By doing so, they ensure other parties have someone to trade with when they want to transact, helping to smooth mismatches between supply and demand. This role is important because it prevents buyers and sellers from being left waiting for a counterparty and helps the market function smoothly even when trading interest is uneven.

Market makers contribute to price discovery and transparency by regularly quoting prices and adjusting them in response to new information. Each trade helps update the market price, ensuring forward prices reflect the latest conditions. Transparent price expectations allow all participants to make informed operational and investment decisions.

2.2 Market makers can benefit but with risk

From the market maker's perspective, the economics of the role depend on the bid-ask spread—the difference between the price at which they will buy and sell.

In highly liquid markets with stable prices, narrow spreads can generate low-risk profits for the market maker because trades are frequent, and any positions they take on are quickly cancelled out by an offsetting trade.

In less liquid or more volatile markets, spreads are typically wider so that the market maker can remain profitable despite lower trading volumes and greater exposure to price movements while holding inventory.[1] However, while wider spreads may help a market maker earn more predictable returns, they also make it harder for buyers and sellers to agree on trades.

2.3 Our market makers can manage some risk

Market makers take on risk whenever they hold inventory—buying from or selling to one party before selling to or buying from another—because the price can change over that time. In principle, a market maker does not require physical generation assets to manage this exposure. If there is a balance of buyers and sellers, the role is essentially to smooth the timing of trades rather than to carry long-term risk. In such cases, the market maker will generally just absorb gains and losses through their balance sheet.

However, because some of New Zealand's market makers are also large generators, their physical assets are directly relevant. To a degree, it is appropriate that risk falls on these parties, as they are best placed to manage generation risk in the wholesale market. If the market makers did not own generation, that risk would instead get passed through to other market participants through subsequent trades.

In stressed market conditions, risk should in principle be allocated to the supplier of last resort, and this can mean dynamic trading in the heat of the moment. However, the supplier of last resort will not always be a market maker, and the party best placed to manage the exposure may choose not to hedge it. There are hence limits to the extent to which risks can be efficiently redistributed through the hedge market, and the design of market-making arrangements becomes a significant factor in where risk lands.

[1] Ordinarily, market makers manage their exposure by adjusting quoted spreads and the volumes they are willing to post. In New Zealand, however, these parameters are regulated, reducing the flexibility market makers would otherwise have to manage their risk.

2.4 Financial traders provide benefits but can impact risk

Financial traders trade without physical positions to hedge, aiming to profit from price movements. They can add liquidity, support price discovery and challenge inefficient prices, and provide supply or demand where it is otherwise lacking. Some may also support small participants in accessing hedges by acting as an intermediary.

But traders operating without a hedging motive can also amplify short-term volatility and concentrate volumes in a few contracts, exposing market makers to bigger positions. Normally, market makers would manage this by adjusting spreads or volumes, but the regulated market-making settings allow only limited flexibility.

This highlights the importance of designing market-making settings that remain effective under different market conditions. At times, obligations may at times create exposures that are difficult for market makers to back with physical assets, so settings need to ensure risks are manageable while still supporting liquidity.

2.5 Determining the balance

In considering market making settings, a balance needs to be struck between:

- the ease and cost of trading for participants that rely on market making for hedging, and the wider benefits of transparent forward price signals
- the commercial sustainability of market making, recognising that potential benefits may be offset by the risk of inefficient, loss-making positions.

Therefore, in considering the four scenarios of interest to the Authority, our analysis aims to identify the balance between these factors that can best support competition, investment efficiency, and operational efficiency in the electricity market.

3 Current market making settings

There are currently three standardised electricity risk management products in the New Zealand market:

- **Baseload:** Currently trades on the Australian Securities Exchange (ASX). It covers 0.1 MW across all half-hour periods across every day, providing a continuous hedge against average system prices. The contracts cover full calendar months which are available for trading 6 months ahead, and calendar quarters which are available for 14-17 quarters ahead.
- **Peak:** Also ASX-traded, this product covers daytime hours between 7am-10pm on business days only. Contracts are 0.1 MW and available for calendar quarters 14-17 quarters ahead.
- **Super-peak:** A recently introduced product traded via fortnightly OTC auctions. It covers the hours of 7:00-9:30am and 5:00-9:00pm every day, including weekends. It was designed in 2024 by an advisory group convened by the Authority.

Settlement for all three products occurs against electricity spot prices at two reference nodes: Ōtāhuhu and Benmore. Options on the baseload product are also available.

Market making obligations currently exist only for the ASX baseload product. Four large generator-retailers are obliged to provide market making services, supplemented by one party contracted by the Authority.

Market makers must quote for all available quarterly and monthly baseload contracts, with:

- 3% bid-ask spread
- minimum quote size of 2.4 MW per market maker (giving 12 MW total), with the option to provide 50% upfront and 50% on refresh
- for a set half hour of each business day
- stress provisions allowing withdrawal from obligations for up to 5 of every 20 trading days, on a rolling basis.

No formal market making requirements currently apply to peak or super-peak products—quoting and liquidity are provided on a voluntary basis.

4 Analytical approach

4.1 We used our electricity price forecasts

For this analysis, we have leveraged Concept's price forecasts. These consist of modelled hourly prices derived from the interaction of supply and demand under a range of future system states, across multiple simulated 'weather years'.

Figure 1 shows an example of our forecast for a single year. The dark line represents the average price outcome across 43 weather years (indicated in blue) for one system state scenario.

Figure 1: Illustration of price forecast[2]

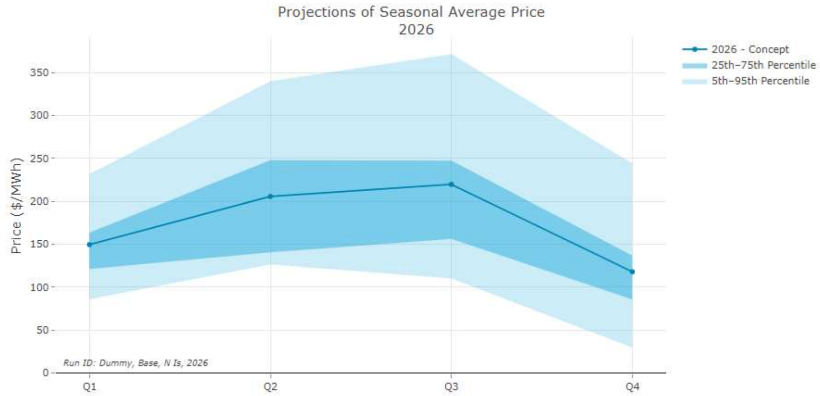

We use the average price from our central scenario as a proxy forward price curve, forming the basis of a notional hedged position. Because the forecasts are constructed at an hourly level, we can aggregate them into hypothetical forward price curves for baseload, peak, and super-peak products.

Against that hedge, we test a range of weather- and scenario-driven demand, supply, and price outcomes to assess potential risk exposures.

Using forecast prices allows us to capture the impacts of an evolving generation mix and load profile, rather than relying solely on historical price patterns. It also allows us to consider the range of outcomes that participants factor in when valuing hedges.

While the forecasts form the foundation of our analysis, they are supplemented with observed ASX futures outcomes where appropriate to ensure the results reflect current market conditions.

4.2 We analysed product profiles in future context

To understand which hours of the day provide the most meaningful price information beyond a simple baseload contract, we analysed the signal-to-noise ratio (SNR) of hourly prices using our forward price forecasts. The SNR is defined as the ratio of:

- the average difference between an hour's price and the baseload price ('signal')
- the variation in that difference across scenarios ('noise').

A high absolute SNR (>1 or <-1) indicates the hour is consistently different from baseload and reasonably predictable. A low absolute SNR (between -1 and 1) indicates the price in that hour either resembles baseload or is highly variable, and therefore provides less additional or reliable price information.

The SNR can be more informative than looking at the average difference or variability alone. For example, a large average difference suggests material exposure if unhedged, but if the variation across years is also high, that exposure is less

[2] Demonstration data only - does not represent actual forecast results.

predictable. By combining the two, the SNR highlights the hours where distinct and reliable price signals align.

This framework helps distinguish between hours that sharpen the informational content of a hedge product—by varying in a similar and predictable way—and those that dilute or complicate the signal by varying in an opposing or unpredictable way.

4.3 We developed earnings at risk models

We used models to analyse the earnings at risk from the perspective of:

- parties using contracts for hedging
- market makers.

The models are explained at a high-level here. The full sets of modelling assumptions and caveats are provided in 35Appendix B and Appendix C.

4.3.1 Modelling earnings at risk from hedging

We use an earnings-at-risk (EaR) model to quantify how different hedging strategies affect the earnings of electricity market participants, such as retailers or solar generators.[3] The model applies a notional hedge to a participant's demand or generation profile, evaluating how baseload, peak, and super-peak contracts influence risk exposure.

For each participant type, the model calculates hourly earnings across multiple simulated weather years, producing a distribution of annual outcomes. This allows assessment of average earnings as well as downside and upside scenarios, with the EaR metric defined as the difference between mean earnings and the 5th percentile outcome.

By testing different quantities and combinations of hedge contracts, the model identifies the hedge structure that minimises earnings at risk. This provides a clear picture of how hedging can mitigate exposure and supports quantitative evaluation of potential benefits for market participants and the design of market-making arrangements.

4.3.2 Modelling earnings at risk from market making

For this report, we developed a model to estimate the potential variation in market maker earnings across a wide range of future price and trading outcomes. Earnings primarily depend on inventory volumes, holding periods, and the magnitude of price movements. To capture this, the model simulates price paths using a 'random walk' across trading days, and randomly allocates trades in time.[4] It runs multiple simulations of both prices and trading behaviour to generate an indicative range of outcomes.

The model incorporates transaction revenue from bid-ask spreads, trading fees, and the cost of financing prudential requirements based on an ASX-style framework. It has been calibrated using historical ASX trading data to ensure outputs are grounded in observed market dynamics. However, there is inherent uncertainty in this type of modelling, and historical relationships may not always hold or apply in different contexts.

A key simplifying assumption is that the five New Zealand market makers are treated as a single entity, collectively taking on the inventory required for the rest of the market to trade. The model assumes the market makers do not take any active steps to

[3] Earnings-at-risk is a measure of potential downside in future earnings across many scenarios, typically expressed as the expected shortfall relative to average earnings. When quantifying the earnings-at-risk, we refer to the average versus 5% probability lowest outcome.

[4] A "random walk" analysis assumes that future price changes are unpredictable and follow a path where each movement is independent of the last, like successive coin tosses.

manage or offload this inventory—including through prices. It makes no assumption about how the inventory (and hence the cost of holding it) is allocated among them, and in practice the distribution may be uneven.

We emphasise that the model and forecast data have limitations. It is not intended to capture the full complexity of trading dynamics, but rather is intended to be illustrative. All figures are estimates.

4.4 We assessed product set economics

The outputs of our price forecasts and earnings at risk models allow us to consider the economic costs and benefits of different market-making arrangements in a structured way. Our analysis has focused primarily on two key benefit areas: supporting retail competition and encouraging efficient generation investment. We have also considered operational efficiency and demand-side investment, though these have been addressed in a lighter-touch manner.

By linking simulated market outcomes to these benefit areas, we can evaluate how alternative arrangements may improve hedging opportunities, strengthen liquidity and transparency, and support overall market efficiency. This ensures our recommendations are grounded in a consistent, quantitative framework while remaining mindful of the uncertainties inherent in future market conditions.

5 Results of our analysis

This section discusses the key results and insights from our analysis.

We note the analysis and modelling presented are intended to provide guidance on potential market-making arrangements rather than to prescribe exact bid-ask spreads or volumes. In practice, optimal settings will depend on evolving market-wide demand, liquidity, and participant behaviour. We have not sought to determine the "right" settings in detail. Accordingly, the figures and scenarios presented should be interpreted as indicative, supporting informed decision-making rather than representing definitive prescriptions.

5.1 Price signals under different product profiles

We used a signal-to-noise ratio (SNR) to highlight the hours where shaped hedge products provide the most useful extra price information compared with a baseload hedge. The analysis treats each hour as if it could be hedged on its own, and then compares the price it would attract to the baseload contract it sits within.

The point of this exercise is to identify groups of hours that consistently move together and therefore tell a strong, reliable story about the profile of prices. These groups can form the basis of shaped hedge products that add meaningful price discovery beyond a simple baseload contract.

An SNR tells us how strong the 'signal' is compared with the 'noise':

- The signal is the average size of the difference between that hour's price and the baseload price.
- The noise is how much that difference varies under different system conditions. For example, if an hour is sometimes materially above baseload and sometimes materially below, the signal is "noisy" and unreliable.

A high SNR suggests the hour shows a clear, consistent difference from baseload. A low SNR suggests the difference is small or too inconsistent to be informative about the drivers of prices. The ±1 boundary is significant because it marks the point where the 'signal' is stronger than the 'noise', meaning the difference is more likely to be meaningful, rather than random fluctuation.[5]

Results for 2026–2028 are shown in Figure 2. The analysis and results are discussed in greater detail in Appendix A.

Figure 2: Signal to noise ratio of each hour relative to baseload

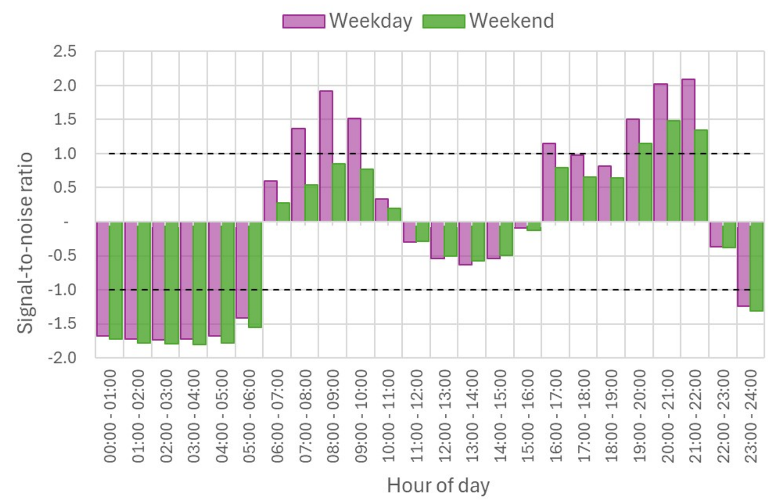

The key insights from this analysis are:

- **Overnight (10pm–6am):** These prices are materially and consistently lower than baseload. Both the peak and super-peak

[5] The ±1 threshold is not an empirical cut-off but a practical guide.

products exclude these hours so provide more distinct signals, but low prices and low variation in these hours means low risk.

- **Morning (7—10am) and evening (4—10pm):** These hours can provide strong price signals. They broadly match the super-peak product though the evening peak is slightly earlier and longer.[6] Including these in the product could provide consistent information that adds to understanding about the profile of prices, but the signal is expected to diminish in the longer-term, and the risk-management value of including them is low.
- **Mid-evening hours 5—7pm:** These hours show a lower SNR, but this reflects very high variability that increases over time. This means the risk is most significant in these hours, so its inclusion in a peak product is imperative for risk management purposes, but the price cannot be as reliably signalled, which will likely drive product risk premiums.
- **Midday (10am—4pm):** Prices in these hours are very similar to baseload prices on average, such that the weak 'signal' is overwhelmed by any 'noise'. This means these prices are well signalled by a baseload product and can be effectively managed with one. Including these hours in a peak product only dilutes the clarity of the signal around hours that drive higher prices. Further, these prices are forecast to 'hollow out' progressively over time with increased solar penetration, so a peak product would price in increasingly divergent signals.
- **Weekends:** Signals are consistently weaker, because weekend prices vary less predictably from the baseload price. For the super-peak product, the inclusion of weekends complicates the product's valuation, but the variability is greatest in winter quarters and grows over time, so the risk management value of including them is likely the prevailing factor. For the peak product, including weekends may actually make the product *easier* to value, as the excluded periods would be limited to those with a clear, consistent signal. The product would then also capture the risk-management value of hedging weekends.

Looking ahead (2032–2037), our forecasts indicate prices will skew further towards the evening peak and become overall much harder to predict. This could, in time, justify separate morning and evening peak products (as has recently been introduced in the ASX products for the Australian NEM[7]), and will likely lead to higher risk premiums.

Overall, the super-peak product aligns best with hours that deliver clear and consistent price signals while also capturing important risk-management value, including on weekends.

The peak product covers relevant hours but provides a weaker, more complex signal. Extending it to weekends could sharpen its price-signal and improve its practical risk-management value.

5.2 Impact of product sets on hedgers

We undertook earnings-at-risk analysis to understand the value of different products to participants. We considered an independent retailer and an independent (merchant) solar generator. We iteratively tested different combinations of hedge products to identify the minimum earnings-at-risk achieved given a wide range of future scenarios.

The results for an independent retailer for 2026-2028 are shown in Figure 3, and those for an independent solar generator are shown in Figure 4. These are discussed in greater depth in Appendix B.

[6] The morning peak for the super-peak product cover 7-9:30am. Our price forecasts are hourly, and hence cannot reflect on the half-hour difference in window.

[7] https://www.asxenergy.com.au/newsroom/industry_news/launching-morning-and-evening

Figure 3: Risk mitigation for independent retailer

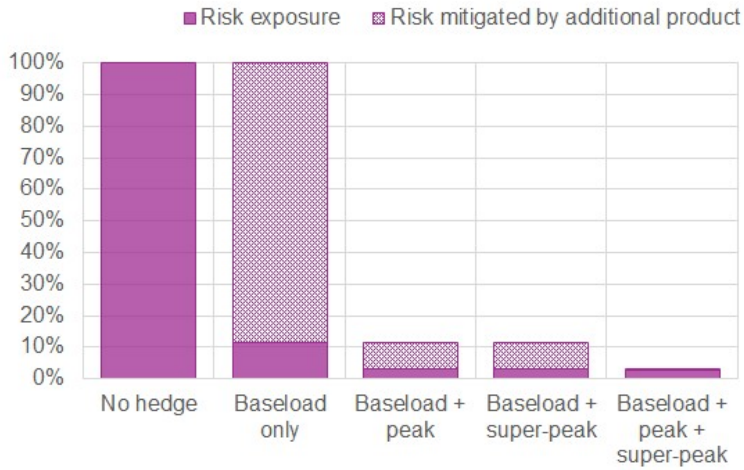

Figure 4: Risk mitigation for independent solar generator

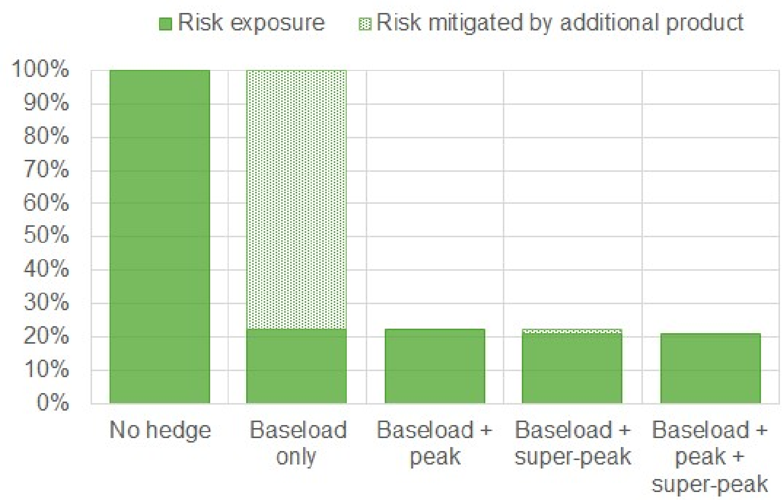

Key insights from this analysis are:

- **Baseload delivers substantial benefit:** The baseload product alone materially reduces risk, potentially lowering earnings-at-risk by around 90% for an independent retailer and 80% for a solar generator (averaged across system states).
- **Peak and super-peak provide similar value for a retailer:** For independent retailers, both products reduce earnings-at-risk to just a few percent (on average). Inspecting the results in more detail indicates that earnings-at-risk is marginally reduced when these products include the weekends. It also suggests that volumes required to achieve the minimum earnings-at-risk are different:
- **Adding a third product delivers little extra benefit to a retailer:** Introducing a third product reduces earnings-at-risk by less than 0.5% compared with two products.
- **Limited benefit to solar from selling peak and super-peak:** Neither peak nor super-peak materially reduce risk for solar generators. For super-peak this is expected, since it covers low-generation hours. For peak, the result is explained by the fact solar output is concentrated in midday hours, where prices already align closely with baseload, while overnight hours carry little inherent risk (as identified in the SNR analysis).
- **Residual solar risk remains, but may be managed through tailored structures:** Around 20% of earnings-at-risk remains for solar after selling hedges using any or all of the three products. We note solar generators could potentially achieve further reductions by combining baseload sales with targeted super-peak *purchases* to better align their hedge books with their generation profile. We have not considered this in detail in our analysis.

These results relate to the years 2026-2028. Analysis of later years indicates that earnings-at-risk (absent hedging) will roughly double over time. For an independent retailer, hedging with baseload alone becomes less effective against the changing risk, but hedging with a combination of the baseload and peak or super-peak products remains effective (with the super-peak becoming slightly more effective than the peak product). However, selling these products becomes even *less* effective for a solar generator over time, and their residual risk with baseload products roughly doubles by 2037. This may justify more tailored hedge products in future.

For each of the four scenarios, the lowest-risk point is achieved with a combination of hedges that is similar on a total GWh basis, with the new products substituting baseload volume.

Figure 5 shows how this translates into a MW basis. We have excluded the scenario involving all three products as the optimal mix of products can change significantly with system conditions.

Figure 5: Optimal hedge mix for independent retailers for three scenarios

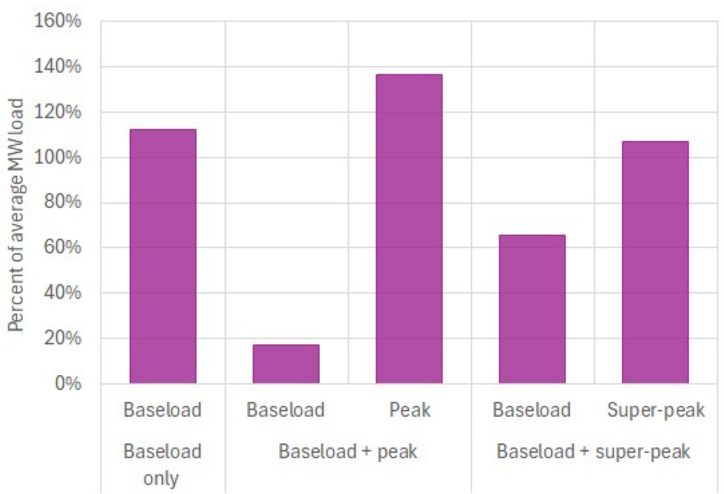

This graph highlights that:

- For the baseload product, optimal risk reduction is achieved by over-hedging by around 12%.
- For baseload and peak, optimal risk reduction is achieved through relatively small baseload volumes and large peak volumes (around a 10–90% split).
- For the baseload and super-peak scenario, optimal risk reduction is more balanced (around 40–60% split).

In interpreting all of the results above, it is important to note that the analysis has effectively assumed equal access and cost for all hedge products. In practice, outcomes will vary depending on spread, liquidity, and frequency of access:

- bid–ask spread: The spread is the direct cost of hedging. Wider spreads reduce the net benefit of risk reduction, though the materiality depends on whether participants can pass costs through to retail tariffs or generation prices while remaining competitive. If peak or super-peak products are more expensive than baseload, the optimal hedge mix would shift toward baseload.
- liquidity (depth): Even if quoted spreads look narrow, limited market depth may mean participants cannot transact required volumes without moving the price. This reduces the practical benefit of additional hedge products relative to the analysis.
- frequency of access: Infrequent or narrow trading windows increase timing risk by making it harder for participants to adjust hedge positions when market conditions change. This again reduces the ability to achieve the "ideal" mix shown in the modelling.
- overall effect: These frictions would not eliminate the benefit of risk management, but they would tend to reduce the incremental value of peak and super-peak products relative to baseload.

Overall, the analysis shows that while the baseload product provides the bulk of risk reduction, peak and super-peak products can deliver meaningful additional benefits for retailers, further lowering earnings-at-risk across a range of scenarios. However, their value will depend on the ease and cost of access. For solar generators, selling baseload remains the key tool, though there is potential for super-peak purchases to help mitigate further risk.

5.3 Impact of product sets and obligations on market makers

We used an earnings-at-risk analysis to assess how different market making arrangements affect potential earnings and risk for market makers across baseload, peak, and super-peak products.

The analysis uses a model that simulates trades and prices under multiple scenarios, drawing on our central price forecasts and plausible ranges of weather-driven outcomes (as discussed in Section 4.1). It does not capture the complexity of real trading, but it illustrates key dynamics of inventory accumulation, spreads, and trading bias. A full explanation of the modelling assumptions is given in Appendix C.

Figure 6 shows the modelled distribution of market maker earnings for three individual product scenarios, represented as box-and-whisker plots (P95, P75, P50, P25, P05). Combined, the purple and green boxes represent the interquartile range (P25–P75), the line inside the box shows the median (P50), and the patterned bars show the 5th and 95th percentiles. Outliers beyond this range are not shown.

The scenarios reflect combinations of bid-ask spreads and volumes available for trading:

- 3% spread with 12 MW of volume
- 5% spread with 12 MW of volume
- 3% spread with 6 MW of volume

The first two scenarios roughly align with current and past baseload settings, while the third is a notional reduction included for demonstration. These scenarios assume an equal number of buyers and sellers across contracts. Results are illustrative.

Figure 6: Distribution of market maker earnings for different products under different service levels

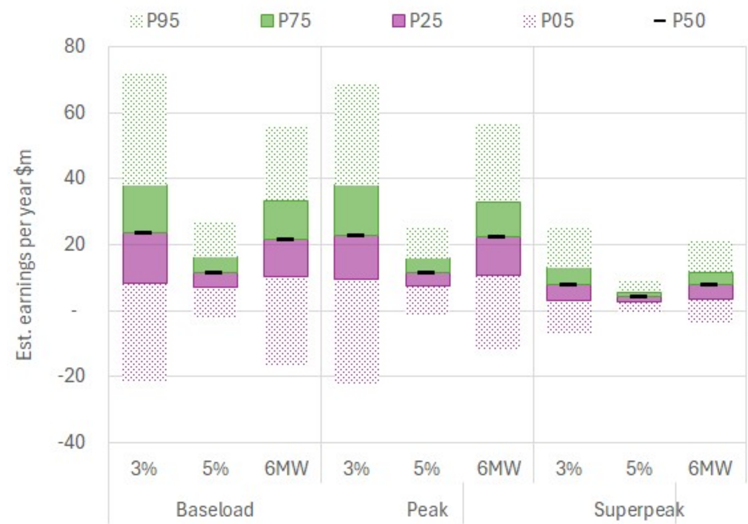

Key insights from this analysis are:

- **Positive earnings across most scenarios:** Under these assumptions, collective market maker earnings are positive, and within a relatively consistent range across the majority of outcomes.
- **Tail risk is significant:** Extreme scenarios can generate much higher or lower earnings, with earnings-at-risk (average minus 5th percentile) around $45m for the baseload contract under a 3% spread.

- **Spread and volume impact risk and returns:** Increasing the spread to 5% reduces earnings-at-risk by around 70%, though prospective earnings also fall. This reflects that the higher spread increases transaction costs for participants which can materially reduce trading activity. This in turn reduces market makers' risk by reducing their inventory accumulation, but at the expense of transaction earnings. Lower volume caps can similarly impact risk and returns by helping manage the volume or timing of trades and hence how inventory accumulates, though this is less sensitive. Overall, both tools can help to reduce market maker risk, but at the expense of participant access—highlighting the need to balance sustainable market making with reasonable liquidity for those seeking to hedge.
- **Peak and super-peak products do not appear inherently risky:** Expected earnings from the peak product are similar to baseload, while super-peak earnings are lower (noting the contract covers fewer hours so has a lower dollar value). Neither product appears inherently more risky than the baseload product under equivalent conditions.[8]

As stated, this analysis assumes a relatively even share of parties wanting to buy from and sell to the market maker. Historical trading data suggests this has been near-true for the baseload product. Market makers (collectively) buy around 95% of what they sell for each contract on average, though there is a bias towards them buying Otahuhu contracts and selling Benmore contracts.[9]

Figure 7 shows that the distribution of earnings becomes heavily skewed by a change in the balance of buyers and sellers, and the tail risks become extreme. When more participants want to buy from the market maker than sell to it, it faces a much greater risk of loss. When more participants want to sell to the market maker than buy, it has a greater potential for profit. This asymmetry arises because prices cannot fall below zero, so losses from concentrated buying can exceed gains from concentrated selling.

Over a large portfolio, gains and losses can offset, but if market makers are exposed to systemic biases or extreme outcomes in a subset of contracts that are unforeseen, then the risk could be significant.

Market makers will seek to manage this exposure through pricing: they can increase offer prices to discourage further buying when inventory risk grows, effectively embedding a price premium that would come at the detriment of other traders. This mechanism invites short-selling, where possible, which can help rebalance flows. These dynamics are important in reality but are not able to be captured in our analysis. The risk presented in our analysis is therefore likely an over-estimate, and could be considered risk that may in practice remain with other participants.

[8] Note the modelling assumes prices for these products during the last 6 months of their trading are 120% more volatile than for a baseload contract.

[9] This is shown later in Figure 8

Figure 7: Change in distribution of market maker earnings given different proportion of buyers and sellers[10]

Overall, the analysis shows that market maker risk is more sensitive to trading imbalances than the intrinsic characteristics of the products. Where these exposures can be backed off through physical assets, market makers may be well placed to manage them. Where they cannot, the risks are harder to absorb and may become inefficient.

The impact on market maker risk from pricing different products will hence depend on expected buying and selling interest, including from financial traders. At a market-wide level, this is a key uncertainty.

5.4 Impact on economic costs and benefits

5.4.1 Retail competition benefits

Improved access to baseload, peak, and super-peak hedges reduces retailers' earnings-at-risk, giving them greater confidence to manage price volatility. Lower earnings-at-risk enables retailers to offer more competitive and differentiated products, encouraging new entrants and strengthening competition. Stronger competition drives efficiency by incentivising retailers to innovate, lower costs, and better respond to consumer preferences. For consumers, this can translate into more choice, products that better match usage patterns, and potentially lower overall electricity costs.

Our analysis of the earnings-at-risk reductions achievable for retailers with baseload, peak, and super-peak products shows that better access to shaped products could reduce their financial exposure by around 8%, relative to relying on the baseload product alone. This benefit is expected to increase over time as prices become 'peakier'. The reduced risk could reduce those retailers' need to hold collateral in reserve to manage it, reducing their cost-to-serve.[11] If other retailers needed to match this efficiency gain to remain competitive, it could translate into a potential market-wide benefit of $2.2m per year, increasing to $3m per year in 2032. While the direct benefit is relatively modest, reduced barriers to participation could significantly amplify the benefits and see the impact of the change drive broader efficiencies that reduce prices more significantly over time.

5.4.2 Generation investment benefits

Investment in new generation ensures sufficient capacity to meet future demand, improves reliability, and reduces the likelihood of

[10] The x-axis here descrbies purchases and sales from the market makers perspective – eg, the market maker buys 120% of what it sells.

[11] Our modelling suggests costs could be reduced by around $40,000 for every 12,000 customers, increasing to $50,000 in 2032, assuming collateral is held to cover the full risk at a 6% WACC.

costly shortages. Efficiently targeted investment also supports lower operating costs, by encouraging generation that is dispatched when it is most needed, and can reduce fuel use and emissions. Together, these effects improve overall market efficiency and the long-term affordability and reliability of electricity supply.

As identified in our SNR analysis, a well-designed shaped product can capture differentiated price signals across the day. With shaped products—alongside baseload—investors can isolate the drivers of price risk, guiding investment to the periods and technologies where it is most needed. The baseload product alone provides a generic signal, but does not distinguish between specific sources of risk. Correctly signalling the timing of this risk is also critical, as generation that arrives too early or too late increases costs for consumers.

More robust information on future price shapes is expected to improve efficiency. Our price forecasts indicate a situation of insufficient generation in 2032 could incur around $65m of shortage costs and an additional $60m of fuel and carbon costs, compared with sufficient generation.[12] With improved information, such outcomes are less likely to occur, avoiding the associated economic costs.

The same scenario could see average wholesale prices $50/MWh higher, corresponding to a one-year wealth transfer of around $2.5b from generators to consumers. This transfer does not represent economic inefficiency and would not be captured in an assessment of costs and benefits. However, if such wealth transfers were passed through in retail prices, they could have social consequences that would be relevant to the assessment—though we have not sought to quantify these.

Our forecasts also anticipate the build-out of over $1b of new solar generation and $0.3b of batteries by 2032. Illustratively, a 1% reduction in earnings-at-risk for these investments could reduce financing costs by $1.2m per year by 2032.[13]

Conversely, higher earnings-at-risk for market makers could increase their financing costs in extreme scenarios—noting that for market makers with physical assets, earnings from market making represent only a small portion of total revenues, which are influenced by other significant sources of risk and may at times be intertwined with hedging activity.

5.4.3 Broader economic benefits

Shaped products may also provide additional, smaller-scale economic benefits beyond risk management and investment efficiency:

- **Fuel management:** While baseload contracts largely cover standard fuel procurement needs, shaped products could allow generators to better align fuel use with periods of highest price risk, potentially reducing costs or enabling more efficient dispatch.
- **Demand-side investment:** By providing clearer price signals at specific hours, shaped products could encourage investment in flexible demand resources, such as behind-the-meter storage or controllable loads, that target periods of peak risk.

[12] Assuming 20,000 Value of Lost Load, $15/GJ gas, and NZ$150/tonne of coal.

[13] Assuming a 1% reduction in EaR translates to a 10 basis point reduction on a WACC of 6%, over 20 years. The assumption that a 1% reduction in EaR corresponds to a 10 bps reduction in WACC is a conservative, illustrative assumption, but is consistent with standard financial modelling practice. The exact relationship is difficult to quantify as it depends on firm-specific risk profiles, capital structure, and market conditions.

- **Demand-side operations:** Operationally, demand response and flexible consumption can be optimized when prices across the day are better signalled. As demand-side participation grows, shaped products are expected to increase in value by allowing more precise response to the drivers of system risk.

Overall, while these benefits are likely smaller than the core retail competition and investment efficiency impacts, they are expected to grow in importance as flexible demand becomes a more significant part of the market.

6 Costs and benefits by scenarios

This section synthesises the results of our analysis and applies it to the four product set scenarios of interest to the Authority:

(a) Baseload only

(b) Baseload + peak

(c) Baseload + super-peak

(d) Baseload + peak + super-peak

6.1 Baseload only

6.1.1 Economic benefit impacts

The baseload market has been operating for some time and has developed progressively as participants have become more familiar with the products. The most significant lift in trading activity and forward liquidity followed the move to a 3% spread, which appears to have unlocked a step-change in hedging, while also driving significant 'churn' between market makers (ie, market makers trading back-and-forth between each-other). Against this backdrop, the scope for further gains from simply expanding baseload volumes now appears limited.

Further, baseload-only arrangements have important limitations. While they provide a useful generic price signal, they cannot distinguish between energy- and capacity-driven risks. As a result, retailers and other participants will continue to face challenges in understanding and managing different risks, reducing the efficiency of their operations and investment decisions. For these reasons, baseload alone is unlikely to represent the optimal market-making scenario.

Instead, we expect additional improvements in market outcomes are more likely to come from greater focus on other products, where price signals and risk management tools remain less developed.

Regarding the trading horizon, on the one hand, longer-dated prices could provide additional information and greater certainty for participants who want to lock in positions further out. Some increase in activity has been observed at the end of the curve, and if buying and selling were balanced, the additional costs to market makers would be modest given long-horizon prices that are relatively stable. We anticipate the potential for a small benefit from a two year extension, but that average earnings would start to reduce beyond that as financing costs outweigh any transaction income[14]–though this depends almost entirely on assumptions around demand, and does not consider resourcing costs for market makers.

Further, for market makers with physical assets, the incremental cost could arguably just be the financing and operational costs associated with daily exchange trading, as they may be the natural counterparty to these trades if they were to otherwise occur OTC. We anticipate the financing costs to be around $5,000 per MW held per year.

On the other hand, long-term forecasts are inherently uncertain for all parties, and most long-dated contracts use indexing to reflect this. Thin liquidity means prices are unlikely to be particularly meaningful, and if trading were one-sided, market makers could face prolonged inventory risk, with low transaction earnings unlikely to offset that risk. Retailers will generally not prefer to lock in long-term positions, as it could leave them unable to compete if prices move. Buying interest may therefore be limited to large industrial consumers (though they may be better served by other products such as PPAs). For most participants, the confidence of accessing

[14] And assuming negligible mark-to-market earnings/loss

fair-value products on a rolling three-year window is likely to provide as much certainty as longer-dated hedges themselves.

Overall, we anticipate the potential for both costs and benefits to be minor. Absent greater certainty on likely buyers and sellers, we consider it unclear whether the costs of extending the horizon would be outweighed by the benefits.

6.1.2 Optimising costs and benefits

The move to a 3% spread in 2020 has delivered clear liquidity benefits, with higher trading volumes, more forward hedging, and greater open interest suggesting evidence of an increase in genuine hedging (alongside an increase in trading 'churn' between market makers).[15]

Our analysis suggests market makers may have collectively faced a modest loss in recent years. This reflects small but significant imbalances in buying versus selling pressure (ie, the 95% identified previously), in a context of high and volatile prices due to scarcity in the underlying physical market. Losses appear to have been greater in contracts covering periods of physical supply stress.

However, the trading imbalance is across multiple contracts, varies by island, and inventory builds progressively over time. This is shown in Figure 8, which shows the net amount of MW traded with market makers – and hence their collective inventory – for each contract from 2021-2024. The shading indicates how far ahead in time the trades were made.

Figure 8: Net trading with market maker by contract and timing of trades

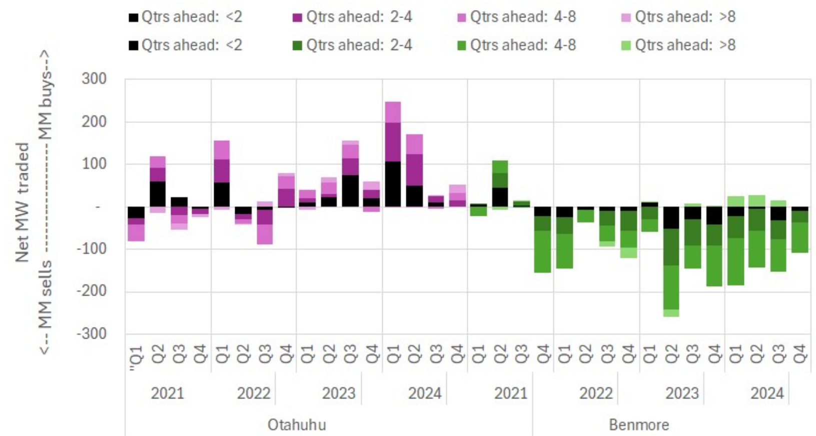

This gradual build-up of inventory suggests it reflects risk that the market makers with physical assets are best placed to manage. We therefore do not consider recent market making losses sufficient to justify a move away from the current 3% spread given its broader liquidity benefits.

That said, narrow spreads and high-volume requirements during periods of volatility can see market makers take on significant inventory, without access to some of the tools market makers would normally use in unregulated environments to manage those situations.

Reduced volumes could ease this pressure, but it would be important to consider the impact of a global change to volumes on

---

[15] As can be observed at https://www.emi.ea.govt.nz/Forward%20markets/Reports/DRERRQ?_rsdr=L10Y&DateFrom=20150101&DateTo=20241231&_si=v|3 and https://www.emi.ea.govt.nz/Forward%20markets/Reports/RYMLIE?DateFrom=20150101&DateTo=20241231&_rsdr=L10Y&_si=v|3

trading more broadly. Trading in excess of market maker volume appears to only occur on around 6% of trading days, primarily for near-term contracts, and on a rolling basis traded volume only represents 30% of total potential volume available from market makers. This suggests there may be some scope for a reduction in volume. However, we have not analysed the most appropriate volume in any detail.

6.2 Baseload and peak

6.2.1 Economic benefit impacts

Our SNR analysis indicates that the peak product does capture some incremental price information relative to baseload. However, this signal is less clear and less valuable than that of the super-peak product. Looking ahead, the usefulness of the peak product is expected to diminish as midday prices continue to hollow out and evening peak prices strengthen. This suggests the overall investment efficiency benefits from improved price discovery from market making in the peak product would be modest.

In terms of hedging, the peak product appears most valuable for demand-side participants, who could reduce overall earnings-at-risk by reducing baseload hedges in favour of peak.

The peak product has limited benefit for solar generators over and above the baseload product. While the peak product may have some value for existing hydro and thermal generators, we do not anticipate there are new-build technologies that would align well with a peak-only profile, which would limit potential investment efficiency benefits.

Consequently, the main economic advantage of introducing a peak product arises from retail competition. As discussed in section 5.4, by reducing earnings-at-risk for demand-side participants, peak contracts enable more confident risk management and support more competitive pricing, which we estimate as having a benefit of $2.2m per year, increasing to $3m by 2032. However, our analysis shows a super-peak product would achieve these same retail competition benefits and have greater investment efficiency benefits, and is therefore likely to be preferable to a peak product.

6.2.2 Balancing costs and benefits

A peak product is currently available for daily exchange trading, and we have assumed any market making arrangements would occur on the exchange.

Estimating an appropriate market-making volume for the peak product is inherently uncertain. It depends on assumptions about how widely the product would be used across the market, and on practical liquidity effects—neither of which we have analysed in detail. Our earnings-at-risk analysis of independent retailers suggests that they would have a strong preference for the peak product, substituting much of their current baseload positions while maintaining broadly the same overall GWh hedge volume. However, these retailers collectively make up only a small share of the overall hedge market and also rely on other hedging options such as bilateral contracts. Assuming they build positions gradually and do not need to make large adjustments at once, their needs could potentially be met with as little as 1-2 MW of peak volume.[16] Given this would largely substitute their baseload purchases, and other analysis suggests the existing 12 MW of baseload volume is rarely exhausted, such a low volume could likely be accommodated by a reduction in volume for the baseload product to 10 MW.

However, this 1-2 MW estimate does not capture liquidity effects, or account for other potential traders (including financial traders), which are uncertain.

We do not expect the peak product to attract much wider selling interest (beyond market makers themselves given their hydro and thermal generation), so reallocating too much baseload volume could reduce access for participants who prefer it. This suggests that peak volume should be at least partly additive rather than fully substitutive.

Balancing these considerations, we have used 10 MW baseload and 4 MW peak in our modelling. These are only indicative assumptions—we have not attempted to determine the "right" volume in detail, which would require more detailed demand and liquidity analysis. We note that the risks of setting volume too low are likely greater than the risks of setting it too high, as market makers can manage significant risk and have other tools available to manage their inventory.

The uncertainties around demand also affect the implications for market makers of extending support to the peak product, noting the potential for this to be imbalanced if the peak product attracts few natural sellers.

Given the uncertainty, we have modelled illustrative demand scenarios to indicate potential outcomes. Figure 9 shows the potential distribution of market maker earnings from both products under five illustrative cases:

- **The counterfactual** – the baseload product is subject to market making with a 3% spread and 12 MW volume.
- **Substitution case** – Independent retailers shift to their lowest-risk combination of baseload and peak, with no other demand changes meaning both products are imbalanced. Volume is provided at 10 MW for baseload and 4 MW for peak. Both are offered at a 3% spread.
- **Balanced case** – Independent retailers shift to their lowest-risk combination of baseload and peak, and other traders emerge to balance supply and demand for both products. Volumes are 10 MW and 4 MW, and spreads are 3%.
- **Unbalanced case** – Independent retailers shift to their lowest-risk combination of baseload and peak. Demand for the baseload product re-balances, but the peak product attracts half as many sellers as buyers. Volumes are 10 MW and 4 MW, and spreads are 3%.

[16] Depending on how much load is presumed to be covered by exchange-based hedging versus other hedging options, and the period over which a hedge position is built. The range we present would cover hedging of anywhere between 25-75% of total load, and a position built gradually each day for between 6-months to 1 year. Different assumptions around this, or factoring the potential need to make significant changes to a position in a short time could see this range double.

- **Suppressed liquidity case** – As above, but spreads for the peak product are increased to 5% so market makers can manage the potential imbalance.

The graph highlights again that earnings are generally positive. They are also sensitive to buying and selling imbalances, but earnings-at-risk is materially reduced in these circumstances by changing the spread. The substitution scenario shows slightly greater variation though this may be over-estimated given real-world price correlations between products that the modelling does not fully capture.

The sensitivity to a situation in which trading of the peak product is imbalanced suggests a prudent approach would be to maintain a conservative bid-ask spread until actual demand patterns emerge and trading behaviour becomes more embedded. Market makers will use prices to encourage a balance of supply and demand, but a higher spread gives them more options for managing inventory, particularly through periods of price volatility.

We do not consider the scale of earnings-at-risk for market makers indicated here would give rise to significant economic costs (only wealth transfers). We note market makers will be best placed to manage a significant portion of this risk, and will likely use price and spread to avoid taking on inventory in ways not captured in our analysis.

Figure 9: Outcomes for baseload and peak scenario under extreme scenarios of demand

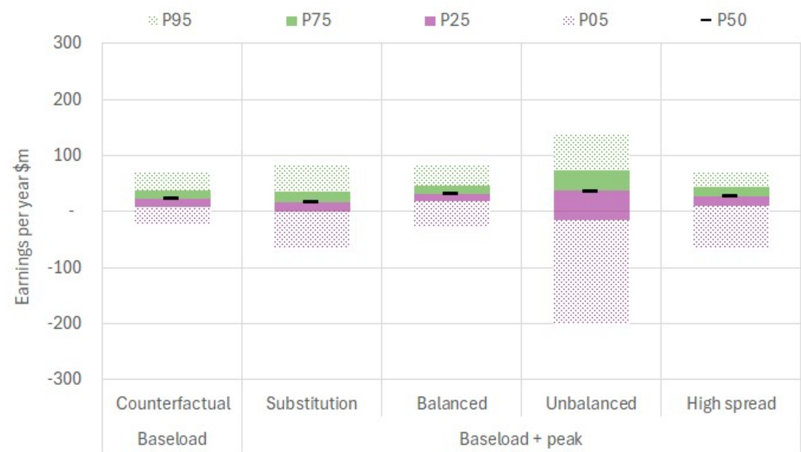

6.3 Baseload and super-peak

6.3.1 Economic benefit impacts

Adding a super-peak product alongside baseload enhances price discovery. It allows energy- and capacity-driven risks to be distinguished and ensures that future ‘duck curve’ impacts from solar remain isolated from the signal, which would not be the case for the peak product. This improved signalling can support more informed operational and investment decisions.

Applying a conservative 5% reduction to the expected shortage, fuel, and carbon costs identified for 2032 gives an illustrative benefit of $6.25m.[17] While this figure reflects avoided costs in a single year, similar benefits could accrue in subsequent years if demand continues to grow and scarcity-driven signals remain relevant, making these investment efficiency gains ongoing rather than one-off.

The super-peak product also delivers tangible hedging benefits. Our earnings-at-risk analysis shows that access to super-peak hedges can materially reduce earnings-at-risk for demand-side participants. As discussed in section 5.4, we estimate this could translate into an annual $2.2m benefit, increasing to $3m by 2032.

While selling super-peak contracts does not directly reduce risk for solar generators, they could potentially benefit from purchasing them. The product is also expected to support battery investors by informing and enabling hedging of charge/discharge strategies, though this has not been formally modelled. While the precise reduction in earnings-at-risk has not been quantified, conservatively assuming a 1% reduction would provide a potential benefit of $1.2m per year by 2032.

6.3.2 Balancing costs and benefits

The super-peak product is currently available for trading OTC, supporting ease of access for independent retailers who are likely to be early beneficiaries. This trading occurs fortnightly, and market making arrangements could be applied to this setting.

As with the peak product, our estimates of an appropriate market-making volume for the super-peak product are inherently uncertain and rely on assumptions about trader interest and behaviour and how the market may develop. Fortnightly OTC trading naturally focuses on physical hedgers—particularly independent retailers, who we expect to be the main source of demand in the near term.

Our earnings-at-risk analysis of independent retailers suggests that they would substitute some of their current baseload positions, but have a significant appetite for a super-peak product (while maintaining similar overall GWh hedge volume). On the same assumption that they build positions gradually and do not need to make large adjustments at once, their needs could potentially be met with around 6-8 MW of peak volume available fortnightly.[18] Noting uncertainty around other traders and liquidity conditions, we have assumed 10 MW per fortnight for our modelling—but again emphasise that we have not analysed this in detail.

Again, such a low volume could likely be accommodated by a reduction in volume for the baseload product to 10 MW, and so we assume as much in our modelling.

In the short term, OTC trading may be prudent given generation scarcity. By constraining trading to fortnightly intervals and limiting financial trader participation, market-making activity is more closely tied to underlying physical hedging needs, which market makers can generally manage through their generation assets. The downside is that fortnightly trading limits the ability of participants to

[17] Research suggests that improved price signals can plausibly avoid a modest share of expected shortage, fuel, and carbon costs. The 5% assumption here is conservative and intended to illustrate potential benefits rather than provide a precise forecast. References: Dixit, A. & Pindyck, R. (1994), Investment Under Uncertainty, Princeton University Press; Green, R. & Newbery, D. (1992), Competition in the British Electricity Spot Market, Journal of Political Economy, 100(5), 929–953; Newbery, D. (2018), Electricity Market Design and Investment Incentives, Cambridge Working Papers in Economics.

[18] As previously, this would cover a range of assumptions around the proportion of load hedged and the period over which the hedged position was built, though more extreme assumptions around this could see this range double.

adjust positions, and with fewer financial traders there is less discipline on prices.

In the longer-term, the full benefits of the super-peak product are likely to be realised from daily trading. Unlike the peak product, the super-peak product has the potential to deliver material price-discovery benefits and support more efficient investment decisions. These potential investment efficiency benefits are likely greater than the competition benefits, but they are harder to realise under fortnightly OTC trading.

Daily exchange-based trading allows financial traders to enhance liquidity and correct imbalances between buyers and sellers by taking opposite positions when hedging is one-sided. This helps ensure prices reflect underlying scarcity rather than market makers' inventory pressures. Financial traders also improve price efficiency by arbitraging inconsistencies and reducing inefficient risk premiums across products.

More accurate prices provide a stronger signal of where additional generation is needed, guiding efficient investment. Based on that clear signal, generation developers can use the super-peak contracts to help them finance investment in capacity, aligning incentives so that the market moves toward an efficient balance of supply and demand. Timely participation by both physical and financial participants is therefore important to ensure that scarcity is reflected in prices promptly, strengthening the investment signal.

Under daily trading, the appropriate market-making volume for the super-peak product would likely need to be higher than under fortnightly OTC trading, reflecting a broader pool of participants including financial traders. In the longer-term, we also expect increasing interest from investors in solar and batteries.[19] The volume required would depend on market-wide demand, spreads and on how liquidity evolves across both baseload and super-peak products. In our modelling, we have assumed an illustrative figure of 6 MW for daily trading, reflecting that we expect higher interest than the peak product. As with the peak product, we emphasise the uncertainty around this figure, and note the consequences of setting volume too low may be more significant than setting it too high.

In either scenario, physical supply constraints at peak times mean that market-making settings should reflect the need for all participants to manage the risks prudently. A looser spread than for baseload would incentivise early hedging and enable better inventory management, supporting both market stability and efficient participation.

We again note that demand for the super-peak product is uncertain, and hence have assessed market maker impacts under a range of potential scenarios. Figure 10 shows the potential distribution of market maker earnings from supporting both products under:

- **The counterfactual** – the baseload product is subject to market making with a 3% spread and 12 MW volume.
- **OTC trading case** – Independent retailers shift to their lowest-risk combination of baseload and super-peak accessed OTC, with no other demand changes meaning both products are imbalanced. Volume is provided at 10 MW for baseload and 10 MW fortnightly for super-peak. The baseload product maintains a 3% spread, while the super-peak product has a 5% spread.[20]
- **Balanced exchange-based case** – Daily exchange-based trading of the super-peak product is established. Independent retailers shift to their lowest-risk combination of baseload and super-peak, and other traders emerge to balance supply and

[19] Referrnig to solar investor's potential interest in purchasing super-peak contracts, which we have not modelled.

[20] We have not considered a 3% spread for the super-peak product if subject to OTC trading.

demand for both products. Volumes are 10 MW and 6 MW, and spreads are 3%.

- **Unbalanced exchange-based case** – As above, but the super-peak product attracts half as many sellers as buyers.
- **Suppressed liquidity case** – As above, but spreads for the peak product are increased to 5% so market makers can manage the potential imbalance.

Again, these scenarios are illustrative, rather than likely outcomes.

The outcomes across these scenarios are very similar to those for the peak product, but less severe largely reflecting the lower face-value of the super-peak product given it covers fewer hours. The analysis suggests greater variation in earnings from an OTC trading scenario, but again, the modelling does not fully capture price correlations between products, so this variation is likely over-stated. The moderating effect on demand of the 5% spread is again evident.

Recognising that market makers will be best placed to manage a large proportion of this risk, and will use price and spread to manage inventory in ways not captured in our analysis, the earnings-at-risk for market makers is unlikely to create economic costs that exceed the potential market-wide benefits of enhanced liquidity for the super-peak product.

Figure 10: Outcomes for baseload and super-peak scenarios

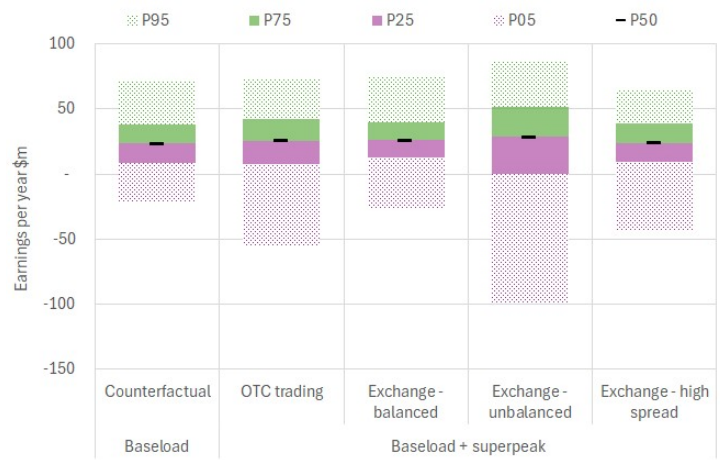

To support both products efficiently, the Authority could consider an arrangement that specifies total market-making volume across all products, while allowing market makers flexibility in how they allocate that volume between baseload, peak, and super-peak contracts. As a purely illustrative example, if the total required volume is 18 MW across all products, the market makers (collectively) could choose to allocate 9 MW to baseload and 9 MW to super-peak (or peak, if applied to that scenario), based on expected demand and their own risk management capacity. This could alternatively be done on a GWh basis.

Such flexibility allows market makers to adjust to changing demand patterns, manage inventory more effectively, and reduce the risk of being over-exposed in any single product. It also ensures that liquidity is provided where it is most needed, supporting efficient price discovery across products. Similar arrangements are used in other energy markets internationally, including certain European intraday markets, where market-making obligations are set in aggregate terms but market makers retain discretion over product allocation to reflect changing conditions.

There are trade-offs to consider. Mandating fixed volumes per product can ensure a minimum level of liquidity, but may force market makers into positions that they cannot risk-manage if actual demand differs from assumptions. Allowing flexibility mitigates this risk, though it requires monitoring to ensure each product maintains adequate liquidity for participants.

6.4 Baseload, peak and super-peak

6.4.1 Economic benefit impacts

Introducing a full suite of baseload, peak, and super-peak products provides only modest incremental price-signalling benefits beyond the baseload and super-peak combination. The super-peak product remains the primary driver of improved price discovery, distinguishing energy- and capacity-driven risks and highlighting periods of high solar variability. The peak product adds limited additional information beyond what is signalled by baseload and super-peak contracts.

From a hedging perspective, the incremental benefit of a third product is minimal. Our earnings-at-risk analysis shows demand-side participants see virtually no further reduction in earnings-at-risk compared with baseload and peak or baseload and super-peak alone, and solar generators similarly gain little additional risk reduction. Overall, introducing a third product is unlikely to deliver greater economic benefits than focusing on baseload and super-peak.

Overall, we would not expect this scenario to deliver greater benefits than the baseload and super-peak scenario.

6.4.2 Balancing costs and benefits

From a market-making perspective, adding a third product alongside baseload and super-peak would increase operational complexity without meaningful benefit. Trading would be spread across three products rather than concentrated in the two that provide most hedging and price-discovery value, diluting liquidity and increasing inventory management challenges and operational costs.

The peak product already trades on the ASX and so it may be possible to achieve liquidity faster than the super-peak product. However, it delivers a weaker price signal, and attempting to develop liquidity across all three products would likely divert market development efforts from the more valuable super-peak product, leading to less efficient outcomes than a two-product approach.

Focusing market-making obligations on baseload and super-peak allows liquidity and spreads to be maintained where they matter most, while keeping volumes manageable under stressed conditions. This approach supports effective risk management and preserves the efficiency of price discovery.

Given the expectation that this scenario would achieve lower benefits at higher cost, we have not modelled the market maker impacts in detail.

6.5 Summary of four scenarios

Overall, our analysis indicates clear differences in price discovery, investment efficiency and retail competition across scenarios, as summarised in Table 1. We have suggested the market making arrangements that may best enable the benefits to be realised, but these should be considered indicative.

Table 1: Summary of benefits and market making arrangements for each scenario

Baseload	
Estimated potential benefit	Negligible
Market making arrangements	3%, 12MW (though a reduction could be considered)
Baseload + peak	
Estimated potential benefit	$3m per year by 2032 (retail competition benefits)
Market making arrangements	Peak: 5% spread, 4 MW volume daily Baseload: 3% spread, 10 MW volume daily
Baseload + super-peak	
Estimated potential benefit	Up to $10m per year by 2032 made up of: • $3m retail competition • $1.2m generator financing • $6.25m avoided shortage/fuel costs
Market making arrangements	Super-peak: 5% spread, 10 MW volume fortnightly Baseload: 3% spread, 10 MW volume daily But should transition in time to exchange trading with appropriate volumes to be determined.
Baseload + peak + super-peak	
Estimated potential benefit	Not explicitly considered
Market making arrangements	Not explicitly considered

7 Overall recommendations

In setting market-making arrangements, a balance must be struck between facilitating cost-effective trading for participants that rely on hedges, supporting transparent forward price signals, and ensuring the commercial sustainability of market makers.

Across the four scenarios examined, our analysis suggests that a combination of **baseload and super-peak** products provides the greatest benefits, and that the risks to market makers of supporting two products are likely manageable through spread and volume settings.

For baseload, improvements in trading justify maintaining the current 3% spread. Volume reductions could be considered to help to mitigate market maker stress during volatile periods, and to better balance volumes with a second product. Both costs and benefits of extending the trading horizon appear modest, though the overall balance remains unclear without better information on likely buyer and seller interest.

For the super-peak product, we consider OTC trading is a practical initial step given constrained physical supply and demand for the product that will likely build over time. We indicatively suggest a 5% spread and 10 MW volume for fortnightly OTC trading, but note we have not assessed these settings in detail and they are subject to uncertainty around trader interest and liquidity impacts. We expect these settings will allow market makers to manage earnings-at-risk from trading the new product.

Earning-at-risk can be more significant for any product with imbalances between buyers and sellers, which may be reflected in product prices. There may be a particular risk of imbalanced trading for the super-peak product in the short term due to constrained generation. Furthermore, OTC trading limits participation by financial traders who can help to balance liquidity and support efficient prices. The super-peak product can play a key role in providing price signals that encourage the right generation to be built at the right time. We expect this to be the primary source of benefit from developing liquidity in the product. We therefore consider that the superior liquidity and price discovery potential of daily trading makes a transition to ASX trading preferable in the longer term.

Overall, these arrangements are expected to support competition, improve price discovery, enhance operational efficiency, and maintain market maker viability, while retaining flexibility for future adjustments.

Appendix A. Product profile analysis

7.1 Signal-to-noise analysis of price-signalling value

This appendix provides more detailed results from our signal-to-noise ratio analysis, which is described in the body of the report.

Figure 11 shows the SNR for 2026-2028. It distinguishes between weekdays and weekends in this analysis.[21]

Figure 11: Signal to noise ratio of each hour relative to baseload

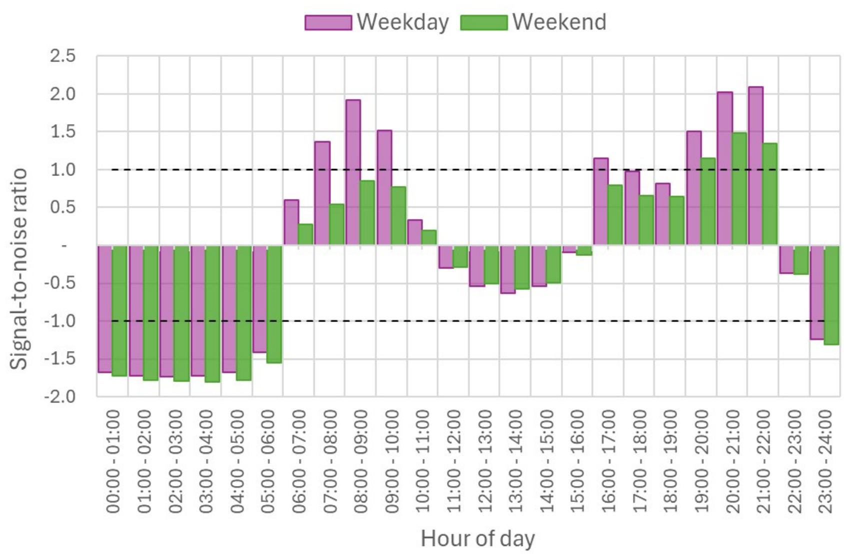

Figures 12 and 13 break this down into 'signal' and 'noise' to interrogate the drivers of the results.

Figure 12: Average difference in price for each hour relative to baseload

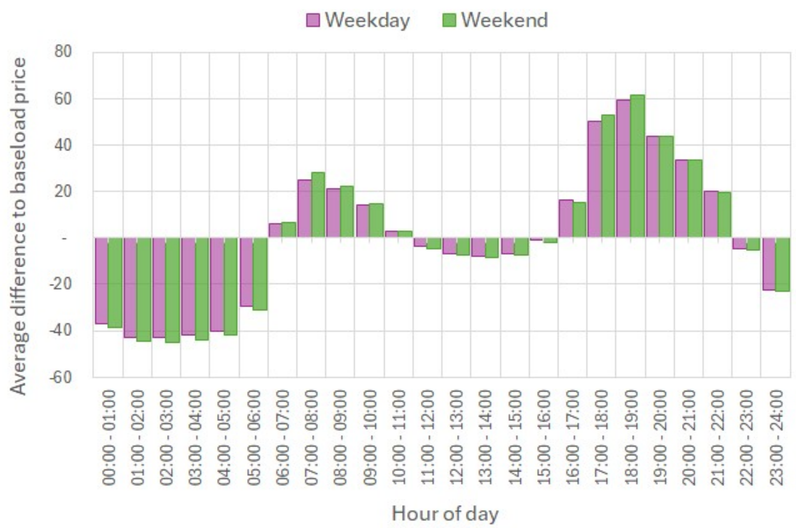

This shows the difference from baseload prices is most significant in the evenings—particularly from 5-7pm—and overnight. There is a moderate difference for morning hours. There is only a small difference for midday hours.

21 We have run the analysis for individual calendar quarters and separately for each year. The results are amplified in some quarters relative to others, and become more amplified over time, but the hours and days are consistently identified.

Figure 13: Variation (standard deviation) in price difference for each hour relative to baseload

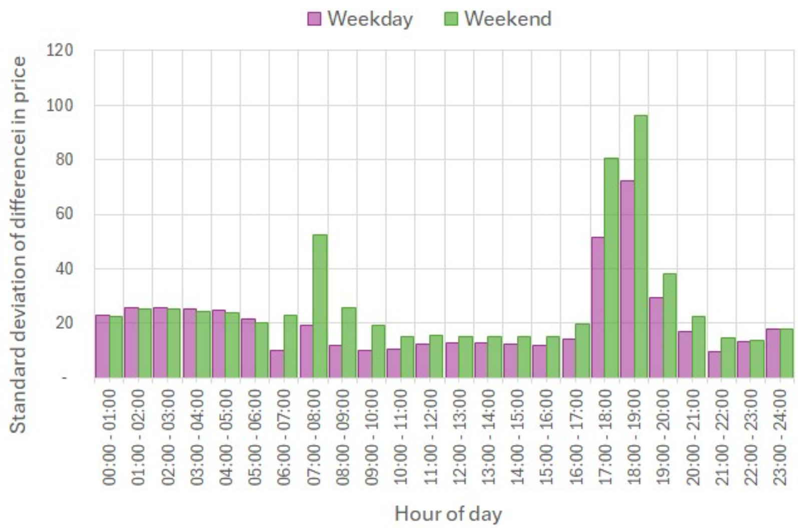

This shows very high standard deviation for the hours of 5-7pm. The deviation is higher in weekends, particularly during peak hours. The deviation is lowest for midday hours.

7.1.1 Longer-term variation in price-signalling

The above analysis reflects our price forecasts for 2026-2028. We also forecast longer-term prices that help highlight challenges to hedging and signalling prices. Figure 14 and Figure 15 show the results of the same SNR analysis for 2032 and 2037 respectively.

Figure 14: Signal to noise ratio for 2032 forecasts

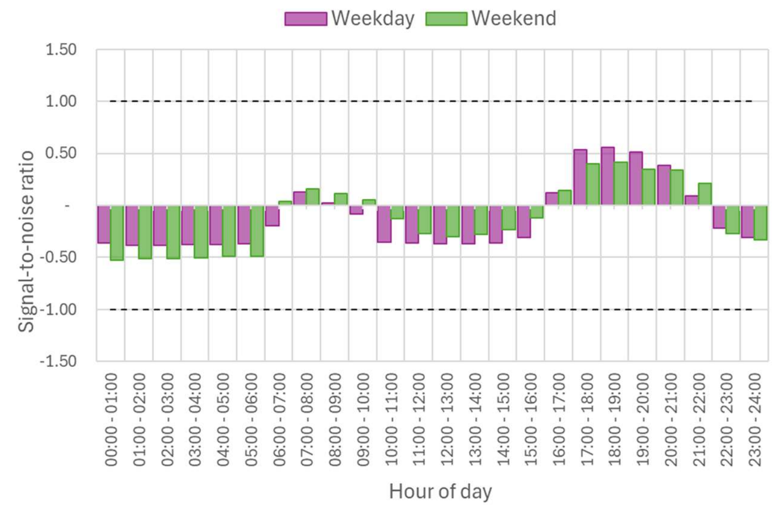

Figure 15: Signal to noise ratio for 2037 forecasts

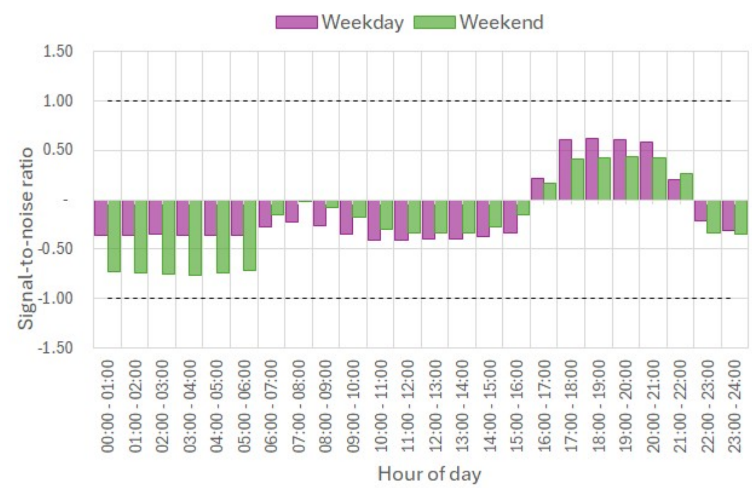

There are two clear effects in these graphs:

- the morning peak becomes much more subdued over time. This primarily reflects the potential for evening peak demand growth from EV charging to become a significant feature of the overall demand (and hence price) profile
- the signal in general is much more muted, with all hours showing a SNR between 1 and -1. This reflects that the variation in prices relative to baseload in all hours becomes much larger as prices become 'peakier', given increasing intermittent generation and demand. The volatility will mean traders find it much harder to price hedge products, and risk premiums will likely increase.

We understand ASX is considering revising the products on offer for Australia's National Energy Market – disestablishing the existing peak product in favour of separate morning and evening peak products. Our analysis would indicate that such an approach could have merit in the longer-term if the dynamics in our price forecasts come to fruition.

For 2037, we again also show the 'signal' and 'noise' separately in Figure 16 and Figure 17—noting that we have not shown this for 2032 as it merely shows a mid-point in a journey toward these outcomes.

Figure 16: Average difference in price for each hour relative to baseload for 2037 forecasts

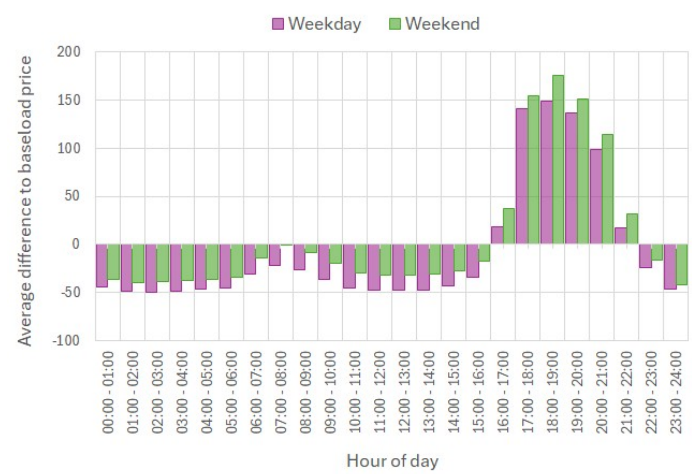

Figure 17: Variation (standard deviation) in price difference for each hour relative to baseload for 2037 forecasts

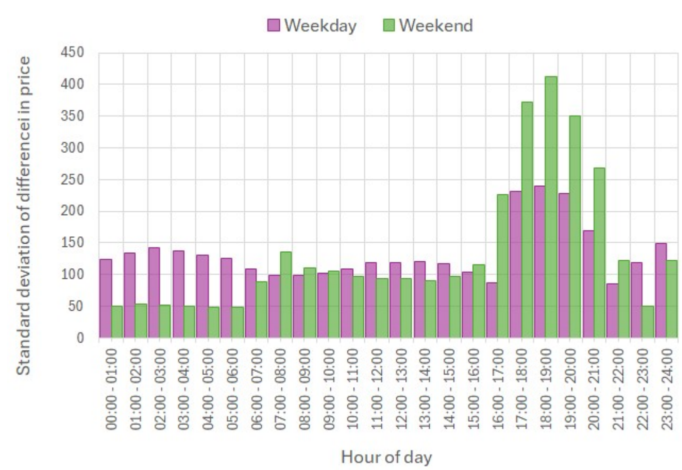

Appendix B. Earnings-at-risk model

The EAR model allows evaluation of the effect of different quantities and types of hedge (base, peak, super-peak) on the earnings of electricity market participants (retailers or generators) for a given future year.

The key data inputs to the model are:

- A price series reflecting the likely range of price outcomes across the year.
  - A number of different price series have been used, coming from Concept's electricity price forecasting model, ORC, modelling different future years and, in some cases, different future system states (eg, a market short of generation). Each future year is represented by 43 'weather years', and with 8,760 hours for each weather year. The weather years are based on actual concurrent hydro inflows, windflow, and sunshine for the historical years 1980 to 2022.
  - In addition to these forecast price series, there is a historical price series using actual spot market prices for the years 1998 to 2024, but with prices for historical years inflated by CPI to convert them all to real $2025.
- A concurrent demand or generation series.
  - These come from the ORC model. The demand series is a so-called 'conforming' demand which simulates a mass-market demand shape. ORC has a model within it which simulates the extent to which there are correlations between renewable flows and weather-driven variations in demand.
  - The model has been structured to enable examination of different demand shapes (eg, flat), or different generation profiles (eg, wind), or different combinations of retail sales and generation (eg, some level of vertical integration). However, it is out of scope for this phase of the engagement to extend the analysis to these situations.
  - To enable some level of consistency in comparing outcomes, each demand or generation shape is factored such that the average level of demand or generation is 10 MW.

For each hour of each of the 43 weather years, the model calculates the amount earned in that hour.

When configured to examine outcomes for an independent retailer, the net earnings equals the sum of:

- retail sales, being the quantity of demand multiplied by the retail tariff. The retail tariff is calculated as being the demand-weighted average price across the 43 weather years for the relevant TOU time block.[22]
- hedge purchases, being the quantity of hedge purchased multiplied by the hedge price.
  - the hedge price varies by quarter, and is the time-weighted average price across the 43 weather years for the relevant time period:
    - baseload (All hours)
    - peak (6am to 9pm)
    - super-peak (6am to 10am and 5pm to 9pm)
  - variation in within-week coverage

[22] A default summer-winter / peak off-peak structure has been used, but different tariff structures can be used (eg, just peak/off-peak, or flat across the year).

- all days for all products, or
- weekday only (only for peak and super-peak
- spot sales / (purchases), being

  ([Hedge quantity]-[Demand quantity]) * [Spot price]

The structure is fundamentally the same for an independent solar generator, except that

- there are no retail sales
- the generator is *selling* a hedge, and
- it is making spot purchases / (sales) calculated as

  ([Hedge quantity]-[Generation quantity]) * [Spot price]

For a given hedge position (ie, quantity and type of hedges purchased) the model:

- sums the net earnings for each of the 43 weather years
- calculates the mean annual earnings across all years
- calculates the range of annual earnings for different percentage likelihoods:
  – P0, being the worst annual earnings
  – P5, being the 5th percentile worst annual earnings
  – P95, being the 95th percentile best annual earnings
  – P100, being the best annual earnings.

To examine the effect of different quantities and types of hedges, the model cycles through lots of different combinations of hedge quantities and hedge structures. The model examines the effect of having access to different combinations of hedge types. Four combinations are considered:

- Baseload only
- Baseload + Peak
- Baseload + Super peak
- All three

For each combination, the model cycles through lots of permutations for the quantity of hedge purchased of each hedge type. As the quantity of hedge varies, the earnings at risk varies. This variation in earnings with quantity of hedge is illustrated by Figure 18 which shows the variation in earnings for different quantities of baseload hedge purchased for a retailer. The MW quantity is expressed as a % of average MW demand for the relevant quarter.

Figure 18: Earnings variation with different quantities of baseload hedge for a mass-market demand shape

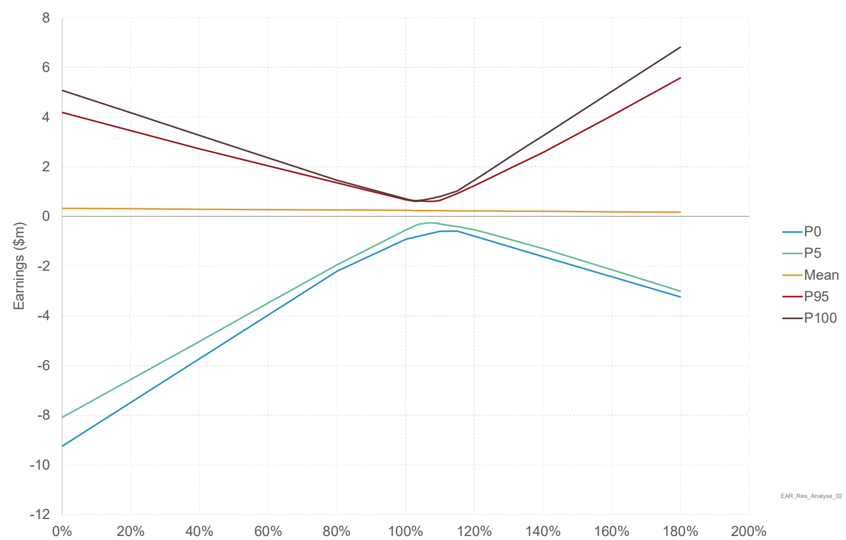

The earnings-at-risk is calculated as the difference between the mean earnings, and the P5 earnings.

The hedge level that achieves the minimum EAR is calculated (eg, 107% in the example above), along with the minimum EAR level.

This is repeated for each of the hedge type combinations. It is not readily possible to graph this variation in EAR in situations with multiple different hedge types, although this ‘surface plot’ shows how the EAR varies with different combinations of the base and super-peak product, and how the model samples different combinations at a relatively coarse level, and then progressively reduces the hedge variation increments around the area of EAR minimum for the coarse sampling to arrive at a more optimum level at a finer hedge granular level.

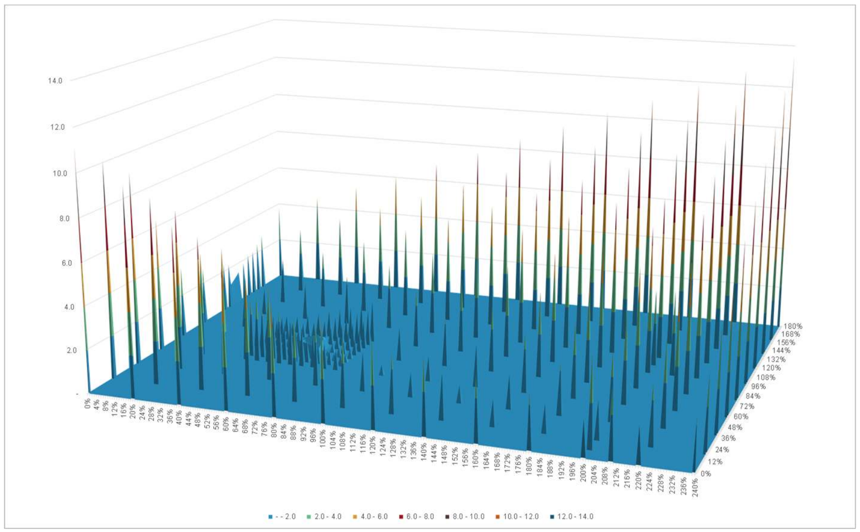

Appendix C. Market maker earnings at risk model

This appendix describes the modelling framework used to simulate the economics of market making. The framework has four linked components: price simulation, trade allocation, earnings calculation, and aggregation. Together these provide a stylised but internally consistent view of market maker performance.

1. Price Simulation

Prices are projected using a 'random walk' from the starting price—being the average price for our central scenario when the forecast was set— through to an end-point being the average forecast price outcome for a quarter, for each weather year and simulation considered.

These forward price paths are generated using a stochastic process designed to capture both systematic market shocks and contract-specific variability. The model assumes that:

- Market-wide 'random' shocks affect all contracts within a weather or scenario year, while per-contract 'random' shocks drive individual deviations.
- Price volatility varies by horizon, reflecting that shorter-dated contracts tend to exhibit greater relative movements. The volatility of peak and super-peak products is simplistically assumed to be 120% of that for baseload.
- Prices gradually converge toward the end-point, with random variation intended to simulate changing market expectations as new information becomes available over the life of the contract.
- No explicit drift is assumed beyond these target levels; randomness is purely driven by volatility and shocks.

2. Trade Allocation

Volumes of product bought and sold are modelled based on a random allocation over time. Assumptions include that:

- Demand is based on historic trading volumes between all market makers and any other party – but excluding trades between market makers. Monthly baseload volume is simplistically captured as an increase in quarterly baseload volume.
- Demand is sensitive to the spread, reflecting historical trading volumes pre and post change from 5% to 3% spread.
- Maximum volumes are sensitive to the volume of trading 'churn' between market makers, which is assumed to occur in advance of other trading, and is also sensitive to the spread.
- Volumes are allocated on four time horizons; 2 quarters ahead, 2-4 quarters ahead, 4-8 quarters ahead; 8+ quarters ahead – in proportions that reflect historic trading.
- Volume is allocated between islands based on historic trading
- If the random allocation on any day exceeds the effective volume available under the cap (ie after accounting for market maker trades), that volume is not relocated.

3. Earnings Calculation

Contract-level earnings are calculated using an ASX-based margining framework, capturing a combination of transaction profits, daily mark-to-market (M2M) valuations for inventory held, fees and initial margin financing requirements. Assumptions include:

- The effective spread captured is simplistically assumed to be 75% of the maximum market maker spread required.

- Margin requirements and fees depend on contract type, horizon, and location, reflecting ASX specifications.[23]
- Financing costs accrue daily at a fixed cost of capital of 6%, providing a stylised view of the cost of carrying inventory.

4. Portfolio Aggregation

Earnings are summarised to provide a representative view of a market maker's performance over a notional portfolio. A significant limitation of the model is the limited dataset, as new contracts do not roll on continuously as others expire. An approximation of the annual earnings is hence made based on:

- The average daily earnings per contract over all contracts modelled.
- Annualised based on 260 trading days per year
- Multiplied by the number of contracts market made when a full portfolio is available (averaged at 30 for a single product over both Otahuhu and Benmore)

It is noted that this approach will over-represent the earnings from contracts in their most volatile period, and under-represent the earnings from contracts in a more stable long-dated-period.

• High-Level Model Assumptions

Across all components, the model embodies several core assumptions:

- Market makers are modelled as a single, passive participant that takes on the inventory given to it by other traders. The market maker does not take any actions to manage or offload inventory, and there is no feedback loop to price.
- Liquidity (in terms of spreads and volumes available) are assumed consistent on each trading day. Trade allocation is rule-based and does not adapt dynamically to market depth or competitor behaviour.
- Price dynamics are independent of trading volumes and are not responsive to market supply and demand or other trading factors.

Overall, the model is intended to illustrate plausible ranges of earnings under assumed spreads, volatilities, and portfolio sizes, rather than precise forecasts, and are for illustrative purposes only. The model was run with actual historic price and trading data in order to calibrate it to within a functional level, while recognising that historic outcomes are not necessarily representative of future outcomes.

[23] https://www.asx.com.au/documents/clearing/asx-clear-futures-energy-margining-example.pdf