PRINCIPAL ECONOMICS

The standardised super-peak hedge product

Technical addendum: response to consultation questions

REPORT TO

Electricity Authority Te Mana Hiko

www.principaleconomics.co.nz

SEPTEMBER 2025

Contact person

Dr Eilya Torshizian, eilya@principaleconomics.com.

Authors

This report was prepared by Dr Eilya Torshizian, with the assistance of Ruyi Jia.

Principal Economics provides independent economic consultancy services to a wide range of public and private clients.

At Principal Economics we help our clients to find practical robust solutions in a timely manner by prioritising clients' problems, using frontier economic thinking and our familiarity with a wide range of data and methodologies.

For details on our work, see our website: https://www.principaleconomics.com/

Address: Level 17, 55 Shortland Street, Auckland 1010, New Zealand

Executive summary

The Electricity Authority tasked Principal Economics to evaluate the appropriate minimum trading volumes and maximum bid-ask spreads for market-making on a standardised super-peak hedge contract. The purpose of our report, available here, was to provide an evidence-based framework to guide regulatory options that enhance price transparency, improve market liquidity, and deliver lasting consumer benefits. By modelling the costs and benefits of different market-making scenarios, the report supports the Authority's consultation process in developing a robust framework that balances market-maker incentives with broader system efficiency, ensuring sustainable and effective liquidity provision in New Zealand's electricity hedge market.

This document is an addendum to our main report and responds to the questions raised in relation to the Authority's consultation paper *Regulating the Standardised Super-Peak Hedge Contract* (19 August 2025). The consultation paper drew on our analysis and advice regarding the appropriate maximum bid-ask spreads and minimum trading volumes for market making. This addendum clarifies the key points and addresses specific queries.

Our responses to the questions do not introduce any additional limitations beyond those already highlighted in the report (pp. 36–38). Instead, they reinforce the importance of the three caveats we identified and explain how these were treated in our analysis:

1. **Small number of standardised super-peak observations**
   The standardised super-peak data provides the only direct evidence of spreads in the fortnightly super-peak sessions. At the time of our report, only a limited number of observations were available (covering the first half of 2025), which restricted the ability to run statistically robust regressions or test behaviour under varying market conditions. We therefore used the data primarily to anchor spread levels (e.g. median ~3.7%, 95th percentile ~9–10%) and to estimate elasticities specific to the super-peak product, while relying on ASX data for validation. Although these regressions are less statistically robust than those from ASX, they provide product-specific evidence that grounds our recommendations. We recommend re-visiting this analysis once a larger sample of sessions has accumulated — for example, after a year of trading — to test the robustness of elasticities and assess behaviour under a wider range of market conditions.

2. **Reliance on ASX data as a proxy for validation**
   We also conducted regression analyses using ASX data, which has a significantly larger sample and encompasses a wider range of conditions. The ASX results showed elasticities consistent in sign and magnitude with those from the regression results using the standardised super-peak data. This improved our confidence in the robustness of the elasticities using the standardised super-peak data despite the small sample size. ASX, therefore, served primarily as a validation dataset, demonstrating that the standardised super-peak findings were consistent with broader market behaviour and reliable enough for policy calibration.

3. **Simplified cost structure**
   Our stylised inventory cost function scales with spread, volume, mid-price, and cost of capital. Volatility was captured indirectly through the spread function, but the cost function did not explicitly model carry risk between sessions, operational frictions, or dynamic inventory management. This simplification was necessary given the data constraints and the need for a tractable cost-benefit framework. We addressed the issue by:

   - Acknowledging this limitation explicitly in our report.
   - Highlighting that future refinements should incorporate explicit volatility-carry terms, regime-dependent effects, and operational frictions.

  - Stress-testing results across volatility bands to demonstrate the direction of the effect (spreads widen significantly under high volatility), which is already reflected in the recommended tiered spread caps.

Overall, these steps ensured that while the analysis was subject to the recognised caveats, it remained robust for the purposes of informing policy design. At the same time, the limitations point to clear areas for further empirical research that could refine and strengthen future work.

Other limitations and topics for future investigation that were already highlighted in our main report

While our report used the best available information and methodology, there are a range of limitations as follows:

- **No direct super-peak transaction data**: Analysis relies on proxies (high-price periods and auction logs), as real-time super-peak contract trading is limited.
- **Unobserved contract positions**: Market-maker risk exposures (e.g. generator-retailer hedges) are not visible, limiting precision in incentive modelling.
- **Simplified cost function**: Inventory cost is based on volatility and depth, but omits behavioural strategies (e.g. quote withdrawal, position hedging).
- **Social welfare impacts not quantified**: While private net benefit is modelled, broader system-wide effects (e.g. on competition or end-user prices) are not included.
- **No dynamic strategy modelling**: The analysis does not capture potential strategic behaviour under different quoting frequencies or cap structures.

Below is a list of topics for future research:

- **Empirical testing with super-peak implementation data**, if/when contracts are actively traded.
- **Integration of physical capacity constraints** to account for system stress effects on liquidity and pricing.
- **Incorporation of risk-aversion and strategic quoting models**, using agent-based or game-theoretic frameworks.
- **Social CBA extension** to measure transparency, competition, and downstream consumer impacts.
- **Behavioural and timing dynamics** using Granger causality or high-frequency event study methods.
- **Instrumental variable approach** to infer causal impacts.

We should also note that once an S% cap is defined, e.g. 8%, market makers may anchor their quotes at or near the cap, even when tighter spreads would be feasible. The Authority should accompany implementation with monitoring tools to detect clustering around the cap and reassess cap calibration as market behaviour evolves.

Contents

1 Response to questions

Electricity Authority (the Authority) commissioned Principal Economics to provide analysis and advice on appropriate minimum trading volumes and maximum bid-ask spreads for market making of a standardised super-peak hedge contract. The report examined the costs and benefits of various regulatory settings under different market conditions, enabling the Authority to determine arrangements that promote price transparency, market liquidity, and long-term consumer benefits. The findings will inform an upcoming consultation on options for implementing market-making on both over-the-counter and exchange platforms.

This addendum to our main report outlines the questions and provides our responses. Our response does not include the limitations and their implications that have already been highlighted in our report (pp. 36-38). The questions are presented in each section below, together with our response.

1.1 Question 1: How has the analysis accounted for the fortnightly cadence of super-peak trading sessions in its appropriateness of spread?

Response:
The standardised super-peak auction log data reflects the fortnightly cadence, with median spreads of ~3.7% and a 95th percentile of ~9.5%, providing direct evidence from super-peak sessions. To strengthen confidence in these results, we cross-checked them against ASX trading data. While the standardised super-peak observations are limited and do not capture variation under different market conditions, the ASX data provides a broader range of outcomes across liquidity and volatility scenarios. The consistency between the two datasets supports the robustness of the standardised super-peak based results, with ASX data used to capture the necessary variation for policy testing.

The standardised super-peak data has *limited observations*, restricting the ability to test spreads under different market conditions. The implication is that our findings should be seen as indicative of cadence, with ASX used to capture variation (limitation: *data sparsity*).

1.2 Question 2: How has the analysis accounted for the volatility of the base–super-peak relationship when using ASX base data?

Response:
The analysis recognises that super-peak contracts display different volatility characteristics from base contracts. To address this, elasticities of spreads with respect to volatility and depth were estimated using the standardised super-peak auction log (standardised super-peak data) (pp. 24–26). These regressions provide product-specific evidence, but the limited number of observations means the results are less statistically robust across different market conditions. To validate these findings, we also ran regressions on the ASX dataset, which has a much broader sample. The ASX results were consistent in sign and magnitude with the standardised super-peak regressions, giving confidence in the robustness of the elasticities based on the standardised super-peak data. As noted in the report (p. 26), this approach grounds the modelling in super-peak-specific conditions while drawing assurance from the ASX results. This balances product relevance with statistical reliability, supporting the Authority's objective of aligning market-making requirements with observed market behaviour.

In short, the elasticities derived from the standardised super-peak data form the basis of the modelling, while the ASX results serve as an important validation step. The limited standardised super-peak data sample size at the time of our analysis is flagged in the report as a key area for further investigation (pp. 36–38).

1.3 Question 3: How does the known need for higher base market-making compensation in NZ align with the guidance for super-peak spreads/volumes?

Response:
The commercial procurement process confirmed that base market making requires significant compensation, indicating that super-peak would ordinarily demand even more. However, the standardised super-peak auction data shows relatively narrow median spreads (~3.7%), suggesting that when market makers are already exposed to spot positions, the marginal cost of quoting super-peak may be lower, provided volumes are modest. The analysis, therefore, recommends volatility-responsive caps rather than fixed thresholds. This approach supports the Authority's policy objective of balancing incentives with sustainable liquidity, ensuring that market makers are adequately compensated while enhancing transparency and efficiency in the hedge market.

In short, narrow spreads suggest lower marginal cost at modest volumes, but this inference rests on small-sample of the standardised super-peak data. The implication for our report is that more empirical testing is needed to confirm whether this holds consistently across regimes.

1.4 Question 4: Why does net benefit increase sharply as volume rises to ~10–15 MW? And why is this relevant in NZ?

Response:
The cost-benefit model shows that net benefit rises non-linearly with volume, as revenue (spread × volume) scales faster than costs, which increase with spread and volatility but not linearly with volume. For market makers already holding firming capacity (e.g. hydro, gas), incremental super-peak sales generate additional revenue with limited added risk. In New Zealand, where gentailers are vertically integrated and already manage super-peak risk, this dynamic is particularly relevant. It reflects realistic commercial incentives and supports the Authority's objective of encouraging natural suppliers to provide sustainable liquidity in the hedge market.

1.5 Question 5: Assumptions about the profit and the inherent trading risk in the volatile super-peak market

> **Question:** A key assumption of the paper outlines a benefit to market makers based on the bid-ask spread multiplied by the volume (Equation 3.3). We would appreciate a more detailed explanation of the assumed benefits to market makers to help us understand this better (including an example). We take that the paper assumes that a market maker gets in and out of trades at a profit. This assumption may be valid in a market with high liquidity and relatively low volatility. In our experience, trading in and out of a position around a stable mid-price is extremely rare in the NZ electricity futures context and is not a valid assumption for super-peak trading events. If the assumed benefits did not hold, how would that alter the findings of the report?

Response:
Equation 3.3 (p. 12) defines revenue as the product of spread and volume. This is not intended to imply that market makers always trade in and out of positions at a guaranteed profit, but rather to formalise how revenues would scale with volumes given observed spreads. The corresponding cost function (p. 15–16), based on elasticities estimated from the **standardised super-peak auction log (SPF data)**, captures how volatility and depth influence spreads. While the SPF sample is limited, these regressions are consistent with results from the broader ASX dataset (pp. 24–26), which provides important validation. The difference between the revenue and cost functions represents an indicative net benefit, not an assumed profit margin. As noted in the report (pp. 36–38), the simplified cost structure remains a key area for further investigation. If the assumed benefits did not hold under all conditions, the implication is that

improved cost modelling would be required, but the overall framework for evaluating the trade-off between liquidity provision and market-maker incentives would remain valid.

Here's a simple worked example that illustrates how Equation 3.3 operates and why net benefit can rise sharply with volume, without assuming "guaranteed" round-trip profits:

Assumptions (illustrative only):

- Mid-price = $200/MWh.
- Observed median spread of standardised super-peak = 3.7% → $7.40/MWh (0.037 × $200).
- Consider three quoting sizes for the same session: 5 MW, 10 MW, 15 MW.
- Costs reflect inventory/hedging and adverse-selection risk that rise with volatility and spread but less than proportionally with volume (consistent with using elasticities estimated on ASX and applying the standardised super-peak levels). For illustration, take indicative total costs of $22, $38, $50 per MWh across the three volumes, respectively (showing sub-linear scaling).

Revenue (Eq. 3.3):

- Revenue per MWh = spread × mid-price = $7.40.
- Session revenue scales with volume:
  - 5 MW → $37.0 per MWh-equivalent of quoted energy
  - 10 MW → $74.0
  - 15 MW → $111.0

Indicative total cost (same session):

- 5 MW → $22
- 10 MW → $38
- 15 MW → $50

Indicative net benefit (Revenue – Cost):

- 5 MW → $15.0
- 10 MW → $36.0
- 15 MW → $61.0

What this shows:

- Equation 3.3 does not assume risk-free round trips; it just formalises how revenue would scale with observed spreads.
- Net benefit depends critically on the cost function (which is where uncertainty lies and where further refinement matters).
- When costs rise less than proportionally with volume (e.g., inventory netting, pre-existing firming), net benefit can accelerate as volume increases—consistent with our finding that benefits rise sharply up to ~10–15 MW.
- The report's framework uses elasticities estimated from the SPF regressions, which are product-specific, with the ASX results used as validation given their broader sample size and consistency in sign and magnitude. This

means that results are anchored in revealed super-peak behaviour, drawing confidence from the ASX data, rather than relying on the assumption of always profiting around a stable mid-price.

We could adapt the same arithmetic to a different mid-price (e.g., $150 or $250/MWh) or to alternative spread points (e.g., the 95th percentile) to further test sensitivity.

1.6 Question 6: The cost function and the consideration of volatility

Question: The cost function represented in Equation 3.4 does not appear to represent any costs associated with volatility. In our experience, high volatility has potential to significantly increase the cost of market making, especially in instances where open inventory is carried from one market making session to another. Are there ways to correct costs for volatility and how would that alter the findings of the paper?

Response:
Equation 3.4 represents a simplified inventory cost, scaling with relative spread, volume, mid-price, and cost of capital. Volatility influences this cost indirectly through the spread function (Eq. 3.1), which is estimated from ASX data and applied to the standardised super-peak observations. However, as acknowledged in the report under "3. Simplified cost structure" (pp. 3–5, 36–38), this specification does not capture all volatility-driven risks — such as carrying inventory between sessions, operational frictions, or regime-dependent effects. This simplification reflects both data constraints (limited standardised super-peak observations and the absence of carry dynamics) and the need for tractability in the cost-benefit framework.

As discussed on page 38, incorporating explicit volatility-carry terms and operational frictions would shift the net-benefit curve downward in high-volatility periods, making tight spread caps less sustainable. This reinforces, rather than alters, our recommendation for volatility-responsive caps over fixed thresholds. The framework remains valid, but the report highlights improving the cost function as a priority for further investigation.

1.7 Question 7: If volatility assumptions are adjusted to reflect the fortnightly cadence of super-peak trading events, would this change the recommended bid-ask spreads?

- Question: NZ electricity products that trade on ASX are extremely volatile. Based on our experience, there is a strong link between ASX baseload pricing and super-peak pricing, implying that super-peak products should be at least as volatile, if not more so, as their baseload equivalents. The recommended approach from the Authority is to hold fortnightly trading events for super-peak products. Holding super-peak trading events at fortnightly intervals effectively increases the risk for market makers from holding inventory from one super-peak event to the next due to the potential for prices to move significantly. If the volatility assumptions of the paper are adjusted for the time between super-peak trading events would the recommendations on the bid-ask spread change?

Response:
The report's recommendations are already volatility-responsive, with tiered spread caps calibrated to observed volatility bands (exec. summary). We addressed the impact of volatility through three complementary analyses, each reinforcing the same conclusion that higher volatility drives wider spreads:

1. **ASX baseline analysis (pp. 24–25):** Regression results show a clear positive elasticity of spreads with respect to volatility. For example, a 1% increase in volatility is associated with a measurable widening of spreads, while depth has a negative effect (elasticity ≈ −0.08). This establishes the general link between volatility and spreads.

2. **High-price ASX analysis (pp. 28–29):** Under stressed, high-price conditions, the effect of volatility is much stronger. Here, a 1% increase in volatility raises spreads by ~0.37%, compared with essentially no effect in other regimes. This demonstrates that spreads are particularly sensitive to volatility shocks in stressed conditions.
3. **The standardised super-peak data (p. 21):** The SPF auction logs provide direct evidence of spreads in fortnightly super-peak sessions, with a median of ~3.7% and a 95th percentile near 9–10%. Although limited in observations, the distribution shows that spreads widen materially during stress events, consistent with the ASX results. The SPF auction data was the *primary anchor* for our final recommendation on the spread caps, because it directly reflects super-peak trading sessions and their fortnightly cadence.

By triangulating these three analyses, the report demonstrates both statistical robustness (ASX) and product specificity (SPF). Together, they provide strong evidence that increased volatility leads to wider spreads and that this effect is amplified in stress regimes. Adjusting volatility assumptions to reflect the fortnightly cadence of super-peak events explicitly would therefore strengthen our recommendations: spread caps should be volatility-responsive rather than fixed, ensuring sustainability of market-making while aligning with observed market behaviour.

In short, the SPF data cannot capture cross-session inventory carry, and fortnightly frequency increases exposure to price shocks. Our report's volatility-responsive caps already account for this in direction, but precise calibration needs further investigation.

2 Conclusion, limitations and recommendations

This section provides a recap of the conclusion, limitations, and recommendations for future research of our main report. Our report developed an evidence-based framework to advise the Electricity Authority on appropriate bid-ask spread caps (S%) and minimum quote depths for market-making obligations on the standardised super-peak hedge contract. Our approach combined econometric analysis of ASX base and peak load futures with observed pricing behaviour from auction log data for super-peak products. Empirical analysis of the auction log data revealed that typical relative spreads for super-peak contracts are tight under normal conditions—median S% was 3%. However, spreads widen markedly during stress periods, with the 90th percentile exceeding 5% and the 95th percentile reaching up to 9–10%. To keep the cap both practical and responsive, we recommend a clear two-tier policy:

- a base cap of **3.5%** under normal conditions; and
- an elevated cap of up to **8%** when volatility exceeds a defined threshold or during designated stress events.

This structure balances cost-effective market-making with dependable liquidity and price transparency for hedge-takers in the electricity futures market.

2.1 Limitations and recommendations for future investigation

While we have applied best-practice methods and used all available data sources to inform this analysis, the following caveats apply. This section provides a summary of the caveats and encourages future work on these topics.

- The consideration of the social benefits of liquidity is beyond the scope of our report.

  We suggest that while **greater liquidity** (i.e. tighter spreads and deeper markets) is often assumed to improve market function, the **social benefit of increased liquidity** depends on:
  - **Who benefits** from improved liquidity (e.g. small retailers, large gentailers, industrial users).
  - **Whether lower hedge costs** for buyers are passed through to **final consumers**.
  - **The structure of the market** — particularly in NZ, where **gentailers internalise generation and retail risk**, possibly leading to **efficient hedging via vertical integration** rather than financial market making.
  - **The marginal cost of achieving that liquidity** — e.g., costs imposed on market makers (possibly gentailers), which could reduce incentives for long-term investment or distort risk allocation.

  Improved access to hedging may reduce retail price volatility for some participants, but it may also encourage entry of less efficient retailers, distort investment signals, or duplicate the risk management already internalised by vertically integrated gentailers. The implication of this for the Authority is that the **social benefit of liquidity is context-dependent**, and **not necessarily positive**. In some cases, **excess liquidity could redistribute risk inefficiently**, or **support market entrants** whose participation leads to **higher system costs** (e.g. through volatility amplification or overreliance on financial hedging). Hence, we suggest a future project to investigate the social value of liquidity. That study should use a general equilibrium approach to capture impacts across markets.

- **Lack of directly observed super-peak data**: The ASX Information Services dataset does not include actual super-peak contracts, requiring the use of high-price proxies and supplemental the SPF auction data. This limits precision when extrapolating to super-peak conditions.

- **Limited visibility of contract positions**: Contractual exposure, such as physical generation hedges or portfolio obligations, can materially affect market-making behaviour. These positions are not observable in the available data and were therefore excluded, but warrant further investigation.

- **Omission of supply-side constraints**: The analysis does not explicitly account for physical capacity constraints during peak periods, which may influence both pricing and quoting behaviour. This could be explored through integrated market simulation or dispatch models in future work.
- **Exponential risk pricing under stress**: Extreme market conditions may result in nonlinear pricing dynamics (e.g. jump diffusion, regime shifts). While we incorporate volatility in our cost modelling, additional work using disaster or rare-event modelling could enhance robustness.
- **Social benefit estimation is simplified**: The analysis quantifies private net benefit to market makers. Broader social welfare impacts — such as price transparency, risk-sharing, or retailer competition — are acknowledged but not modelled. These require value judgments and assumptions beyond this report's scope.
- **Trading frequency and behavioural effects not directly modelled**: While we distinguish between daily (exchange) and fortnightly (OTC) regimes, strategic behaviours (e.g. quote fading, market timing) are not explicitly modelled, and could influence real-world performance.

2.2 Further discussion of the implications of the limitations

While the analysis is grounded in robust data and modelling, several limitations were acknowledged. These limitations do not invalidate the findings but help frame their **appropriate interpretation and use** for policy development. In this section, we further explore the potential impacts of the limitations on our findings and recommendations.

1. Lack of actively traded super-peak contract data

- **Implication**: Spread and depth estimates for super-peak products rely on auction log data and proxies (e.g. high-price ASX contracts), not continuous market observations.
- **Effect on findings**:
    - The estimated **S% range (3.7% median, 9.5% at 95th percentile)** reflects indicative behaviour rather than observed liquidity during active trading.
    - The proposed **tiered spread cap structure** (e.g. 4%, 6–8%, up to 10%) provides **scenario-based** evidence and should be further empirically calibrated from live markets.
    - Policy should allow **flexibility for refinement** once trading data becomes available.

2. Unobserved contract positions and internal hedging

- **Implication**: Market-makers may already be hedging super-peak exposures internally (e.g. via vertical integration), which reduces the need for quoting compensation.
- **Effect on findings**:
    - Net benefit and cost simulations may **overstate required incentives**, particularly for gentailers.
    - The **optimal frontier** is still informative but should be interpreted as a **ceiling of required compensation**, not a minimum.
    - Break-even social benefit thresholds (e.g. $1.5/MWh) may be **conservative estimates** of what's needed.

3. Simplified cost structure (inventory risk only)

- **Implication**: The model assumes spread costs arise from inventory risk based on volatility, midprice, and depth, but excludes:
    - Operational constraints

- Strategic behaviour (e.g. quote shading or cancellation)

- **Effect on findings**:
    - The calculated **marginal net benefit curves** are directionally correct but may **underestimate tactical quoting incentives**.
    - The identified **efficient quote depth range (10–15 MW)** may shift slightly once operational factors are accounted for.

Further refinement of the market-making cost function could draw on sector-specific input price escalation forecasts, such as those provided in Principal Economics Limited (2024)[1]. This would ensure that future adjustments to spread caps remain aligned with evolving cost structures for liquidity provision.

4. No quantification of downstream social welfare effects

- **Implication**: While we estimate private net benefit, broader impacts on competition, hedge accessibility, or consumer pricing are not modelled.
- **Effect on findings**:
    - The **break-even social benefit analysis** (e.g. $1.5/MWh to justify a 2–3% S% cap under high volatility) provides a **threshold**, not a full CBA.
    - Any decision to set tighter caps must be backed by **external evidence** of market or consumer benefit.

5. No endogenous modelling of market behaviour

- **Implication**: The model treats volatility, depth, and spreads as inputs, not jointly determined by strategic interaction among traders.
- **Effect on findings**:
    - The recommended spread caps are structurally sound, but **not equilibrium-tested** in a dynamic setting.
    - This suggests the need for **ongoing monitoring and potential re-calibration**.

These limitations do not undermine the usefulness of the results but indicate that the findings provide a strong initial calibration for spread caps and quoting obligations. They should be implemented with flexibility and sunset provisions, and re-evaluated as live super-peak trading data becomes available. The Authority should consider this project as a first-stage cost-benefit screen, with further development focused on empirical validation and social welfare quantification.

[1] Principal Economics. (2024). *Sector-specific input price escalation forecasts for Electricity Distribution Businesses (EDBs)*. https://principaleconomics.com/reports/sector-specific-input-price-escalation-forecasts-for-electricity-distribution-businesses-edbs/