Claire Preston
Director, Network Expenditure
Australian Energy Regulator
Melbourne VIC 3000

By email: AERInquiry@aer.gov.au

28 February 2025

Jemena Electricity
Networks (Vic) Ltd
ABN 82 064 651 083

Level 16, 567 Collins Street
Melbourne, VIC 3000
PO Box 16182
Melbourne, VIC 3000
T +61 3 9173 7000
F +61 3 9173 7516
www.jemena.com.au

Dear Claire

Submission on AER's consultation on potential refinements to opex cost functions

Jemena Electricity Networks (Vic) Ltd (**JEN**) welcomes the opportunity to provide feedback on the Australian Energy Regulator's (**AER**) consultation of Quantonomics' memorandum regarding potential refinements to opex cost functions for Distribution Network Service Providers (**DNSP**). We appreciate the AER's willingness to explore modelling refinements in response to increasing monotonicity violations and non-convergence issues in Translog models.

Given the holistic nature of this consultation and the impact it can have on base year opex assessment, it is important that the AER gives itself and stakeholders sufficient time to consult to ensure that any refinements to its long-established approach do statistically improve the performance of its models consistently and over time. It is also important that any refinements avoid unexplained shifts in DNSPs' relative benchmarking performance or material changes in efficiency outcomes.

In this consultation, the AER is exploring an alternative approach that introduces jurisdiction-specific time trends, which it believes could better capture the systematic differences in opex between jurisdictions.[1]

To test this approach, the AER and Quantonomics examined two model specifications:

- Jurisdictional Time Trend (**JTT**) – which allows the time trend to vary across all three jurisdictions.
- Australian Time Trend (**ATT**) – which maintains a common trend for New Zealand and Ontario but applies a separate time trend for Australia.

[1] AER, *2024 Annual Benchmarking Report - Electricity distribution network service providers, November 2024*, Pg. 64

While these approaches reduce monotonicity issues in some models, they have some unintended consequences that we share below:

- While JTT and ATT improve monotonicity outcomes in Translog models, they also introduce monotonicity violations in the historically stable SFA CD model.
- The improvements observed in 2024 results do not hold consistently across previous years. Applying JTT to past benchmarking reports (e.g., 2021 or 2022) would have produced similar or even higher levels of monotonicity violations.
- The Australian-specific time trend also captures 'catch-up' efficiencies apart from the technological change. This could result in an underestimation of the efficient level of opex required in revenue determinations, particularly when used in the current AER's roll-forward model. We discuss this in more detail in *Annexure A*.

The AER, in its 2024 benchmarking report, is also considering narrowing the definition of monotonicity violations to only models where the negative relationship between opex and output is statistically significant. Narrowing the definition of monotonicity violations may overlook important issues. It could mask the monotonicity issue rather than address it, as it allows more unreliable estimates to be used in the efficiency assessment.

We have provided more details on our analysis of the 'catch-up' efficiency and the monotonicity definition issues in *Annexure A.* In summary, we suggest that the AER carefully assess whether the overall statistical performance of JTT and ATT justifies their adoption and the unintended consequences of changing the monotonicity definition. Given the complexity of these refinements and their potential impact on DNSP efficiency assessments, we suggest allowing additional time for thorough evaluation and consultation before implementing significant changes.

We look forward to further engagement with the AER on these matters. Please contact [redacted] if you would like to discuss this submission further.

Yours sincerely

[Signed]

Sandeep Kumar
Group Manager Regulatory Analysis, Pricing and Strategy

Annexure A

Issues of 'catch-up' efficiency in Australia-specific time trend

The observed differences in time trends between Australian and overseas DNSPs appear to be largely driven by catch-up efficiencies rather than fundamental jurisdictional differences. This distinction is important because incorporating catch-up efficiencies into a jurisdiction-specific time trend could misrepresent the true drivers of opex reductions.

'Catch-up' efficiency drives the different time trends between Australian and overseas DNSPs

In Figure 1 below, we compare the opex per unit of outputs[2] over time across two groups of DNSPs. The first group contains DNSPs considered to be more efficient by the AER[3], which represents the firms near the efficiency frontier (blue line). The second group contains all other firms, which are further away from the efficiency frontier and more likely to have some 'catch-up efficiencies' embedded over time (orange line). It shows that the *less efficient firms have had much faster reductions in opex per unit of output* (especially since 2012) on average compared to the more efficient firms.

Figure 1: Time series of opex per unit of output

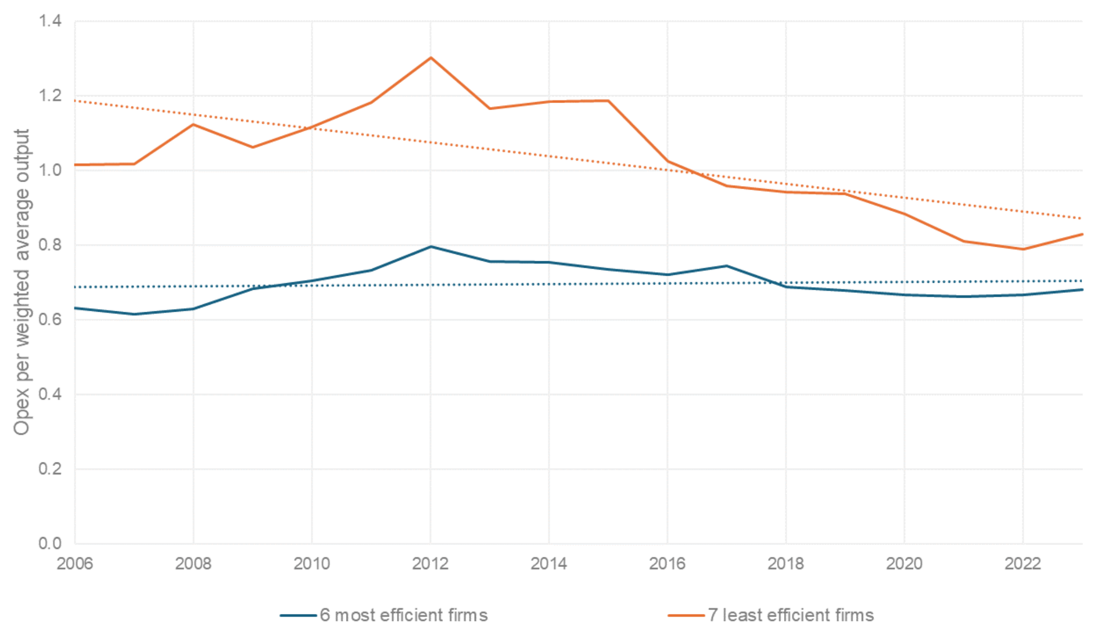

2 Using an estimated weighted average output for each DNSP based on the average of the seven output elasticities (excluding the non-convergent short run SFATLG model) from the 2024 benchmarking report.

3 The six more efficient firms are PCR, SAP, UED, TND CIT and AND, as identified in Figure 14 of the AER's 2024 benchmarking report.

When assessed on an individual DNSP level, we observed a similar pattern, where the DNSPs that were more inefficient in 2006 have generally improved their efficiency more significantly over time from 2006 to 2023.

This suggests that the difference in Australia-specific time trend may be caused by DNSPs undergoing catch-up efficiency improvements rather than experiencing a different underlying technological shift or jurisdictional differences.

Australia-specific time trend risks overestimating the true 'frontier shift'

Adopting an Australian-specific time trend risks misattributing catch-up efficiency as technological progress, which could materially overstate the rate of change at the efficiency frontier. This distinction is critical because overstating the 'frontier shift' has direct implications for how efficient opex is estimated in revenue determinations.

To illustrate the nature of the problem, consider a stylised scenario where all DNSPs have the same and constant outputs and operating environments (the same OEFs). The true technological change is zero. There are two broad groups of DNSPs:

- half of the DNSPs have **always operated efficiently** and remain on the frontier throughout the estimation period ('always efficient DNSPs').
- the other half of DNSPs **were inefficient at the start**, but improve their performance over time and reach the frontier in year 10 ('DNSPs with catch up efficiency').

We apply a regression with a common time trend across these DNSPs and plot them below in Figure 4.

Figure 4: Stylised example of opex cost functions for frontier and catch-up DNSPs

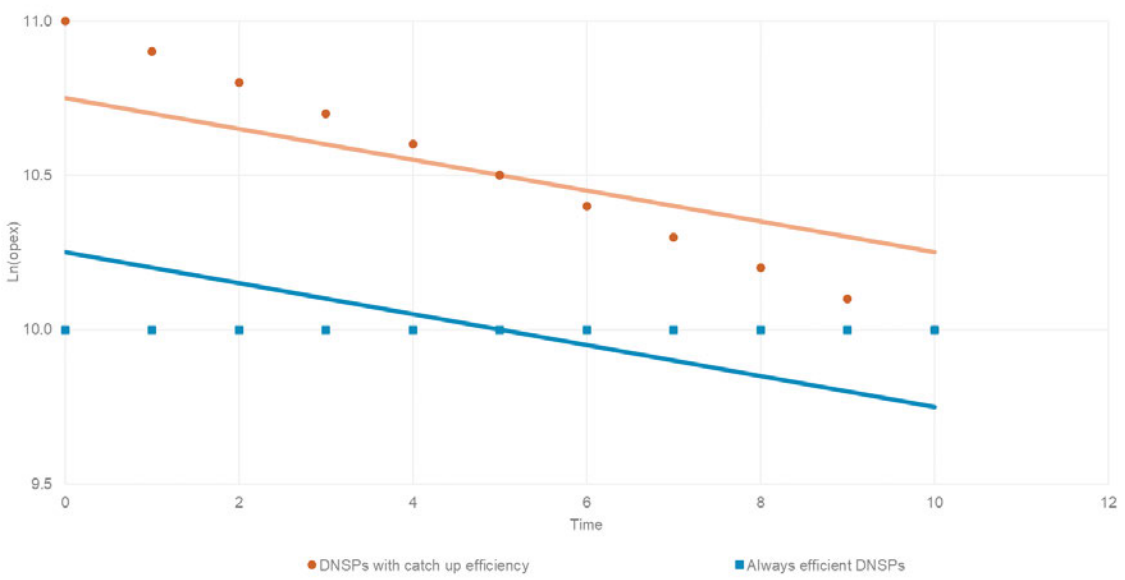

This illustration shows:

- the *always efficient DNSPs* (in blue dots) have a constant *ln(opex)* of 10.

- the *DNSPs with catch-up efficiency* (in orange dots) are 1.71 times less efficient at the beginning of the period (i.e., 2.71 times higher opex) with *ln(opex)* of 11.
- the *DNSPs with catch-up efficiency* improve their efficiencies by 10% each year, such that, after 10 years, they have the same opex (the same efficiency) as the 'always efficient DNSPs'.

While the true 'frontier shift' is zero in this example, the regression assumes the same time trend for all DNSPs and estimates a time trend of 5%, the average of:

- 0% for the always-efficient DNSPs
- 10% for the catching-up DNSPs

This results in overstating productivity growth for the efficient DNSPs, implying they should achieve further opex reductions when they are already at the frontier.

This model misspecification could have significant implications in the way the AER applies the regression results to roll forward opex to the base year for revenue determinations, which we explain further below.

Embedding 'catch-up' efficiencies in time trend risks underestimating base year opex in the AER's roll-forward approach

The AER's current opex assessment approach takes the regression outputs and uses them in its opex benchmarking roll-forward model to determine efficient base year opex in revenue determinations.

With the JTT and ATT models, where the Australia-specific time trend captures catch-up efficiencies, the AER's roll-forward model risks double-counting efficiency improvements and significantly underestimating the efficient base year opex for DNSPs.

The AER's opex roll-forward approach consists of two key steps:

- **Adjusting mid-period opex to an efficient level** – For DNSPs with an efficiency score below 75%, the AER's approach first reduces their actual average opex to a benchmark-efficient level.
- **Applying a rate of change, including time trend, to roll forward efficient opex** – The model then projects efficient opex from the mid-period to the base year using the estimated time trend.

If the time trend includes catch-up efficiency, the roll-forward model effectively applies the same required efficiency improvement twice:

- Once when adjusting mid-period opex down to the efficiency frontier.
- Again when rolling forward opex using a time trend that embeds further efficiency improvements.

This can lead to an artificially low estimate of efficient base-year opex, particularly for DNSPs that have already made substantial efficiency improvements.

We illustrate this further using the same stylised example in Figure 4, with an additional step to roll forward opex – as illustrated in Figure 5.

Figure 5: Stylised example showing AER's efficiency adjustment and opex roll-forward

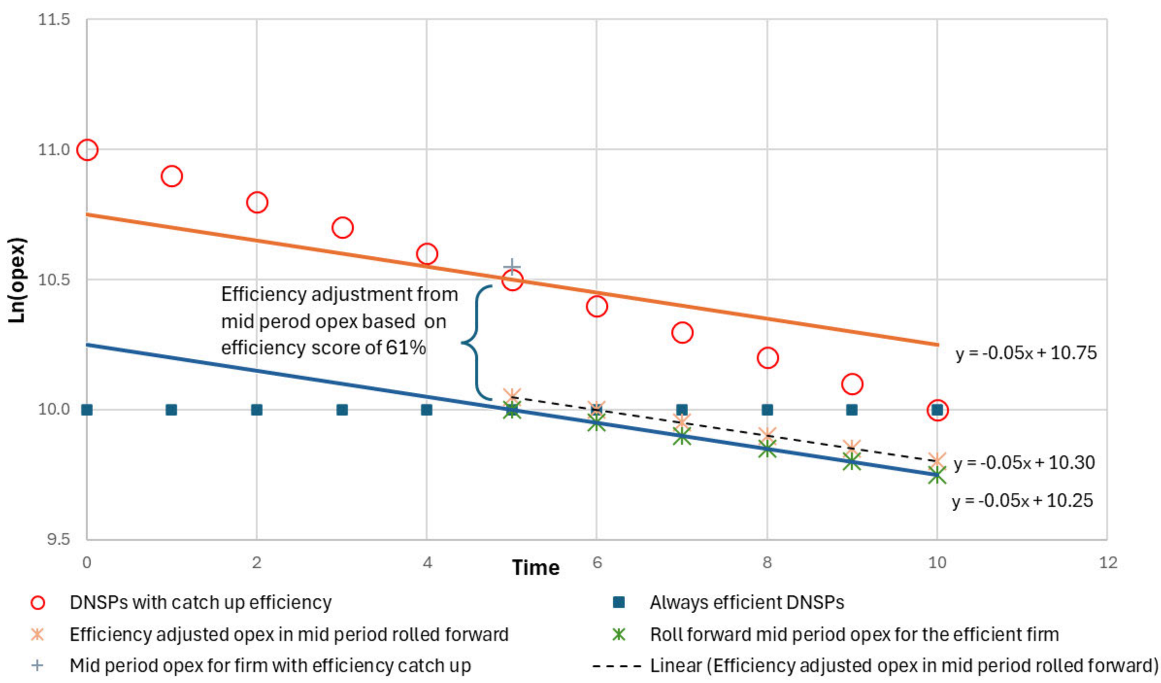

Figure 5 applies the AER's opex roll forward method to the *DNSPs with catch-up efficiency* in the following steps:

1. Determine the efficiency score: the model calculates an efficiency score of 60.7%, calculated by the difference in intercepts in the two regression lines shown in Figure 5 above (i.e. 60.7% = *exp(10.25)/exp(10.75)*).
2. Calculate the average actual opex over 10 years: in this example, this equals year 5 opex (*ln(opex)=10.5*)
3. Adjust mid-period opex: the mid-period actual opex is reduced to align with the efficiency frontier (shown as the yellow dot in year 5)
4. trend the mid-period opex to base year (year 10) (shown as the yellow dot in year 10)

In this example, the final rolled-forward opex is systematically underestimated by 18%:

- The true efficient opex in year 10 should be *ln(opex) = 10*
- However, the AER's roll-forward model estimates *ln(opex) = 9.8*, underestimating efficient opex by 18% (i.e., *1 - exp(9.8)/exp(10)*).

This means DNSPs with catch-up efficiency are being penalised twice:

1. Their opex is already adjusted downwards to the efficiency frontier (from red to yellow dot in year 5).
2. The roll-forward model further reduces opex by applying a time trend that embeds additional efficiency improvements. (from yellow dot in year 5 to year 10).

This underestimation has significant consequences if applied to revenue determinations, as efficient DNSPs may receive allowances below sustainable levels, affecting their ability to maintain service quality. Also, benchmarking outcomes could unfairly force DNSPs that have already improved efficiency to continue reducing opex at an impractical rate. We recommend the AER explore ways to separate catch-up efficiency from frontier shifts before changing its current model specification.

Considerations for modelling catch-up efficiencies

The improvements in monotonicity violations observed with an Australian-specific time trend could also arise from model mis-specification that captures catch-up efficiency rather than an actual improvement in reliability. While addressing the treatment of catch-up efficiencies in benchmarking models is an important area for further development, it presents significant complexities that must be carefully considered.

Unlike technological change, which affects all DNSPs similarly, catch-up efficiency varies significantly across DNSPs, over time, and in magnitude. To model the catch-up efficiencies for Australian DNSPs, the key challenges include:

- Variability in efficiency improvements – Some DNSPs experience substantial opex reductions due to internal efficiency improvements, while others show more gradual changes. These trends are not consistent across all DNSPs or years.
- Changing relative efficiencies – The relative efficiency of DNSPs shifts over time, meaning that some firms improve faster than others at different points in time, rather than having static efficiency position over the years.
- Potential Model Instability – Highly flexible models designed to capture catch-up efficiencies risk introducing new statistical challenges, such as data convergence issues, overfitting or further monotonicity violations.

Given these complexities, modelling catch-up efficiency effectively may require changes beyond simply adjusting the time trend. It is important to ensure that –

- the underlying assumptions in the time-varying inefficiency term in the model specifications align with the actual observed Australian DNSP efficiency trends.
- the improvement in statistical properties and reliability is consistent across all four econometric models.
- if a new approach is introduced to capture time-varying efficiencies, the existing roll-forward approach would also require further consideration and changes to avoid unintended distortions in revenue determinations.

Issue of monotonicity violation definition

Monotonicity is a key economic property required for the AER's econometric models. It requires that an increase in output is only achieved with an increase in inputs (i.e., opex), holding all else constant. In its 2024 benchmarking report, the AER is considering narrowing the definition of monotonicity violations to only models where the negative relationship between opex and output is statistically significant.

Ensuring the definition of monotonicity violations supports more reliable benchmarking outcomes

Under the new definition of a 'statistically significant' monotonicity violation, it would require both:

- that the point a DNSPs located on the cost function to have a negative slope (negative change in opex with increasing outputs); and
- that this slope was statistically significantly different to zero.

This would involve a material narrowing of (reduction in) the circumstances that would be treated as a monotonicity violation.

While we recognise the AER's intent to refine the treatment of monotonicity violations, we believe the definition of monotonicity violation should be broadened instead of narrowed to produce more reliable benchmarking outcomes.

A monotonicity violation (negative relationship between opex and output) anywhere on a cost curve suggests the overall model is unreliable. At a minimum, all points near a monotonicity violation are unreliable – even if the slope of the curve at those points is positive in relation to output. Limiting exclusions to only statistically significant cases risks retaining models that generate implausible relationships between opex and output.

We provide an example on a fairly common translog model shape below to illustrate the issue further, as shown in Figure 6 below.

Figure 6: Illustration of a typical translog cost curve

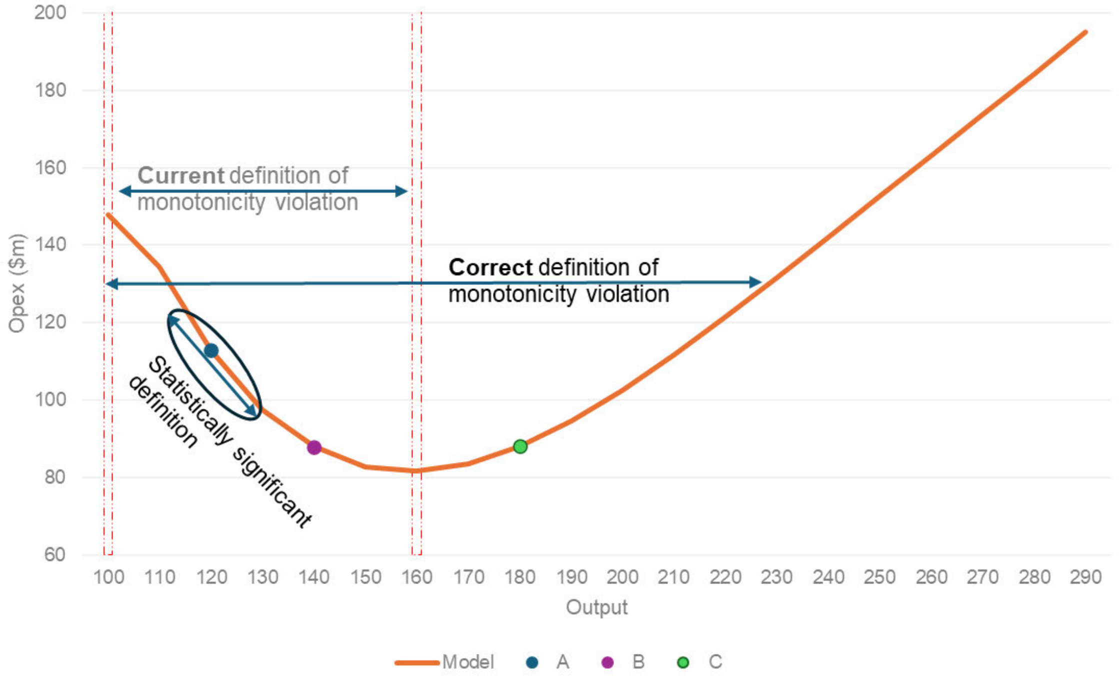

In this model, opex costs are initially falling with one output (e.g., fall with RMD) before increasing again. Under the AER's current test:

- DNSPs positioned on the left of the minimum point, where opex declines as output rises, are identified as having monotonicity violations (points A and B).
- DNSPs positioned on the right of the minimum point, where opex rises again, are treated as not violating monotonicity (point C), even though point C has *higher outputs but lower opex* compared to point A. Point C technically violates monotonicity but is not treated as a violation under the AER's current test.

This approach implicitly assumes that the minimum point is valid when, in reality, the presence of any turning point undermines the economic credibility of the entire cost curve. The minimum point represents the culmination of all of the prior unrealistic negative slopes up to that point. It is, therefore, resulting in an unreliably low estimate of the efficient opex in that entire 'valley' region.

If the AER further narrows the definition by only excluding 'statistically significant' violations, the issue becomes more pronounced. In Figure 6, when adopting the narrower definition of 'statistically significant' violations:

- only point A is treated as a violation, but
- points B and C are treated as satisfying monotonicity, even though both points have higher output but lower opex than point A, and should technically be treated as violations.

If this narrower definition is adopted, models with unrealistic cost curves could still be used for benchmarking simply because their violations are not statistically significant. This risks undermining the economic validity of the results and runs against the intention of making the models more reliable and robust.

We recommend adopting a broader definition of monotonicity violation

Instead of narrowing the definition, we suggest a broader definition of monotonicity violations to ensure that all retained models generate reliable cost curves (refer to '*correct definition of monotonicity violation*' in Figure 6 above).

If the estimated cost curve has any negative slope, the entire curve should not be used. At a minimum, the curve should not be used to test efficiency for DNSPs located in both:

- on a negative slope (i.e., on the left (right) hand side of a minimum (maximum) turning point); or
- Close to the trough, if a ± 20% change in any one output would shift that DNSP to a negative slope.

An alternative approach, which avoids the need to assess individual cost curves for each DNSP and output, could be to exclude translog models for DNSPs with a history of persistent monotonicity violations or those more prone to such issues.

For instance, if a DNSP has had a translog model excluded in more than 50% of assessments over the past five years, it may be more appropriate to remove that model from its benchmarking assessment. Under this criterion, as shown in Table 1, Ausgrid, JEN, and United Energy would be excluded from both Translog models, while CIT would be excluded from one model.

Table 1: Monotonicity violations over the last five years

	Violations over last five years		**Exclude if >50% of models are violated**	
	SFATLG	LSETLG	SFATLG	LSETLG
EVO	30%	0%		
AGD	60%	60%	Exclude	Exclude
CIT	50%	80%		Exclude
END	50%	20%		
ENX	50%	50%		
ERG	50%	0%		
ESS	30%	0%		
JEN	70%	70%	Exclude	Exclude
PCR	50%	0%		
SAP	40%	0%		
AND	50%	20%		
TND	20%	0%		
UED	70%	80%	Exclude	Exclude