MONASH University

MONASH ENERGY INSTITUTE

Submission to the Artificial Intelligence and Data Centres Inquiry and Report consultation

HOW TO GROW THE DIGITAL ECONOMY WITHOUT SENDING THE BILL TO HOUSEHOLDS

SEPTEMBER 2026

Prepared by:
Dr Changlong Wang
Dr Hao Wang
A/Prof Roger Dargaville
Dr Graham Palmer
Dr Mehdi Ghazavi Dozein
A/Prof Stuart Walsh
Dr James Wang

Monash University

PREFACE

The Monash Energy Institute welcomes the opportunity to contribute to the Senate inquiry into artificial intelligence and data centres.

The Institute brings together researchers from across Monash University to address the social, technical, and economic dimensions of the energy transition. We work with government agencies, industry and communities to produce independent research-based evidence.

AI and data centres present potential opportunities for Australia. They will also become a major source of electricity demand, with implications for energy infrastructure, system reliability, consumer costs, water, the environment and local communities. Decisions being made now about where and how data centres are built and connected will shape these impacts for decades.

This submission asks how Australia can support data-centre growth without shifting disproportionate costs and risks onto households and other electricity users. Our authors' modelling indicates that inflexible growth could substantially increase system costs and prices, while modest flexibility in the timing of some computing workloads could deliver significant savings.

Important knowledge gaps to inform decisions remain, including how demand and flexibility will evolve and perform in practice; how power purchase agreements, firm supply, batteries and backup generation could be treated; and how connection standards, tariffs and incentives can be designed. Equally important are complex social and environmental questions that shed light on impacts on households, communities and nature.

This is a relatively new and rapidly changing field. Some future challenges can already be anticipated. Others will emerge as AI technologies, business models and energy systems develop. Australia will therefore need sustained, interdisciplinary research and innovation to identify emerging risks and opportunities, test policy and market options, and provide independent evidence as circumstances change.

The Monash Energy Institute is committed to helping build this evidence base and translating it to inform sound policy and practice. We welcome the opportunity to work with governments, industry and communities on this important national challenge.

CONTENTS

ACKNOWLEDGEMENTS

The authors thank Professor Yolande Strengers, Associate Professor Julie Karel, Annelieke de Wit, Mara Jorgovic, Shreejan Pandey, and colleagues at the Monash Energy Institute for their constructive comments, and discussions. Their input helped improve this submission.

SUMMARY

Australia's data centres are about to become one of the largest new sources of electricity demand in the country. On AEMO's central forecast, their consumption grows from about 4.7 TWh in 2025-26 to almost 34 TWh by 2050, approaching one tenth of all underlying electricity demand in the national grid.

That growth can be managed well or poorly, with the difference being potentially worth billions of dollars each year. Monash Energy Institute's modelling of the national electricity market, built to align with AEMO's 2026 Integrated System Plan, finds:

1. If data centres grow as forecast and run *inflexibly*, the power system incurs an extra **A$18.5 billion**[1] in cost. That is a 9.1% increase in total system cost from today to 2050.
2. Ordinary consumers [2]may pay most of it. By 2050, households and businesses that have nothing to do with data centres will pay an extra **A$6.1 billion** a year in wholesale electricity costs, about **58%** of the total increase. The cost burden is heaviest exactly when people can least afford it: in winter, and in the periods of greatest system stress. These figures describe data centres supplied from the shared grid, as AEMO's forecast assumes; a facility that arrives with its own firm generation adds less to the burden, but only to the extent its supply is still there in those tightest hours.
3. A modest amount of flexibility changes the picture. If 20% of IT computing work is delay-tolerant and can be rescheduled within a 24-hour completion window, the system saves **A$3.5 billion** on the same present-value basis, and roughly two thirds of the resulting bill relief lands with consumers. At 30% the saving rises to **A$4.7 billion**. And the leeway needed is short: allowing even six hours of delay already delivers more than half of what a full week of leeway would, so the savings come from broad participation on moderate terms, not from demanding extreme deferral.
4. The requirements placed on operators are concentrated and manageable. In the most stressed hours of the year, flexible facilities deliver 96% of their promised demand reduction, and every deferred job still finishes on time.
5. Timing matters. Flexibility commitments made now, while facilities are being designed, are worth about 28% more than commitments phased in gradually as the fleet builds out.

Flexibility reduces the need for additional renewable generation, energy storage and transmission infrastructure, lowering the overall system investment required to support data centre growth. Our recommendations, in brief:

- make flexibility a condition of connection for new, large, data centre loads;
- act sooner, rather than deferring action until demand pressures become more acute;
- pair any clean-energy matching obligation with recognition of flexibility;
- make what is inside data centres visible to connection standards;
- and where system planning relies on committed flexibility, establish clear incentives for delivery and consequences for non-performance.

[1] Expressed in present value, i.e., the sum of all extra capital and operating costs from now to 2050, discounted back to today and expressed in 2026 dollars.

[2] "Ordinary consumers" here means all electricity users other than the data centres themselves; using digital services does not put a household on the data-centre side of this ledger. The split is the wholesale-market outcome under current arrangements, before any policy that reallocates costs between users.

1. The scale of future data centre demand

Data centres are buildings containing clusters of computers that run cloud services, web services, artificial intelligence (AI) training, and inferencing. Australia already has more than 160 of them in operation, concentrated in Sydney and Melbourne. AEMO's March 2026 connection pipeline includes eleven transmission-connected projects totalling more than 5 GW.

AEMO's central (Step Change) forecast has data centre electricity use rising from 4.7 TWh in 2025-26 to 11.8 TWh by 2029-30, and 33.8 TWh by 2049-50. By mid-century, data centres are expected to reach one tenth of all underlying electricity demand in the national grid.

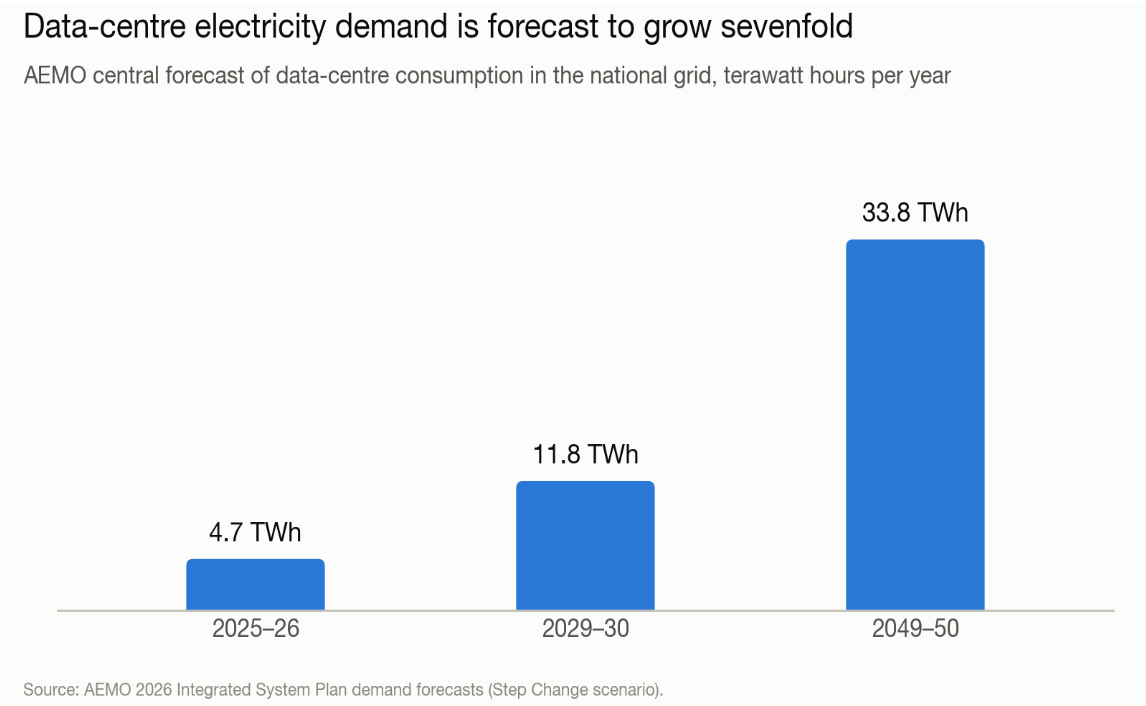

Figure 1. AEMO's central forecast. The 2049-50 figure approaches one tenth of underlying electricity demand in the national grid.

Two features of the official forecast matter for policy:

- AEMO treats data centre demand as **non-flexible**. Its 2025 reliability outlook showed that developers prioritised uptime, so the bespoke data-centre forecast did not include demand flexibility. The AEMO Integrated System Plan (ISP) does carry a separate, industry-unassigned allowance for general demand-side participation, which could in part be supplied by data centres; that allowance is unchanged across all our scenarios.
- AEMO does not count the backup equipment inside data centres (batteries, uninterruptible power supplies, standby generators) as resources available to the grid, and it publishes no hour-by-hour data-centre demand trace that would let outsiders test alternatives. Not counting on-site backup equipment is common practice in other jurisdictions as well. Section 6 returns to both gaps: what it would take for planners to count flexibility they can rely on, and whether the equipment already inside facilities should be visible to the planning framework; work quantifying both is under way at Monash.

In other words, the national planning framework currently plans *around* data centres rather than *with* them. This submission asks whether that is the right default, and what it costs.

2. What inflexible growth costs

We model the national electricity market with Monash University's own optimisation model (MUREIL-ISP), aligned with the 2026 Integrated System Plan: same regions, same candidate technologies, same carbon constraints, and hourly operation. We compare a world where data-centre demand grows on AEMO's forecast and runs inflexibly (the AEMO-aligned non-flexible case) against a world where data-centre demand stays at its 2026 level.

Serving the forecast growth inflexibly adds **A$18.5 billion** to the cost of power system operation and expansion. Throughout this submission, system costs and savings of this kind are stated on one consistent basis: the present value of the total cost difference from today to 2050. That is, we add up every extra dollar of capital and operating cost the growth causes in each year from 2026 to 2050, discount those dollars back to today, and express the total in real 2026 dollars. On that basis the no-growth world costs A$203 billion in total, so inflexible data-centre growth adds 9.1% to the cost of the entire system, attributable to a single demand category.

The costs counted are generation, storage and grid-security plant (capital and operating) plus the nine major transmission corridors AEMO's plan treats as candidate upgrades, which the model builds as needed. Grid-security plant is the equipment that keeps the power system stable as inverter-connected supply grows; the ISP's planning rules require it to accompany new wind, solar and batteries, and its cost is included in our totals. The costs not counted are local connection assets, distribution networks and retail costs. The direction of the excluded network effects is genuinely open, and we do not claim it. Steady, high-utilisation loads can improve network utilisation and dilute fixed costs across more energy, as transmission-connected industrial loads have historically done; additions to coincident peak demand instead drive augmentation whose costs are recovered from all users, and network charges are roughly half of a typical household bill. Which effect dominates for data centres depends on where they connect and how they operate in peak periods, and flexibility in the tight hours pushes toward the benign outcome. Quantifying this is part of our ongoing work.

Where does the money go? Mostly into extra capacity:

- 7.9 GW more wind and 9.9 GW more solar;
- 3.9 GW more grid batteries, plus the accompanying grid-enhancing technologies, such as synchronous condensers.

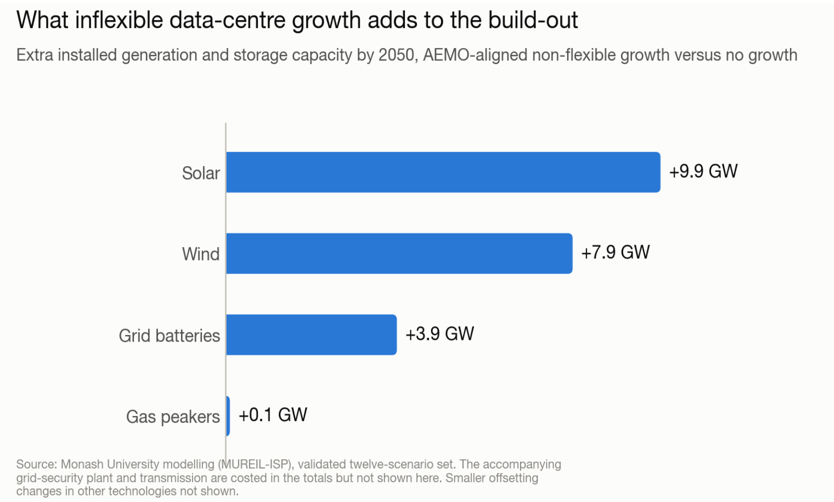

Figure 2. The comparison is the AEMO-aligned non-flexible growth case against a world where data-centre electricity stays at its 2026 level, measured at 2050.

3. WHO PAYS

Electricity costs reach people through prices. Because ordinary consumer demand is identical across all our scenarios, we can read the bill changes as pure price effects, with no change in how much anyone uses. Two notes on units. First, unlike the present-value totals in section 2, the figures in this section are annual amounts for the year stated, undiscounted, so "A$6.1 billion a year in 2050" means exactly that. Second, every bill and price in this submission is a wholesale spot-market figure: each hour's consumption valued at the wholesale price for that hour. Network charges, retail margins and other levies, which make up much of a final retail bill, sit on top of these amounts and are not modelled. Who ultimately bears these costs is a policy choice as much as a market outcome; what the modelling pins down is the size of the cost pool and its default incidence under current arrangements.

By 2050, inflexible data-centre growth raises the total annual electricity bill by A$10.4 billion. Of that:

- 58% is paid by ordinary consumers**:** A$6.1 billion a year, or an extra $31.2/MWh. If passed through in full, that is about $155 a year in wholesale costs for a household using about 5 MWh a year, before retail margins.
- Data-centre operators carry most of the remainder, about A$4.4 billion a year. Lack of flexibility in data centre operation is expensive for the industry too, because a larger, always-on fleet bids up the price of the very hours it depends on. (The small remainder nets against a modelled hydrogen-production load that is marginally better off, which closes the arithmetic.) The hydrogen load is part of the AEMO Step Change demand projections, not an addition of ours.

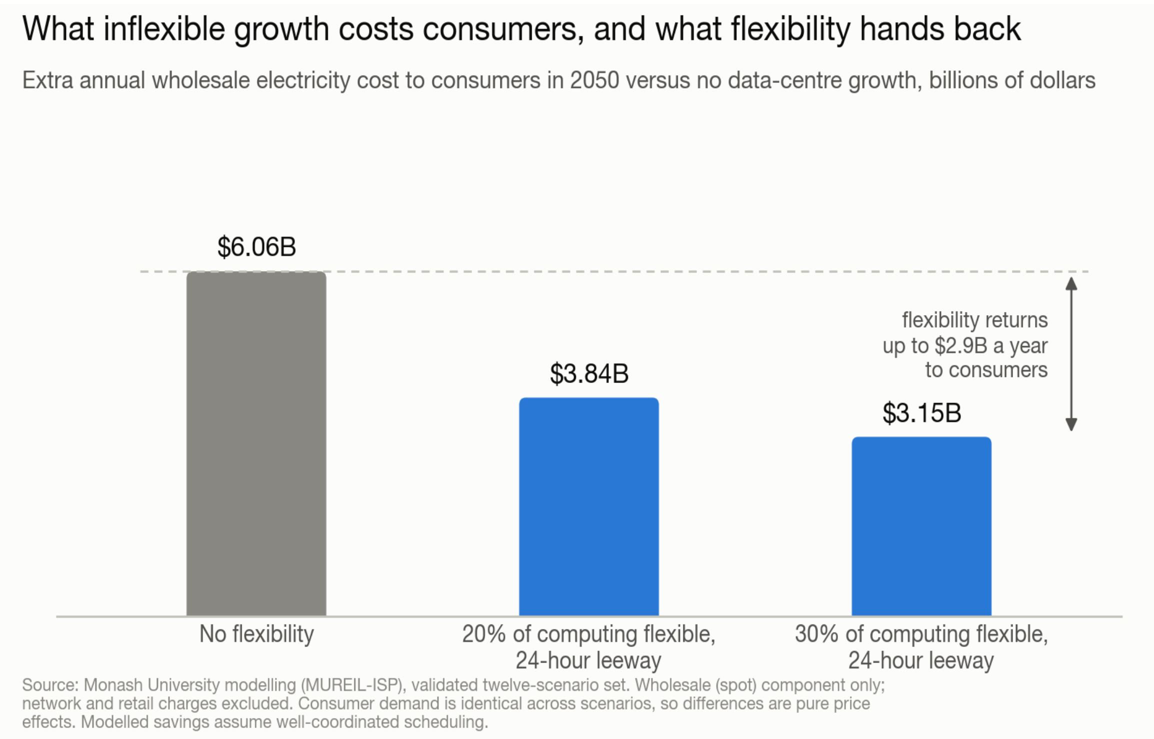

Figure 3. Consumer demand is identical in every scenario, so the differences shown are pure price effects, not changes in how much anyone uses.

The consumer burden is not spread evenly throughout the year. The winter price increase, at $102.7/MWh, is more than three times the annual average increase. In the top few percent of stressed hours, price rises by around $341/MWh. Inflexible data-centre growth costs consumers most in exactly the hours when the system, and household budgets, are already under the most pressure.

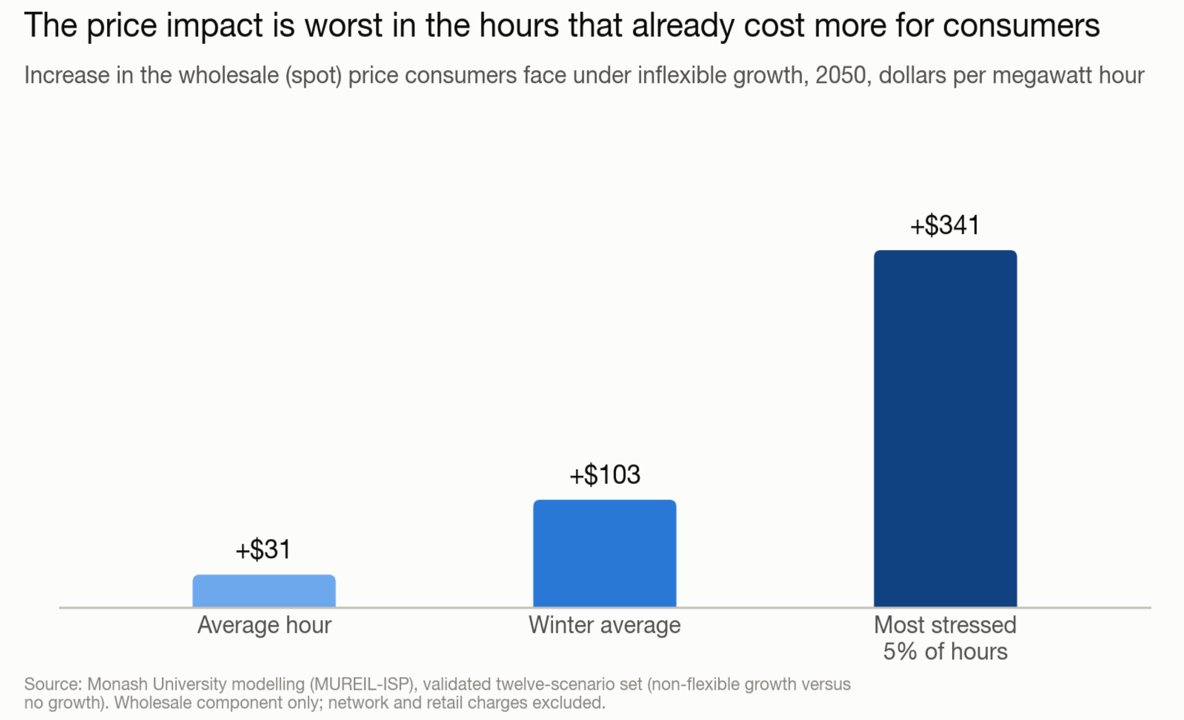

Figure 4. The same inflexible-growth comparison as Figure 2, expressed in the wholesale price consumers face at different times.

This escalates over time. The consumer cost of inflexibility rises nearly sixfold between 2026 and 2050, tracking the growth of the fleet. Decisions taken in the next few years about how the coming wave connects will set the trajectory.

One condition applies to all of these figures. Following AEMO's forecast, we model data-centre demand as supplied from the shared grid. A facility that instead arrives with its own dedicated generation adds less net demand, and the burden on consumers shrinks with it. How much it shrinks depends on when that self-supply is available, not just how much energy it produces. Because the consumer burden is concentrated in winter and in scarce hours, self-supply that reduces in exactly those hours (solar without associated storage, for example) leaves most of the burden in place, while firm self-supply, such as on-site generation paired with storage or contracted dependable capacity, genuinely relieves it. We have not modelled self-supplying data centres explicitly; the direction of these effects follows from where the costs arise. It is one more reason connection frameworks should treat what a facility brings with it as part of the deal (Recommendations 1 and 5).

A closely related question, which we raise for the inquiry rather than answer, is how power purchase agreements (PPAs) should be treated. Data centres increasingly contract with new wind and solar projects, and those contracts genuinely help finance new supply. However, a PPA is a financial arrangement: by itself, it does not align a facility's electricity consumption with the output of the contracted generator or remove its reliance on the grid when that generator is unavailable. This distinction is particularly important during high-demand winter periods, and when wind or solar output do not coincide with the facility's load. Whether a PPA-backed facility should be treated as relieving other consumers, and what firmness standards that treatment should require, is a question the presented results do not address. It is a priority in our ongoing work, and we suggest the inquiry treat it as a live design question alongside the connection issues above.

4. The opportunity: a little flexibility is worth billions

Not all computing has to happen the moment it is requested. Some workloads (overnight batch processing, model training, some data management) can be deferred for hours or even days with little or no noticeable impact. We tested what happens if a slice of data-centre computing can be shifted in time, with every job still completed by a deadline and total annual computing unchanged.

The central case gives 20% of computing work up to 24 hours of leeway. The results are summarised below:

- The system saves **A$3.5 billion**, again as a present value of all avoided capital and operating cost from now to 2050, in 2026 dollars. Raising the flexible share to 30% lifts the saving to **A$4.7 billion**.
- Consumers capture roughly 60% to two thirds of the bill relief in every flexible case we ran; the remainder goes to the data-centre operators providing the flexibility. This happens because ordinary consumer demand is about six times data-centre demand, so a price improvement is worth more in aggregate to consumers than to the operators themselves. (Part of the system saving is avoided capital that never enters any bill, which is why the bill relief is smaller than the headline saving.) Consumer demand follows the AEMO Step Change projections, which embed AEMO assumptions on rooftop solar, home batteries and other consumer energy resources, and is identical across our scenarios.
- In our central and deepest cases, the flexible slice of computing buys its electricity for **$24 to $38/MWh**, against roughly $140 if it ran at the moment jobs arrive. Under deep flexibility, a data-centre fleet nearly six times today's size buys its energy more cheaply than today's fleet does.
- The savings are not a transfer from generators. Under deep flexibility, total payments to generators fall because less scarcity pricing occurs; the balance of the saving is avoided capital that was never built at all. In the central case the avoided build at 2050 is concrete: about 1.0 GW of grid batteries, 1.6 GW of wind, 0.4 GW of solar, and the grid-security plant that would have accompanied them.

What the 20% / 24-hour case means. The 20% refers to the share of IT computing work that is assumed to be delay-tolerant. The remaining 80% is treated as fixed in time. The 24 hours is the maximum completion window, not the assumed delay: eligible jobs may run immediately or be moved by only a few hours depending on system conditions. Deferred work must still be completed, so this is workload shifting rather than load shedding. It should not be interpreted as 20% of total facility electricity being continuously interruptible. The flexible share is defined over total IT workload and is technology-agnostic: AI training is the archetypal delay-tolerant workload, but batch analytics, backups and offline inference also qualify, while latency-critical services sit in the fixed 80% and are never rescheduled. These are modelling scenarios rather than measurements: real workloads span a spectrum of delay tolerance, from a few hours to several days, and Table 1 and Figure 7 show how the value scales across that range. Part of the flexible share could also be supplied without touching computing at all, through auxiliary systems such as cooling and on-site storage that shift a facility's electricity draw while applications run unchanged. The share is applied uniformly across the modelled fleet; in practice hyperscale campuses are likely to offer more flexibility than edge or co-location sites, so 20% should be read as a fleet-wide average, with differentiation by facility class part of our ongoing work. Catch-up running is capped at 25% above the reference IT load, a declared scenario assumption, and the operator-side cost of providing any extra computing headroom is not counted: the savings quoted are grid-side only.

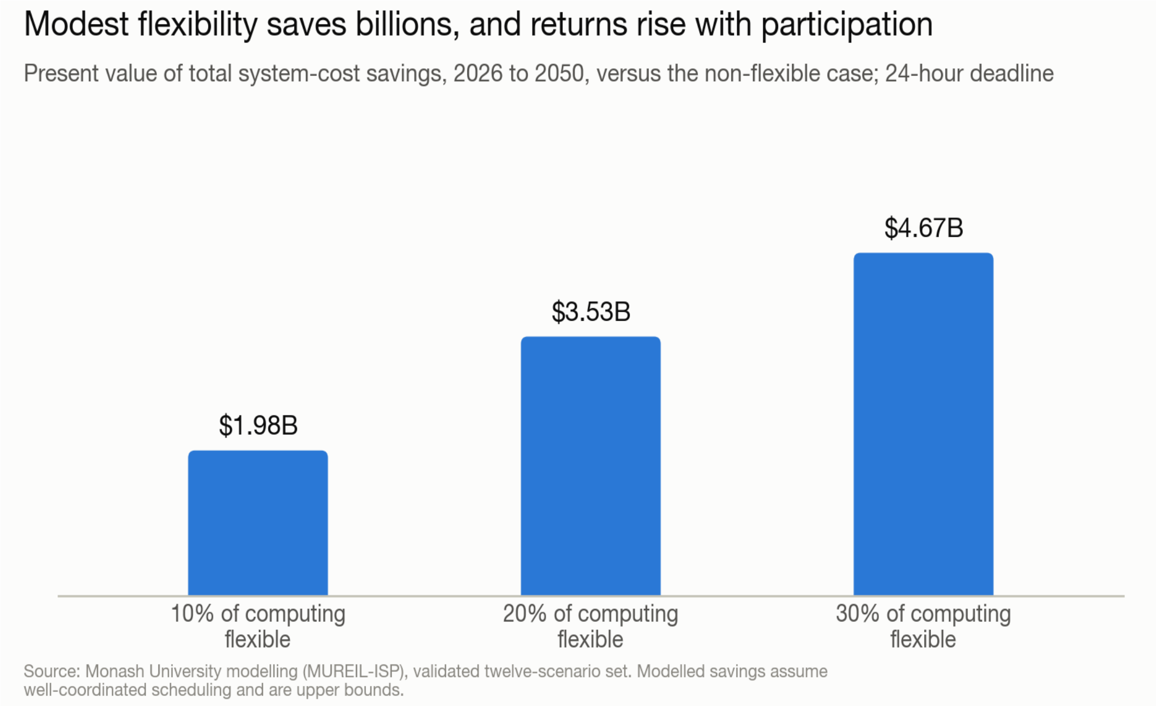

Figure 5. Each additional step of participation adds a similar amount again. The modelled savings assume well-coordinated scheduling and are upper bounds.

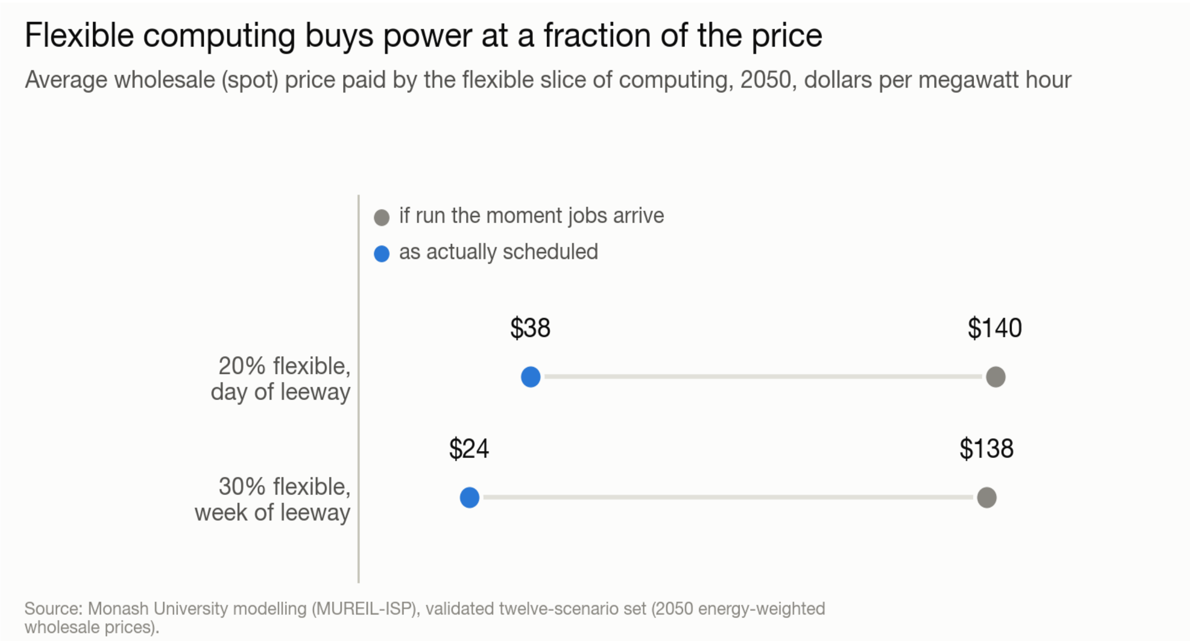

Figure 6. Energy-weighted wholesale (spot) prices for the flexible slice of computing only; the gap between the dots is the raw material of the system savings.

Two design findings matter for policymakers:

The first day of leeway does most of the work. Extending the deadline from six hours to a full week adds progressively less value: the first six hours of allowable delay capture 54% of the value available at a week. By contrast, each additional ten percentage points of flexible share adds roughly constant value. Policy should therefore chase breadth (more participating load) before depth (extreme deferral windows), while recognising the two reinforce each other.

Figure 7 and Table 1 lay this out across all nine solved combinations. Savings rise steadily with the share of computing that participates, at every completion window. They flatten as the window lengthens: at every share, a six-hour window already delivers about half of what a full week would (51% at a 10% share, 54% at 20%, 58% at 30%). The ratio edges up with participation because longer windows are worth more when more load can use them, which is the sense in which breadth and depth reinforce one another. The practical point stands at every share: most of the value of data-centre flexibility is available on terms measured in hours.

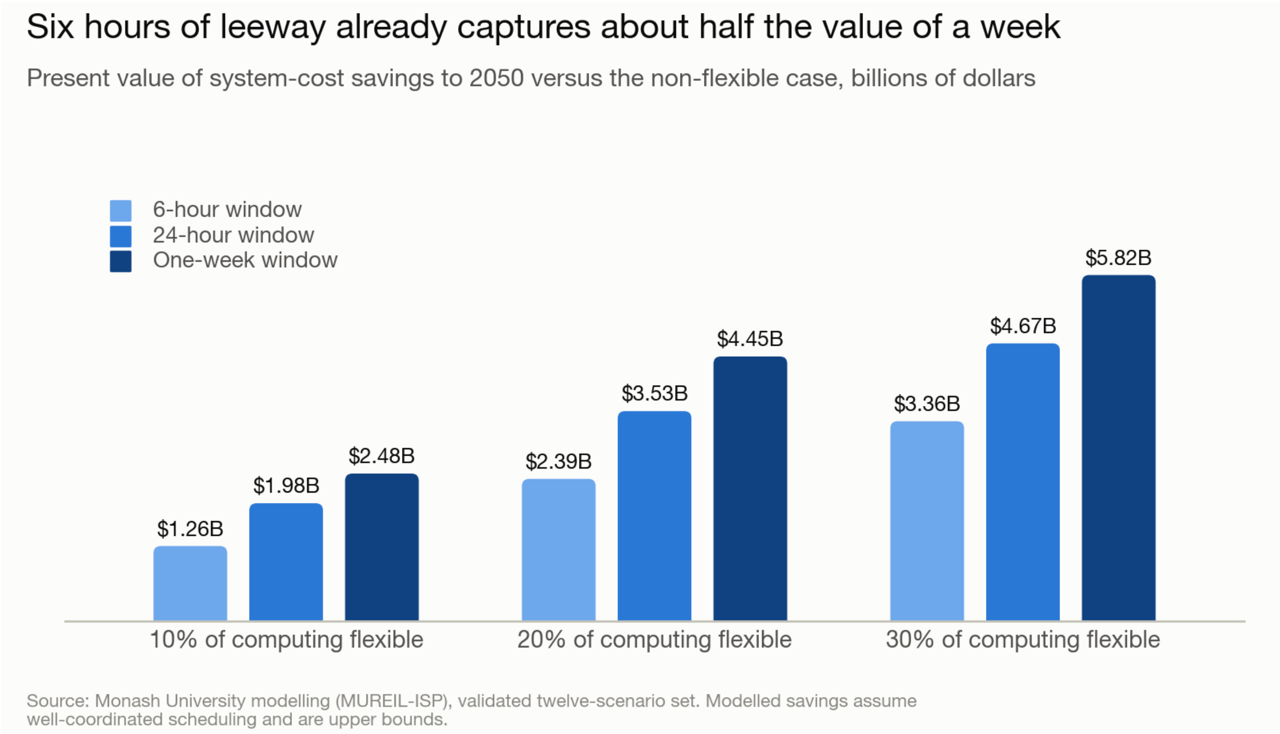

Figure 7. Savings across every solved combination of participation and completion window. Returns rise almost linearly with the participating share, and flatten quickly as the window lengthens: the step from a day to a week adds less than the step from six hours to a day.

Table 1. Savings at six-hour and one-week completion windows, and the six-hour share of the one-week value, for each solved participation level.

Flexible share	Saving, 6-hour window	Saving, one-week window	Six-hour share of week value
10%	A$1.26 billion	A$2.48 billion	51%
20%	A$2.39 billion	A$4.45 billion	54%
30%	A$3.36 billion	A$5.82 billion	58%

Flexibility delivers when it matters. In the most stressed 5% of hours, flexible facilities removed 96% of their nominal flexible demand from the grid. Averaged over the whole year, about 61% of flexible work moves out of its arrival hour, so stress-period delivery is markedly higher than routine-hour rescheduling: the binding constraints act on ordinary weeks, not on the crises. And because the flexible slice is a fifth of computing and two fifths of that slice never moves at all, most of a facility's total workload runs exactly when it otherwise would. For an operator, the commitment amounts to: keep every deadline, and get out of the way in the handful of hours that set everyone's bills.

If flexibility is this valuable, why is it not already routine? Mostly because facilities rarely see the value. Most data centres buy electricity through flat-priced contracts and are held to uptime-first service agreements, so the hourly price signal stops at their gate. Much of the value is avoided system capacity, which no energy price pays them for under current arrangements. And with no framework for a planner to count, verify and reward the capability, offering it earns nothing. Some operators already shift work for cost or carbon reasons, and job-scheduling software is standard in the industry; the missing piece is institutional rather than technical, which is what the recommendations address.

The modelled record shows the mechanism directly. In the most expensive week of 2050 at the largest data-centre node, a winter week in the Sydney region where the wholesale price spikes to $25,000/MWh, flexible computing drops to zero across the price spike, bursts to about 2.5 GW in the cheap hours either side, and clears its backlog of deferred work (peaking near 15 GWh nationally) without missing a single deadline. That week was selected by a pre-registered rule, the highest-priced week at the largest data-centre node, not chosen for effect.

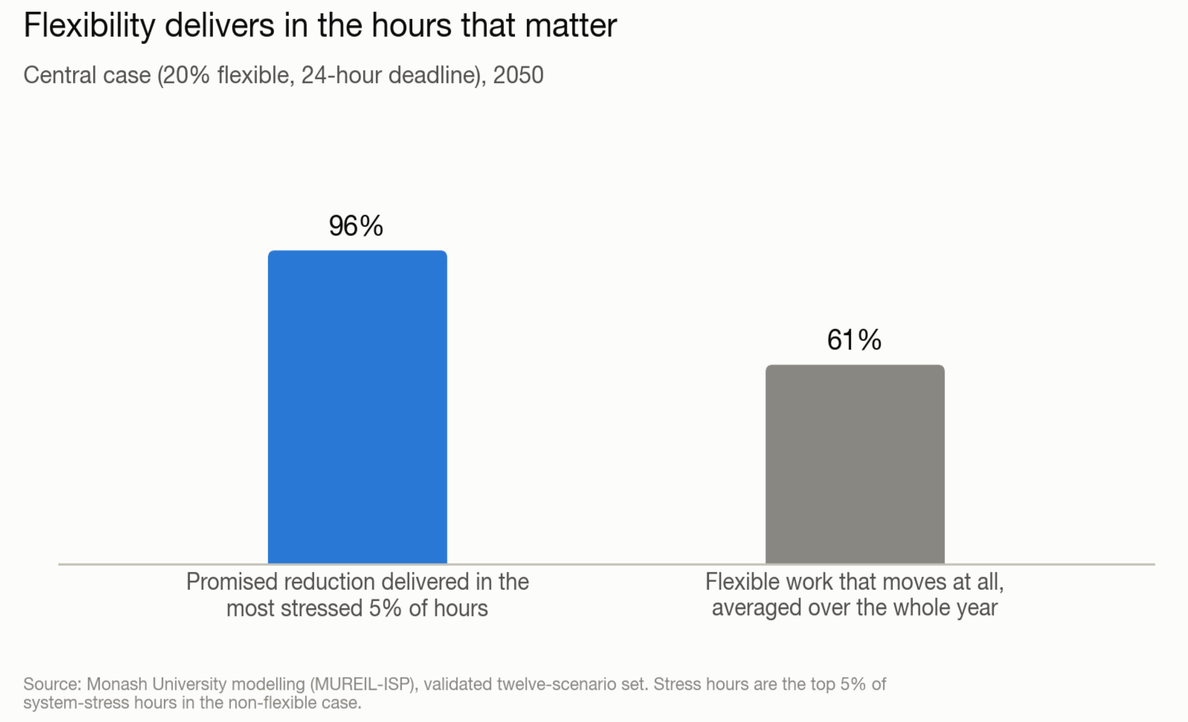

Figure 8. **Delivery is near-complete in the hours that set the system's costs, even though across the year almost 40% of flexible work never needs to move at all.**

We should be clear about one limitation: these savings are computed with the model coordinating flexible workloads perfectly against system conditions. They are upper bounds on what real markets and contracts will capture. Real institutions will deliver a share of them; the size of that share is a design question for this consultation, not a modelling question.

Flexibility is also not a free lunch in every hour. It lowers the extreme peaks substantially, and it raises prices in previously quiet hours as shifted computing fills them; in the deepest flexibility case, 36% of the year's hours end up more expensive than they would have been without flexibility. The net annual effect for consumers is still clearly positive, but the honest description is that flexibility flattens prices rather than lowering them everywhere. The interaction with storage deserves particular attention in further work: flexible computing competes with grid batteries for the same price spreads even as it substitutes for them, so deeper data-centre flexibility reshapes the battery investment case. We recommend the inquiry treat that interaction as a modelling priority; it is one of ours.

5. Why acting early matters

We tested a realistic adoption path: flexibility ramping from zero today to 30% of computing by the mid-2040s. That ramp delivers only about as much value as a constant 21% had delivered from the start. The early, under-deployed years forgo savings precisely while the fleet, and therefore the prize, is growing fastest.

Put simply: a flexibility commitment negotiated now, while facilities are on the drawing board, is worth about 28% more than one phased in as the fleet builds out. For data centres the technical barrier to retrofitting flexibility is lower than for industrial plant, because job-scheduling software is already core to their operation; the harder retrofit is contractual, since customer service agreements written around uninterrupted running must be renegotiated. Committing at design time avoids both. The window in which policy can shape the fleet cheaply is the next few years.

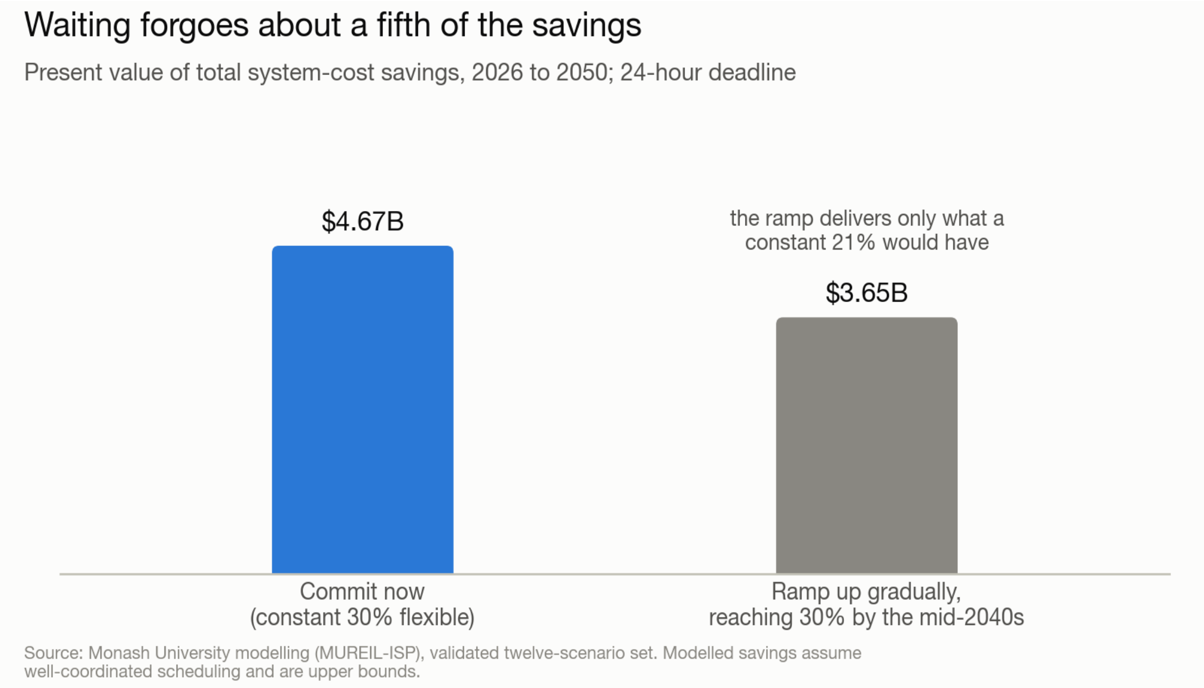

Figure 9. Both paths reach 30% flexibility; the difference is purely when the commitments arrive.

6. What we recommend

Recommendation 1: Make flexibility part of the connection deal. Connection frameworks for large new loads should recognise, reward and, where appropriate, expect a declared flexible capability. Our central case assumes that 20% of IT computing work can be rescheduled within a 24-hour completion window (rather than 20% of total facility load being interruptible). In practice, any flexibility credited through a connection agreement should be specified and verified as deliverable MW reduction at the facility meter, including its duration and recovery requirements. Flexibility need not be a mandate from the system operator. For developers it is also a lever to lower connection costs and shorten connection timeframes, because a facility that can reduce demand in the tight hours needs less network and system capacity reserved for it; connection frameworks should let that bargain be struck explicitly.

Recommendation 2: Act in this decade. Because early flexibility is worth about 28% more than late flexibility, governments should not wait for the fleet to build out before acting. Every connection agreement signed without a flexibility provision is a small, permanent loss to consumers.

Recommendation 3: Keep the ask proportionate. The modelling shows most of the value comes from modest commitments: a minority of workload, a day of leeway, delivered in a small number of critical hours. Policy should not demand extreme deferral windows; it should secure broad participation in moderate ones. Even a six-hour window captures more than half of the week-long value (Table 1), so the entry ask can be very modest.

Recommendation 4: Pair any 24/7 clean-energy matching obligation with flexibility recognition. If governments impose or encourage hourly clean-energy matching for large loads, they should recognise workload flexibility as a compliance pathway. Under a binding national carbon budget, procurement obligations change who pays for the transition rather than total emissions, so an obligation designed without flexibility would be expected to raise compliance costs without climate benefit. The package of matching plus flexibility aligns private and system interests; either alone does not.

Recommendation 5: Count what is already inside the building. Data centres contain substantial batteries, uninterruptible power supplies and standby generation that the planning framework currently treats as invisible. Connection standards for new large loads should address whether and how this equipment supports the grid, for example through grid-forming capability requirements per MW of new load. Our quantitative work on this is under way; the direction is consistent with the grid-security cost channel described in section 2, and we would welcome the consultation treating it as a live design question.

Recommendation 6: If the planner counts promised flexibility, price the promise and price the breach. A system planner who defers building capacity because large loads have promised flexibility is carrying delivery risk on consumers' behalf. Promised flexibility should be credited prudently, verified against delivery, and backed by penalty or collateral arrangements when it fails to show up. We are completing modelling on prudent crediting levels; the accountability principle stands independently of the final numbers. In other words, if the planner builds less because data centres promised flexibility, someone carries the risk that the promise fails. Crediting prudently means counting only part of what is promised until delivery is proven, and backing the promise with penalties or collateral, as capacity commitments are treated in other markets.

7. What this submission does not claim

This submission does not claim:

- That flexibility reduces emissions. Under the legislated-style carbon budget in our modelling, cumulative emissions are identical in every scenario. Flexibility changes the cost of meeting the budget and who pays, not the tonnes.
- That our data-centre demand profile is AEMO's. AEMO publishes no hour-by-hour data-centre trace, so we built a documented synthetic profile whose annual totals match AEMO's forecast exactly by

construction. Testing how sensitive the results are to the assumed hourly shape is part of our ongoing robustness programme, and fine hourly detail should not be quoted from this work.

- That all data-centre demand will be grid-supplied. Our scenarios follow AEMO's forecast, which nets off no self-supply. Facilities arriving with firm generation of their own would reduce the modelled burden; we have not modelled that case, and the size of the reduction depends on whether the self-supply is dependable in the scarce hours where the burden concentrates. Nor are the data centres' own grid-connection costs counted anywhere; flexible operation would tend to reduce the connection capacity a facility needs and shorten its connection process, a private benefit additional to the system savings modelled here.
- That the savings will be captured in full. As stated in section 4, the modelled savings assume well-coordinated scheduling and are upper bounds.
- Anything about named companies or sites. All modelling uses archetypes, and nothing here should be read as endorsing or criticising any specific facility or proposal.
- Final numbers on connection standards, tariffs, or crediting rules. Recommendations 5 and 6 rest on completed system-level evidence for their direction; the quantitative design work is in progress and can be shared with the consultation as it completes.

8. About the Evidence

As with any long-term electricity-system model, these results should be interpreted as scenarios rather than forecasts. The model simplifies a complex system and necessarily relies on assumptions. Its value is not in predicting a precise future system or cost, but in comparing scenarios on a consistent basis and showing how additional data-centre demand changes the scale and composition of the investment required under the assumptions used.

The findings in Sections 2 to 5 come from 12 solved scenarios of MUREIL-ISP, an hourly optimisation model of the national electricity market developed at Monash University, aligned with AEMO's 2026 Integrated System Plan (Step Change), with data-centre demand represented explicitly at the facility level and workload deadlines enforced hour by hour. Every scenario passed a pre-registered set of validation checks, including exact conservation of annual computing work and energy across scenarios, which is what allows bill differences to be read as pure price effects. System costs and savings are reported as net present values of the 2026-to-2050 cost difference in real 2026 dollars; bills and prices are undiscounted annual amounts for the year stated, computed as hourly consumption valued at wholesale spot prices; network and retail cost components are outside the model. The cost base covers generation, storage, grid-security plant and the nine candidate transmission corridors in the ISP; it excludes local connection assets, distribution networks and retail costs. AEMO figures are drawn from the 2026 ISP demand forecasts, the 2025 Electricity Statement of Opportunities, and the March 2026 connections pipeline. Full technical documentation, including the complete scenario set and validation record, is available on request.

We would welcome the opportunity to discuss any part of this evidence with the consultation team, including running additional scenarios against specific policy designs under consideration.

ANNEX: THE KEY NUMBERS AT A GLANCE

All dollar figures below are in real 2026 dollars. System costs and savings are present values: the total cost difference from today to 2050, with each year's dollars discounted back to 2026. Bill and price figures are annual amounts for the year stated, undiscounted, and refer to the wholesale (spot) market component only; network charges and retail margins are not modelled.

What we measured	Result
Cost of inflexible data-centre growth to 2050	+A$18.5 billion (+9.1% of system cost)
Extra capacity it forces by 2050	7.9 GW wind, 9.9 GW solar, 3.9 GW batteries, plus accompanying grid-security plant
Consumer share of the 2050 bill increase	58% (A$6.1 billion a year; +$31.2/MWh; winter +$102.7/MWh; tightest hours +$341/MWh)
Saving from 20% flexible IT-work share with a 24-hour completion window	A$3.5 billion (A$4.7 billion at 30%)
Who keeps the flexibility bill relief	Consumers, roughly 60% to two thirds in every flexible case; the rest goes to operators
What flexible computing pays for electricity	$24 to $38/MWh in the central and deepest cases, against about $140 at arrival time
Delivery in the tightest hours	96% of nominal flexible demand removed in the top 5% of hours
Value of the first day of leeway	The first 6 hours of allowable delay capture 54% of week-long value
Cost of adopting late	A ramp to 30% by 2050 delivers only a constant-21% equivalent
Emissions	Identical in every scenario, because the carbon budget binds in every modelled year