

Storage and Signals: Solar and Battery Storage, Time-of-Use Pricing, and Household Peak Demand in the Australian NEM

VEPC Working Paper 2609

Bruce Mountain*
September 2026

Abstract

Time-of-use (TOU) electricity tariffs are intended to encourage households to shift consumption away from periods of peak network and generation stress, but their effectiveness remains uncertain and is complicated by the rapid uptake of behind-the-meter solar and battery storage. This paper uses a panel of 5,096 Australian households across four states, drawn from a major retailer and spanning no-solar, solar-only, and solar-plus-battery cohorts, observed at 30-minute resolution over three years, to examine two questions: whether households respond to the size of the price differential a TOU tariff embeds, and whether owning solar or battery storage changes peak-period consumption independently of price. Household fixed-effects models exploiting genuine within-household tariff-price variation find no evidence that households of any type respond to the magnitude of the peak-to-off-peak price differential. A “shadow peak” measure, classifying consumption by standard TOU windows regardless of a household’s actual tariff, shows that solar-plus-battery households consume significantly less during peak hours than other households on the same tariff - around 13% at the sample mean - but only when actually enrolled in a TOU tariff. On flat tariffs, battery households are statistically indistinguishable from non-solar households, and solar ownership alone shows neither pattern. These results are inconsistent with both continuous price elasticity and an unconditional structural effect of storage, pointing instead to a discrete, tariff-triggered mechanism. The findings corroborate Burns and Mountain (2021) in a new panel setting and suggest TOU tariffs may be most effective when paired with battery storage rather than pursued as a stand-alone demand-response policy.

Key words: Time-of-use pricing; battery storage; demand response; price elasticity; residential electricity demand

JEL codes: D12, L94, Q41, Q42, Q48

*Victoria University, bruce.mountain@vu.edu.au.

1. Introduction

Time-of-use (TOU) electricity tariffs are a central plank of Australian and international retail market reform. Time variant prices give households a price signal that is intended to reflect the higher cost of supplying electricity during periods of peak network and generation stress, and thereby to encourage a shift of discretionary consumption away from those periods. Whether households actually respond to these signals has been tested repeatedly since the earliest TOU pricing experiments of the 1970s, generally with modest and inconsistent results (Section 2).

A newer and largely separate question has emerged alongside the rapid uptake of rooftop solar photovoltaics (PV) and, more recently, behind-the-meter battery storage across the Australian National Electricity Market (NEM): does the physical capability to store and dispatch electricity change how much a household draws from the grid during peak hours. And, if so, is that change driven by households responding to the price they face, or does it occur largely independently of price, as a structural consequence of owning storage?

This paper addresses both questions using a panel of approximately 5,096 households across four Australian states (New South Wales, Victoria, Queensland and South Australia), observed at 30-minute resolution over a three-year period. Each household's actual retail tariff, including whether it varies by time of day and by how much, is known precisely. The panel includes households with no solar, solar only, and solar-plus-battery, allowing a direct comparison of how each group's grid consumption is distributed across peak, shoulder and off-peak periods, and how that distribution relates to the price differential each household actually faces.

The paper's main methodological contribution is to separate these two questions empirically rather than treating them as a single test of “demand response”. First, exploiting genuine within-household variation in retail tariff prices over the panel (a majority of sample households experience at least one price change, around half of those mid-year rather than only at an annual reset) the paper estimates household fixed-effects models of the elasticity of substitution between peak and off-peak grid purchases with respect to the peak-to-off-peak price ratio, interacted with solar and battery ownership.

Second, and independently of whether a household is actually exposed to a time-varying price at all, the paper constructs a “shadow peak” measure that classifies every household's actual half-hourly consumption according to each state's standard time-of-use windows, regardless of the tariff that household is actually on. This allows a comparison of peak-period consumption across households with no solar, solar only, and solar-plus-battery, cross-classified by whether each household is actually exposed to a time-of-use tariff at all. This produces, in effect, a difference-in-differences-style comparison across ownership status and tariff exposure simultaneously, even though the panel's own within-household variation in tariff status is too limited to identify this with household fixed effects (Section 5.1).

The results are, in combination, sharper and more specific than either question would produce alone. They point to a mechanism that is neither pure price elasticity nor an unconditional structural effect, but something in between. There is no significant evidence, across several specifications, that households of any type (including those with batteries) alter peak-period grid purchases in response to the size of the peak-to-off-peak price differential they face (Section 6.1). Separately, and this is the paper's central finding, solar-plus-battery households show a significant reduction in shadow peak-period consumption relative to other households on the same tariff exposure, but only once tariff exposure is held fixed. Battery households on a time-of-use tariff show materially and significantly lower peak-period consumption than other time-of-use households, while battery households on a flat tariff look statistically indistinguishable from non-solar flat-tariff households. Solar ownership alone, without battery storage, shows neither pattern.

Taken together, these results are not consistent with a purely structural, tariff-independent mechanism (which would predict a reduction for battery households regardless of tariff), nor with continuous price elasticity (which the paper directly rejects in Section 6.1). Instead, they point to a discrete, threshold-triggered mechanism: batteries appear to be configured, plausibly at installation, or by automated dispatch software responding to a known tariff structure to discharge around peak hours specifically when the household is on a tariff that

distinguishes peak from off-peak pricing, regardless of how large that price differential actually is. Time-of-use pricing, in other words, appears to matter a great deal for how a battery is used, without households responding at all to the size of the price signal a given tariff embeds.

The study period also has a specific characteristic that shapes how the results should be read. The battery-owning households observed here installed their systems during an early phase of Australian residential battery uptake, when typical system capacities were modest by current standards. Battery capacity in the residential market has grown substantially since 2025. A materially larger-capacity cohort of installations has followed in the roughly twelve months preceding this paper. If the mechanism identified here is a discrete, tariff-triggered configuration decision rather than a continuous price response, it raises a question this paper's own data cannot answer but a future extension can: does a materially larger battery amplify the same tariff-conditional pattern (a bigger reduction specifically among TOU-exposed battery households), or does sufficient storage capacity eventually decouple discharge behaviour from tariff status altogether, as self-consumption maximisation dominates regardless of the household's specific price signal? This paper's findings serve as a baseline against which that question can be tested once comparable data covering the larger-capacity cohort becomes available (Section 7).

This paper makes three contributions to the literature on household electricity demand response. First, it estimates household-level time-of-use price elasticity using genuine within-household variation in retail tariff prices over a multi-year panel of revealed consumption, rather than relying on a designed pricing experiment, a single cross-section of bills, or a simulated dispatch strategy. Second, it separates price-driven from storage-driven behaviour by constructing a “shadow peak” measure that classifies every household's actual consumption according to standard time-of-use windows regardless of the tariff that household is actually on, allowing a direct empirical comparison between households with and without a live price signal. Third, it does so across four Australian states with materially different tariff structures and time-of-use definitions and explicitly distinguishes solar-only from solar-plus-battery households throughout, a distinction the price-elasticity literature has not generally made central to the empirical design.

The remainder of the paper proceeds as follows. Section 2 reviews the relevant literature on time-of-use pricing and battery storage. Section 3 sets out the institutional background of Australian retail electricity tariffs. Section 4 describes the data and the construction of the shadow peak measure. Section 5 sets out the empirical strategy. Section 6 presents the results. Section 7 discusses the findings and their policy implications. Section 8 concludes.

2. Literature Review

The literature on residential response to time-of-use (TOU) electricity pricing dates to the large-scale U.S. and European pricing experiments of the late 1970s and 1980s. Early econometric work - Caves and Christensen (1980), Herriges, Caves and Kuester (1989), and Filippini's (1995) application of the Almost Ideal Demand System to Swiss household data established that residential customers do shift consumption between peak and off-peak periods in response to relative price, but that estimated elasticities are typically modest, sensitive to the specific pricing structure tested, and vary considerably across studies and populations. This body of work forms the empirical baseline against which more recent studies, including the present one, continue to be judged: a genuine price response is neither a foregone conclusion nor, where found, typically large.

A recurring theme in this literature, formalised by Ito (2014), is that households often respond to the average price implied by their bill rather than the marginal price associated with a specific consumption decision at a specific time. This distinction matters directly for interpreting a null or weak response to a TOU tariff's marginal peak/off-peak price differential, since such a null result need not imply households are inattentive to price altogether, only that they are not optimising against the specific marginal signal a TOU tariff is designed to convey.

More recent work has emphasised that price responsiveness is far from homogeneous across the customer base. Cappers and Todd-Blick (2021), using data from a U.S. utility randomised-control trial published in Utilities Policy, find that peak-period elasticity varies substantially by usage level, air-conditioning ownership and notably for the present study's identification strategy, by whether a customer was defaulted onto a TOU rate or opted in voluntarily, with defaulted customers exhibiting only around a quarter of the elasticity of voluntary adopters. This heterogeneity result is a direct caution against treating a household's presence on a TOU tariff as if it were

randomly assigned and motivates the state-composition and consumption-scale controls used in the present analysis, as well as the more fundamental point that any cross-sectional comparison between TOU and flat-rate households risks conflating a genuine tariff effect with selection into tariff type.

Choo, Jang and Kim (2025), using a difference-in-differences design around the introduction of an opt-in TOU tariff on Jeju Island, South Korea, find genuine average load-shifting among enrolled households, but also evidence of strategic sorting: high-consumption households enrolling in the tariff to reduce their bills without materially changing their consumption timing. This is a useful cautionary parallel: a positive cross-sectional association between TOU enrolment and some outcome of interest need not reflect a behavioural response to the tariff at all.

Within the Australian National Electricity Market, Burns and Mountain (2021) provide the closest direct precedent, estimating the elasticity of substitution between peak and off-peak consumption using 6,957 household electricity bills from Victoria. They find essentially no response to the peak/off-peak price differential - an elasticity of substitution closely matching estimates from US TOU pricing experiments conducted four decades earlier - and find this null result is unaffected by rooftop solar PV access, with lowest-socio-economic-status households showing no response at all. They conclude that market deregulation, smart metering and improved consumer access to usage data have not delivered any measurable improvement in price responsiveness, and that their findings do not support the imposition of TOU tariffs as a default pricing policy.

The present paper's price-elasticity finding i.e a similarly robust null, now using within-household panel variation across four states rather than a single cross-section of bills corroborates this result in a distinct empirical setting and time period. It extends Burns and Mountain's PV-access check, conducted prior to widespread battery adoption, by distinguishing solar-only from solar-plus-battery households explicitly, and by separating price-driven from storage-driven behaviour via the shadow-peak construction. This allows the paper to show that while price responsiveness remains negligible even for battery-owning households, storage ownership is nonetheless associated with a substantively large reduction in peak-period grid imports as a result of structural changes (the installation of a battery) rather than in response to peak/off-peak price differences.

A recent empirical contribution addresses part of this gap in a different institutional setting. Qiu, Xing, Patwardhan, Hultman and Zhang (2022) (2022), using individual-consumer smart-meter data from 56 solar-plus-battery adopters and a matched control group of non-battery solar households in Arizona, find that battery adoption reduces electricity delivered from the grid during late-afternoon and early-evening hours, but that the degree to which this reduction aligns with the economic predictions implied by each household's specific tariff is highly heterogeneous: some time-of-use customers show discharge patterns broadly consistent with bill-minimising behaviour, while others show reductions not concentrated in peak hours at all, and even households on a flat, non-time-varying rate show a comparable evening-hours reduction in grid draw. This is the closest existing empirical counterpart to the shadow-peak approach used in the present paper, and its central finding that substantial heterogeneity in whether battery discharge timing tracks the economic incentives embedded in a household's specific tariff motivates the present paper's approach. Where Qiu et al. examine whether individual households' discharge timing matches the specific incentives of the tariff each is on, the present paper instead asks a population-level question made possible by a multi-cohort, multi-state dataset: does the mere presence of a time-varying tariff, independent of its magnitude, systematically shift battery households' peak-period consumption relative to otherwise comparable households, and does this differ from solar-only and non-solar households? This reframes the question from individual bill-optimisation to a population-level comparison and adds two further points of distinction from Qiu et al.: it tests responsiveness to the size of the peak-to-off-peak price differential directly (Section 6.1), which Qiu et al.'s design does not address, and it does so across four Australian states with materially different tariff structures rather than a single US utility's rate plans.

The literature specifically addressing behind-the-meter battery storage in a TOU context is considerably thinner than the pure price-elasticity literature and is more often engineering- or optimisation-oriented than econometric. Sharma, Aziz, Haque and Kauschke (2022), studying South Australian households find via a cost-minimisation simulation that battery storage paired with a TOU-with-demand-charge tariff can reduce peak-feeder load by up to 35% relative to an ordinary TOU tariff without storage, a reduction the authors attribute to optimised charge/discharge scheduling rather than active household price-responsiveness per se. Restel and Say (2025),

using dispatch-optimisation modelling on 400 real Melbourne household load profiles each fitted with a hypothetical 9kWp PV and 12kWh battery system, find that time-varying export tariffs have substantially more influence on household grid-interaction patterns than time-varying import tariffs, and that under import-only time-varying pricing, PV-battery households continue to rely on self-consumption in the evening regardless of the price signal. This is consistent with the present paper's result that battery discharge behaviour during the evening peak appears largely independent of the specific price differential a household faces. Both studies, being simulation-based, model an assumed or optimised dispatch strategy rather than observing revealed household behaviour, which is the specific gap the present paper's empirical approach is designed to fill.

This paper contributes to these literatures in three ways. First, it estimates household-level TOU price elasticity using genuine within-household variation in the peak-to-off-peak price ratio over a multi-year panel of revealed consumption, rather than relying on a designed pricing experiment, a single cross-section of bills, or a simulated dispatch strategy and, unlike Qiu et al. (2022), it tests responsiveness to the magnitude of the price differential directly, rather than only whether discharge timing is broadly consistent with a given tariff's structure. Second, it separates price-driven from storage-driven behaviour by constructing a "shadow peak" measure that classifies every household's actual consumption according to standard TOU time windows regardless of the tariff that household is actually on, allowing a direct empirical comparison between households with and without a live price signal — extending the revealed-behaviour approach of Qiu et al. (2022) from a single-utility, single-state setting to a multi-cohort comparison expressly designed to separate price-conditional from unconditional storage effects. Third, it does so across four Australian states with materially different tariff structures and time-of-use definitions, and explicitly distinguishes solar-only from solar-plus-battery households throughout, a distinction the price-elasticity literature, including Burns and Mountain (2021), has not generally made central to the empirical design.

3. Institutional Background

In the ten distribution (network service provider) areas in the four regional markets in Australia covered by this study, households have been free to choose their electricity tariffs and retail offers from numerous competing retailers. The penetration of rooftop solar and behind-the-meter batteries is higher in Australia, on average, than in any other country. There has long been, and continues to be, significant policy support for behind-the-meter rooftop solar and batteries. Australia's regulators have also long promoted "cost-reflective" network tariffs, with encouragement for temporal price variation. Regulated network service providers, through their regulated network tariff structures, effectively determine the structure of prices offered by competing retailers. The selection of time-of-use tariffs by households has been on an opt-in basis in some distribution areas but more recently has become the default on installation of interval meters or the installation of rooftop solar and/or batteries with the option typically, but not always, to opt out of such temporally varying price structures.

To research the effect of distributed energy in this context, considering particularly the growing uptake of batteries, the author sought access to household data for customers supplied by AGL Energy, one of the largest retailers in all the regional markets in this study. The author requested deidentified data for around 500 households in each of three cohorts (households with neither solar nor battery, households with solar, and households with solar and batteries) for three years from 15 October 2023) in each of the four NEM regions. AGL Energy accordingly randomly selected households they supplied, for inclusion in this study.

4. Data

4.1 Sample and panel structure

The dataset covers approximately 5,096 households across four States (NSW, VIC, QLD, SA), with 30-minute grid import and export volumes collapsed to a daily panel of 3,049,653 site-day observations over a three-year period to August 2025. The data was supplied by AGL Energy, a major Australian electricity retailer with a large number of customers in each State. Households are split by solar and battery ownership status (no solar; solar only; solar plus battery), fixed for the whole study period. No household in the sample installed solar or battery mid-panel.

4.2 Tariff classification

Each household's actual electricity tariff is known, including peak, shoulder and off-peak prices where applicable. Households are classified into three tariff types: flat rate (a single price at all times), two-rate time-of-use (peak and off-peak only), and three-rate time-of-use (peak, shoulder and off-peak). Each network/retailer combination defines peak, shoulder and off-peak hours differently:

- NSW (two-rate): peak 3pm–9pm, seasonal (Nov–Mar and Jun–Aug only); all other times off-peak.
- NSW (three-rate): peak 2pm–8pm weekdays; shoulder 7am–2pm and 8pm–10pm weekdays, 7am–10pm weekends; off-peak 10pm–7am all days.
- VIC: peak 3pm–9pm every day, year-round; all other times off-peak (no shoulder rate exists in this state).
- QLD (Energex): peak 4pm–9pm; shoulder 9pm–9am; off-peak 9am–4pm.
- SA: peak 6am–10am and 3pm–1am (wrapping past midnight); shoulder 10am–3pm; off-peak 1am–6am.

4.3 The shadow peak measure

For the price-responsiveness question (Section 5.1), only households on a two-rate or three-rate tariff are included, using their own tariff's definition of peak. For the physical/behavioural question (Section 5.2), a “shadow” peak measure is constructed by applying each state's standard TOU time windows to every household's actual consumption, regardless of what tariff that household is on. So, a flat-rate household's electricity use during what would be peak hours can still be measured and compared to solar/battery households on the same basis. Because NSW has two live, differing peak definitions (the two-rate and three-rate products), two shadow measures are constructed in parallel: Shadow A applies NSW's two-rate definition, Shadow B applies NSW's three-rate definition; the other three states have only one native definition and are identical under both. A third, narrower definition (5pm–9pm, year-round, every day) is used as an additional robustness check.

4.4 Household counts by state, distributor and tariff type

Table 1 reports the number of households in each state by tariff type, aggregated across distributors within each state; the underlying distributor-level breakdown is given in the Appendix (Table A1). Four NSW households, for whom all three price fields are recorded as zero (no tariff data at all), are excluded from the classification and from the analysis; the remaining 5,092 households are used throughout, with 5,096 the count including these four.

Table 1. Households by state and tariff type

State	Flat rate	two-rate	three-rate	Total
NSW	867	61	450	1,378
VIC	584	723	0	1,307
QLD	1,016	26	101	1,143
SA	185	2	1,077	1,264
Total	2,652	812	1,628	5,092

QLD figures reflect a single distributor (Energex); no Ergon Energy (the other Queensland distributor) households are present in the sample since households in that distribution area are not able to choose their retailer.

Tariff structure varies markedly across states, driven by differences in retailer product design and, plausibly, differences in smart-meter rollout history and default tariff assignment across jurisdictions (Section 3). Victoria has no three-rate product in this sample at all, consistent with the absence of a shoulder-rate concept in the state's standard time-of-use tariff (Section 4.2); South Australia and Queensland are dominated by three-rate and flat-rate tariffs respectively; New South Wales shows the most even three-way split.

Table 2 reports the prevalence of controlled load and demand-charge components among households on a time-of-use tariff (2-rate or 3-rate) in each state. “Controlled load” a separate, timed or network-switched circuit typically used for hot water, is present for a meaningful share of TOU households in every state except Victoria, where it is negligible (under 2% of Victorian TOU households). Demand charges, by contrast, are confined almost entirely to a single distributor: 48 households in the Ausgrid area on three-rate TOU tariffs carry a demand

charge, of which 15 also carry controlled load. No demand-charge households are present in VIC, QLD or SA in this sample.

Table 2. Controlled load and demand charges among time-of-use households, by state

State	TOU households	w/ Controlled load	w/ Demand charge	w/ Both
NSW	511	136 (26.6%)	48 (9.4%)	15 (2.9%)
VIC	723	13 (1.8%)	0	0
QLD	127	39 (30.7%)	0	0
SA	1,079	336 (31.1%)	0	0

TOU households = 2-rate + 3-rate households in that state (Table 1). "Both" = households with both controlled load and a demand charge (all 15 in the Ausgrid distributor's area).

4.5 Solar and battery system size

The dataset does not include registered or nameplate solar PV or battery capacity for each household. Following the same approach used to estimate each household's grid import capacity (Section 4.3), solar export capacity is estimated from the household's own metering data: the household's highest recorded half-hourly grid export volume is doubled to obtain an implied kW export capacity, and the maximum of this measure across the household's full observed history is taken as its estimated export/system capacity. This is a lower-bound proxy for nameplate solar capacity. A household's peak observed export will understate its true PV capacity to the extent that some on-site generation is self-consumed rather than exported at the moment of peak output. The measure also does not distinguish solar generation capacity from battery discharge capacity where both contribute to a single export reading.

A natural concern is that “max_export_capacity” and “annual_sold” are constructed from the same underlying export metering stream that also generates “daily_sold”, the export-based variables used elsewhere in the broader project. This does not create circularity for the specifications reported in Section 5, however, since none of the paper's dependent variables (ihs_daily_peak and its shadow variants) are themselves export based. Instead, they are constructed entirely from import-side metering. Excluding export-side scale controls while including import-side ones (“annual_bought”, “max_import_capacity”) would be inconsistent without a substantive reason to treat the two sides of the meter differently, and no such reason is apparent. Accordingly, “annual_sold” and “max_export_capacity” are included alongside their import-side counterparts as consumption-scale controls throughout Section 5. Because these variables are mechanically zero, not missing, for every non-solar household, they are collinear with has_solar and has_solarbattery to some degree, sharpening the interpretation of the solar/battery coefficients to “conditional on how much this household exports, does ownership status still predict peak-period import behaviour”.

4.6 Summary statistics by cohort

Table 3 reports household-level summary statistics by cohort, collapsing the panel to one row per household, with “annual_bought” and “annual_sold” calculated as each household's average daily grid purchases/exports over its full observed panel, annualised to 365 days, rather than summed within calendar year, since the panel's start and end dates (15 October 2023 to 15 October 2025) leave both 2023 and 2025 as partial years for essentially every household.

Table 3. Household characteristics by cohort

Variable	No solar	Solar only	Solar+battery
Annual grid purchases (kWh), mean (SD)	5,380.0 (4,310.2)	5,231.7 (4,193.3)	5,453.4 (3,884.0)
Max import capacity (kW), mean (SD)	2.11 (1.84)	2.35 (1.99)	2.41 (2.21)
Annual grid exports (kWh), mean (SD)	0.0 (0.0)	4,882.9 (3,677.1)	4,663.5 (3,091.0)
Max export capacity (kW), mean (SD)	0.00 (0.00)	3.97 (2.96)	3.88 (3.07)

Share with controlled load	28.3%	25.7%	16.1%
Share on flat rate	64.3%	47.4%	39.7%
Share on two-rate TOU	4.7%	19.5%	28.8%
Share on three-rate TOU	31.0%	33.2%	31.5%

N = 1,950 (no solar), 1,973 (solar only), 1,173 (solar+battery), 5,096 total. `annual_sold` and `max_export_capacity` are by construction zero for all no-solar households.

Correctly annualised, the gap in total grid purchases between cohorts is modest. Solar-plus-battery households purchase around 1% more than non-solar households on average, and solar-only households around 3% less. Max import capacity still shows a modest, genuine difference across cohorts (12–14% higher for solar/battery households than non-solar households).

Two other patterns in Table 3 remain worth flagging for the empirical strategy in Section 5. First, solar and solar-plus-battery households naturally show substantial annual grid exports and export capacity, entirely absent for non-solar households, and hence is the basis for including `annual_sold` and `max_export_capacity` as controls in Section 5 alongside their import-side counterparts (Section 4.5). Second, and directly relevant to the identification concerns raised by Capperaud and Todd-Blick (2021) and Choo, Jang and Kim (2025) in Section 2, solar and solar-plus-battery households are far more likely to be on a time-of-use tariff than non-solar households — 28.8% of solar-plus-battery households are on a two-rate tariff versus only 4.7% of non-solar households. Evidently TOU exposure (“`faces_tou`”) is not randomly assigned with respect to solar/battery ownership, and reinforcing the importance of the state and consumption-scale controls used in the physical/behavioural specification (Section 5.1).

5. Empirical Strategy

5.1 Identification

Solar and battery ownership are fixed for each household for the whole study period, so they cannot be identified using household fixed effects, which would absorb any time-invariant characteristic entirely. Two different identification strategies are therefore used, matched to the two research questions and to how much genuine within-household variation each requires.

For price responsiveness, peak and off-peak prices vary over time for a majority of households in the sample (roughly half see a price change at some point, about half of those mid-year rather than only at an annual reset), so a household fixed-effects specification is used, exploiting within-household changes in the peak-to-off-peak price ratio. Solar/battery main effects cannot be estimated in this specification, since they are absorbed by the household fixed effect, but their interaction with the price ratio can be, since price varies over time even though ownership status does not.

For the physical/behavioural response, whether a household is on any TOU tariff at all never changes within the panel so fixed effects are not viable for this question. A random-effects specification is used instead, drawing on cross-sectional variation across households. State fixed effects are included throughout this specification: TOU uptake varies enormously by state in this sample (from 7.9% of observations in Queensland to 89.1% in South Australia), and an earlier specification without state fixed effects showed a spurious positive association between TOU exposure and shadow peak consumption among non-solar households that was substantially, though not fully, resolved once state composition was controlled for. A residual, unexplained version of this association persists even with state fixed effects and is discussed as a limitation in Section 6.2 and Section 7.

Because the random-effects specification interacts ownership status with tariff exposure, two distinct comparisons can be recovered for each ownership cohort from the same set of coefficients: the effect of ownership relative to other households with the same tariff exposure, and the effect of ownership relative to the non-solar, flat-tariff baseline. These need not agree in sign or significance. And, as Section 6.2 shows, do not always agree in this application because the second comparison also picks up whichever relationship time-of-use exposure itself has with the outcome for non-solar households, discussed above.

Both specifications use the inverse hyperbolic sine (asinh) transform of the relevant consumption measure as the dependent variable, since a meaningful share of households have zero purchases in a given period on a given day and a standard log transform would drop these observations rather than treat them as genuine, informative zeros. Standard errors are clustered by household throughout. Controlled loads, which is present for a meaningfully different share of solar-plus-battery households (16.6%) than other households (26.4%), is included as a control given it represents largely non-discretionary consumption.

5.2 Specifications

The IHS transform used throughout is defined as:

$$\text{asinh}(y) = \ln(y + \sqrt{y^2 + 1})$$

Price responsiveness (household fixed effects, TOU households only):

$$\text{IHS}(Peak_{it}) = \beta_1 \ln(P_{it}) + \beta_2 [\ln(P_{it}) \times Solar_i] + \beta_3 [\ln(P_{it}) \times Battery_i] + \tau_t + \alpha_i + \varepsilon_{it} \quad (1)$$

where P_{it} is the household's peak-to-off-peak price ratio, τ_t denotes month fixed effects, and α_i denotes household fixed effects. `annual_bought`, `annual_sold`, `max_import_capacity` and `max_export_capacity` are omitted from this specification: each is constructed as a single value per household across the full panel, so each is collinear with, and fully absorbed by, α_i .

Physical/behavioural response (random effects, full sample, state FE):

$$\begin{aligned} \text{IHS}(Peak_{it}^{shadow}) = & \beta_1 Solar_i + \beta_2 Battery_i + \beta_3 (Solar_i \times TOU_i) + \beta_4 (Battery_i \times TOU_i) \\ & + \beta_5 TOU_i + \beta_6 (CL_i \times TOU_i) + \beta_7 CL_i + \mathbf{X}_i' \boldsymbol{\gamma} + \delta_s + \tau_t + u_i + \varepsilon_{it} \end{aligned} \quad (2)$$

where $TOU_i = \text{faces_tou}$, $CL_i = \text{has_controlled_load}$, $\mathbf{X}_i$ is the vector of remaining scale/confound controls (`annual_bought`, `max_import_capacity`, `has_ev`), $\boldsymbol{\gamma}$ their coefficient vector, δ_s denotes state fixed effects, τ_t month fixed effects, and u_i the household-level random effect. The interaction between CL_i and TOU_i is included because controlled load's own association with peak-period consumption is not uniform across tariff types (Section 6.2): a single main-effect term for controlled load, without allowing its relationship with the outcome to vary by tariff exposure, was found to materially affect the estimated β_5 coefficient and is retained throughout as part of the preferred specification.

As discussed in Section 5.1, two related but distinct quantities can be recovered from equation (2) for the solar-plus-battery cohort (and, equivalently, the solar-only cohort using β_1 and β_3 in place of β_2 and β_4): the effect relative to other households on the same tariff exposure, and the effect relative to the non-solar, flat-tariff baseline: The relative effect is β_2 plus β_4 ; the absolute effect is β_2 plus β_4 plus β_5 .

Both are reported in Section 6.2, since they answer different questions and, as the results show, are not always the same sign or the same level of significance.

6. Results

6.1 Price responsiveness

Table 4 reports the household fixed-effects regression of peak-period grid purchases on the log peak-to-off-peak price ratio, restricted to households on a two-rate or three-rate tariff. Because `annual_bought`, `annual_sold`, `max_import_capacity` and `max_export_capacity` are all constructed as fixed household-level values across the full

panel, each is collinear with, and absorbed entirely by, the household fixed effect; none is included as a separate regressor in this specification, and none would contribute anything the fixed effect does not already capture.

Table 4. Dependent variable: *ihd_daily_peak* (TOU households, household fixed effects)

Variable	Coef.	Std. Err.	t	p-value
<i>ln_price_ratio</i>	-0.0478	0.1833	-0.26	0.794
<i>ln_price_ratio</i> × <i>has_solar</i>	0.3569	0.2313	1.54	0.123
<i>ln_price_ratio</i> × <i>has_solarbattery</i>	0.1996	0.2633	0.76	0.448

Robust standard errors clustered by siteid. N = 1,467,005; 2,418 households. Month fixed effects included but not shown. R²: within = 0.075, between = 0.002, overall = 0.023. * p<0.10, ** p<0.05, *** p<0.01.

None of the price-ratio terms are statistically significant, and a joint test cannot reject that both interaction terms are zero (F(2,2417) = 1.20, p = 0.302). This is not simply a lack of power from pooling households with very little genuine price variation: restricting the sample to the quartile of households with the greatest within-household price movement (549 households) produces materially the same result. The price-ratio coefficient remains close to zero (−0.043, p = 0.830) and neither interaction term is significant (*has_solar*: p = 0.373; *has_solarbattery*: p = 0.647).

6.2 Physical/behavioural response

Table 5 reports the random-effects regression of shadow peak-period consumption on solar/battery ownership interacted with TOU exposure, using the full sample and the Shadow B (NSW 3-rate style) peak definition, with state fixed effects, EV ownership (*has_ev*), and controlled load interacted with TOU exposure, for the reason given below, all included. Following the discussion in Section 4.5, *annual_sold* and *max_export_capacity* are not included as controls: because these are near-deterministic functions of solar/battery ownership itself (exactly zero for every non-solar household, and continuously positive for nearly every solar household as shown in Table 3). Hence including them alongside *has_solar* and *has_solarbattery* is a support/identification problem rather than a genuine control for an independent scale factor and was found to produce highly unstable estimates when tested. EV ownership, by contrast, is a well-behaved control (VIF 1.11, correlation 0.21 with *has_solarbattery*) and changes the results only modestly.

Table 5. Dependent variable: *ihd_daily_peak_shadowB* (full sample, random effects, state FE)

Variable	Coef.	Std. Err.	z	p-value
<i>has_solar</i>	0.1978	0.0308	6.41	0.000 ***
<i>has_solarbattery</i>	0.0510	0.0331	1.54	0.124
<i>has_solar</i> × <i>faces_tou</i>	-0.2179	0.0583	-3.74	0.000 ***
<i>has_solarbattery</i> × <i>faces_tou</i>	-0.1904	0.0529	-3.60	0.000 ***
<i>faces_tou</i>	0.0882	0.0294	3.00	0.003 ***
<i>has_controlled_load</i> × <i>faces_tou</i>	-0.0607	0.0727	-0.83	0.404
<i>has_controlled_load</i>	0.0562	0.0351	1.60	0.110
<i>annual_bought</i>	0.0001429	0.0000055	25.75	0.000 ***
<i>max_import_capacity</i>	-0.0013	0.0044	-0.29	0.773
<i>has_ev</i>	-0.1445	0.0451	-3.20	0.001 ***

Robust standard errors clustered by siteid. N = 3,049,653; 5,096 households. State dummies (ref. NSW): QLD +0.62, SA +1.08, VIC +0.67, all p<0.001 (not shown in full). R²: within = 0.040, between = 0.567, overall = 0.287. * p<0.10, ** p<0.05, *** p<0.01.

These coefficients support two distinct comparisons for solar-plus-battery households, and the paper's central finding turns on the difference between them. The first (the effect relative to other households on the same tariff exposure) is obtained as *has_solarbattery* + *has_solarbattery*×*faces_tou* and is −0.139 (z = −3.87, p < 0.001): among households on a time-of-use tariff. Evidently solar-plus-battery households consume significantly less during peak hours than non-solar households on the same tariff type. The second (the effect relative to the non-solar, flat-tariff baseline) is obtained by adding *faces_tou* itself (*has_solarbattery* + *faces_tou* + *has_solarbattery*×*faces_tou*) and is −0.051 (z = −1.53, p = 0.125). Evidently battery households on a TOU tariff are not statistically distinguishable from non-solar flat-tariff households. Critically, the effect for battery

households on a flat tariff (has_solarbattery alone) is small and not significant (0.051, $p = 0.124$). From this it is evident that battery ownership without TOU exposure shows no association with shadow peak consumption at all.

Read together, these three sets of results suggest that a purely structural explanation i.e. storage capacity mechanically displacing peak-period grid demand regardless of context, would predict a reduction in peak grid demand for battery households on a flat tariff too. This is not observed. Continuous price elasticity is separately and directly rejected in Section 6.1, where battery households show no response in their grid demand to the size of the peak-to-off-peak price differential. What remains consistent with the data is a discrete, tariff-triggered pattern: battery discharge is configured to avoid peak hours specifically when the household is on a tariff that distinguishes peak pricing, essentially independent of how large that differential is, and not otherwise.

Solar ownership alone shows a different and less coherent pattern, i.e. a significant positive association with shadow peak consumption on flat tariffs (0.198, $p < 0.001$) that is not present on TOU tariffs (relative-to-TOU-baseline effect $-0.218 + 0.198 =$ essentially zero net; absolute effect vs. non-solar flat baseline: 0.068, $p = 0.147$, not significant). This is consistent with no battery-specific dispatch decision being made either way, though the elevated flat-tariff baseline itself is not explained by this paper's data.

A limitation in these results is that faces_tou (i.e. customers on time-of-use tariffs) remains significantly positive among non-solar households (0.088, $p = 0.003$) even after state fixed effects, distributor fixed effects (tested separately, reported below), and the controlled-load $\times$ faces_tou interaction are all included. This means non-solar households on a time-of-use tariff show higher shadow peak-period consumption than non-solar households on a flat tariff, for reasons this paper's controls do not fully explain. Distributor fixed effects (in place of state fixed effects) reduce this association only modestly (to 0.061, $p = 0.025$) and do not eliminate it; splitting non-solar households by two-rate versus three-rate tariff type shows the pattern is not simply driven by one product (two-rate: -0.117 , $p = 0.046$; three-rate: -0.008 , $p = 0.824$, in a specification without state fixed effects), and testing whether controlled load's own relationship with the outcome differs by tariff type - which it does (Table 5) - does not resolve it either. The most plausible remaining explanation is residual selection into time-of-use tariffs among non-solar households on dimensions not observed in this dataset (household size, occupancy, or air-conditioning ownership are natural candidates), rather than a mechanical or compositional artefact this paper is able to identify and correct for. It is increasingly the case that households are automatically placed on time-of-use tariffs when they are placed on interval meters, although they are able to opt-out from this in some cases.¹ This pattern does not affect the paper's central, relative comparison for battery households, which nets out faces_tou's own level entirely, but it does mean the absolute comparison against the flat-tariff baseline should be read with this caveat in mind.

6.3 Interpreting the coefficients

Because the dependent variable is inverse-hyperbolic-sine (IHS) transformed rather than log transformed, the regression coefficients above are not directly interpretable as percentage effects in the way a log-linear semi-elasticity would be. A widely used shortcut treats an IHS coefficient as approximately equivalent to a log-linear semi-elasticity (since $\text{asinh}(y) \approx \ln(2y)$ for y away from zero, implying $\% \Delta y \approx 100 \times (e^\beta - 1)$); this approximation is unreliable, however, precisely when the dependent variable has substantial mass at or near zero - the situation IHS was chosen to accommodate in the first place, and the situation that applies throughout this paper's dependent variables.

Following Bellemare and Wichman (2020), the correct percentage effect from an IHS-linear specification is level-dependent rather than constant, and is evaluated at a chosen point - conventionally the sample mean, $\bar{y}$ - using the identity $\sinh(a+b) = \sinh(a)\cosh(b) + \cosh(a)\sinh(b)$:

$$\% \Delta y \text{ at } \bar{y} = 100 \times \frac{\sinh(\beta) \sqrt{1 + \bar{y}^2} + \cosh(\beta) \bar{y} - \bar{y}}{\bar{y}}$$

¹

For the paper's central, relative-comparison effect for solar-plus-battery households ($\beta = -0.1395$), with $\bar{y} = 5.288$ kWh (the sample mean of daily_peak_shadowB over the regression's estimation sample), this evaluates to a reduction of approximately 13.3% in shadow peak-period consumption, relative to non-solar households on the same tariff type, at the mean. Because the percentage effect is level-dependent, it should be understood as applying specifically at the sample mean; evaluated at a different point in the distribution (for example, the median, given the right-skew common to the daily consumption measures used throughout this paper) the implied percentage would differ.

6.4 Robustness to the definition of “peak”

Table 6 presents the coefficient on the solar+battery variable in the regressions for peak periods defined as Shadow A, Shadow B and an author-specified “narrow window”, using calculations of relative and absolute differences.

Table 6. Solar+battery effect across peak-window definitions: relative and absolute comparisons

Definition	Window	Relative	Absolute
Shadow A	NSW 2-rate style (3pm–9pm, seasonal)	-0.132 ***	-0.062 **
Shadow B	NSW 3-rate style (2pm–8pm weekdays)	-0.139 ***	-0.051
Narrow window	5pm–9pm, year-round, every day	-0.032	+0.102 ***

Relative = has_solarbattery + (has_solarbattery $\times$ faces_tou), i.e. vs. non-solar households on the same tariff exposure. Absolute = relative + faces_tou, i.e. vs. non-solar, flat-tariff households. All specifications include the has_controlled_load $\times$ faces_tou interaction, annual_bought, max_import_capacity, has_ev, month and state fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The relative comparison - solar-plus-battery households against other households on the same tariff exposure - is the paper's central and most robust result: negative and highly significant under both peak-window definitions grounded in actual retail TOU products (Shadow A: $z = -3.94$; Shadow B: $z = -3.87$), and null under the narrower, non-tariff-linked 5pm–9pm window. This is a coherent pattern: the effect is detected specifically when the peak window used to classify consumption corresponds to a window a household could actually be billed against and is not detected under an arbitrary window with no tariff counterpart (i.e. the narrow window), consistent with a mechanism triggered by tariff structure rather than time of day alone. The “narrow window” corresponds to the most likely hours of simultaneous peak demands in the four regional markets, and hence the times that spot market prices in these markets are likely to be higher than in other hours.

The absolute comparison (solar-plus-battery households against the non-solar, flat-tariff baseline) is considerably less stable: significant but modest under Shadow A ($p = 0.035$), not significant under Shadow B ($p = 0.125$), and significantly positive - the opposite sign - under the narrow window ($p = 0.007$). This instability arises because the absolute comparison also embeds the unexplained faces_tou association discussed in Section 6.2, which is itself larger and more consistently significant under the narrow window than under either shadow definition. The absolute comparison is reported for completeness but should not be treated as the paper's primary result; the relative comparison, which nets out the faces_tou term entirely and is stable across both tariff-grounded definitions, is preferred throughout the remainder of this paper.

7. Discussion

There is no evidence, across three separate specifications and a power-focused robustness check, that households of any cohort alter peak-period grid purchases in response to the size of the peak-to-off-peak price differential they face, corroborating Burns and Mountain (2021) in a distinct panel setting.

There is robust evidence that solar-plus-battery households consume significantly less electricity from the grid during peak hours than other households on the same time-of-use tariff, converging tightly across the two peak-window definitions grounded in an actual TOU tariff product (relative effect: Shadow A (0.132), Shadow B (0.139), both $p < 0.001$), equivalent to approximately a 13% reduction at the sample mean (Section 6.3), and null under a narrower, non-tariff-linked evening window, consistent with a mechanism triggered by tariff structure rather than time of day per se.

This effect is specific to households exposed to a time-of-use tariff: battery households on a flat tariff are statistically indistinguishable from non-solar flat-tariff households (Table 5), and solar ownership without battery storage shows neither pattern. Combined with the price-elasticity null result (Section 6.1), this rules out both a purely structural, tariff-independent mechanism and continuous price elasticity, and instead points to a discrete, threshold-triggered mechanism: batteries appear to be configured plausibly at installation, or by automated dispatch software, or by customers manually, responding to a known tariff structure to discharge around peak hours specifically when the household is on a tariff that distinguishes peak pricing, regardless of the size of that price differential. This interpretation is consistent with the simulation-based findings of Sharma et al. (2022) and Restel and Say (2025) and extends them by showing a related pattern in revealed, rather than simulated, household behaviour. The present paper's contribution is to show that this behaviour appears conditional on time of use tariff exposure specifically, not present regardless of it.

These results have significant implications for policy. The finding that time-of-use tariffs matter for peak reduction only once a household also owns a battery points to the value of complementarity between TOU tariff rollout and battery subsidy/incentive schemes, which may need to be considered jointly rather than as substitutes. If a policy objective it to reduce peak consumption there is a good argument that TOU tariffs should be marketed or defaulted specifically to battery-owning households rather than the general customer base particularly since Section 6.1's finding that non-battery households show no price response at all.

A limitation of this study is that the physical/behavioural comparison is cross-sectional for the solar/battery main effects. Since ownership is time-invariant within the panel, the usual caveat about unobserved household differences correlated with solar/battery adoption applies notwithstanding the consumption-scale, state, controlled-load, and EV controls included. A residual, significant association between TOU exposure and shadow peak consumption among non-solar households (Section 6.2) could not be resolved despite testing state fixed effects, distributor fixed effects, and a controlled-load $\times$ tariff-exposure interaction. While this does not affect the paper's preferred, relative comparison for battery households, but is disclosed as an unexplained pattern likely reflecting selection into TOU tariffs on household characteristics not observed in this dataset.

It is also noted that export-side scale measures (annual_sold, max_export_capacity) are deliberately excluded as regression controls given their near-collinearity with solar/battery ownership itself (Section 4.5). A research design capable of separately identifying export volume/capacity effects from ownership effects - for example, variation in system size among solar-owning households only - is left to future work.

A further point of comparison worth addressing directly is Qiu et al.'s (2022) finding that even non-time-of-use (flat-rate) battery households in their Arizona sample show a reduction in grid-delivered electricity during evening hours after adopting a battery. This result, at first glance, sits uneasily against this paper's central finding that flat-tariff battery households show no significant shadow-peak effect at all. Two considerations are worth noting in reconciling this. First, Qiu et al.'s flat-rate evening reduction may in part reflect the physical timing of solar generation tapering off as evening approaches, prompting battery discharge to smooth self-consumption regardless of the household's billing structure. This pattern that need not be tariff-conditional in the sense tested here. Second, and more fundamentally, the two studies estimate different quantities: Qiu et al. measure the change in grid-delivered electricity for a given household following battery adoption, over time, whereas the present paper's random-effects specification compares peak-period consumption across households conditional on total annual consumption and other scale controls, which could in principle net out a reduction that is present but not concentrated during peak hours specifically. Distinguishing between these explanations, and more generally establishing whether the absence of a flat-tariff effect in the present sample reflects a genuine institutional difference between the Australian and US settings studied, or a difference in empirical approach, is not possible with the data available here and is flagged as an open question for future work.

Battery scale and a natural extension

The battery-owning households in this sample installed their systems during an earlier phase of the Australian residential storage market, when typical battery capacities were smaller than those being installed at the time of writing; a subsequent, materially larger-capacity cohort of installations has followed from July 2025 (the data used in this study was in homes that had batteries installed before October 2023). This is a scope condition the paper's own findings suggest a way to test. If the mechanism identified here is a discrete, tariff-triggered

configuration decision rather than continuous price elasticity, which Section 6.1 rejects, or an unconditional structural effect, which Section 6.2 also rejects (since flat-tariff battery households show no effect) a natural question for future work is whether this pattern holds as battery capacity grows. Two competing predictions are both plausible under the same broad interpretation offered here: a larger battery might amplify the same tariff-conditional pattern (a larger reduction, but still concentrated among TOU-exposed households), consistent with more capacity giving the same dispatch logic more room to work with. Alternatively, sufficient capacity might eventually decouple discharge behaviour from tariff status altogether, as self-consumption maximisation dominates regardless of the household's specific price signal, in which case the sharp TOU/flat-tariff distinction documented here would be expected to narrow or disappear in a larger-capacity cohort. Distinguishing between these two possibilities is a direct, falsifiable extension of this paper's design, comparing the present, smaller-capacity cohort against a subsequent larger-capacity cohort using the same empirical approach. Data covering a larger-capacity cohort, drawn from a more recent period, is expected to become available as a natural next step for this research.

8. Conclusion

This paper set out to test two related questions about household response to time-of-use electricity pricing in a setting of growing solar and battery adoption: whether households respond to the size of the price differential a time-of-use tariff embeds, and whether owning solar or battery storage changes peak-period grid consumption independently of price. The answer to the first question is a robust no: across a household fixed-effects specification exploiting genuine within-household price variation, and a further restriction to the households with the greatest such variation, there is no evidence that households of any type (including battery owners) adjust peak-period consumption in proportion to the size of the peak-to-off-peak price differential they face. Of course, this is a statistical result that will not necessarily be true for individual customers.

The answer to the second question is more specific than a simple yes or no: solar-plus-battery households show a significant, robust reduction in peak-period consumption relative to other households on the same tariff, but only when they are actually enrolled in a time-of-use tariff, and not in proportion to the size of that tariff's price differential. It should be stressed however that this conclusion does not reflect more recent development in innovative prices such as those that offer free electricity in the middle of the day. The conclusion will of course not also be true for all customers in the dataset. Battery households on a flat tariff, and solar-only households under either tariff type, show no comparable pattern. Together, these results are inconsistent with continuous price elasticity and with an unconditional, price-independent structural effect of storage, and instead point to a discrete, tariff-triggered mechanism plausibly reflecting battery discharge configured, at installation or through automated or manual adjustment, around a known tariff structure rather than active optimisation against its magnitude.

These findings corroborate Burns and Mountain's (2021) conclusion that Australian households show little genuine price responsiveness to time-of-use tariffs, in a new panel setting that also incorporates battery storage, and extend the wider literature by showing, using revealed rather than simulated household behaviour, that a battery's effect on peak consumption depends on tariff enrolment rather than the price signal's size. The results carry a direct policy implication: TOU tariffs are *not* effective as a general instrument for the broad customer base but *are* effective specifically for the battery-owning sub-population, suggesting value in considering time-of-use tariff rollout and battery incentive policy jointly rather than as independent levers.

The paper's most direct suggested extension follows from its own central finding. Because the battery-owning households studied here installed relatively modest-capacity systems during an early phase of the Australian residential storage market, and a materially larger-capacity cohort has followed more recently, comparing the present results against an equivalent analysis of that larger-capacity cohort offers a natural, falsifiable test of whether the tariff-triggered mechanism identified here strengthens with scale or instead gives way to unconditional self-consumption maximisation once storage capacity is large enough. Secondary directions for future work include decomposing the shoulder- and off-peak-period components of the consumption response alongside the peak-period results reported here; further investigation of the unresolved association between time-of-use exposure and shadow peak consumption among non-solar households; and, should future data capture a meaningful number of households switching tariff type or installing storage mid-panel, a difference-in-differences

design capable of identifying these effects with household fixed effects rather than the cross-sectional comparisons relied upon in this paper.

Appendix

Table A1. Households by state, distributor and tariff type

State	Distributor	Flat	Two-rate	Three-rate	Three-rate w/ Controlled Load	Three-rate with Demand Charge
NSW	Essential	26	0	90	46	0
NSW	Ausgrid	270	26	345	86	48
NSW	Endeavour	571	35	15	4	0
VIC	Citipower	40	19	—	1 (2-rate w/ CL)	0
VIC	Ausnet	74	270	—	0	0
VIC	Powercor	136	172	—	12 (2-rate w/ CL)	0
VIC	Jemena	112	94	—	0	0
VIC	United	222	168	—	0	0
QLD	Energex	1,016	26	101	39	0
SA	SAPN	185	2	1,077	336	0

Only households in the AusGrid distribution area in NSW on three rate TOU carry a demand charge (48, of which 15 also carry controlled load); no other distributor has any.

Table A2–A4. Full regression results for the physical/behavioural specification, all three peak-window definitions

The following three tables report the complete estimation output, including all month and state fixed effects, for the physical/behavioural random-effects specification (equation (2), Section 5.2) under each of the three peak-window definitions compared in Table 6. Table A3 reproduces the coefficients shown in abbreviated form in Table 5 (Section 6.2), with the month and state fixed effects shown in full here rather than summarised in a footnote.

Table A2. Dependent variable: *ihs_daily_peak_shadowA* (full sample, random effects, state FE)

Variable	Coef.	Std. Err.	z	p-value
has_solar	0.2088	0.0257	8.12	0.000 ***
has_solarbattery	0.1157	0.0280	4.13	0.000 ***
has_solar × faces_tou	-0.1982	0.0473	-4.19	0.000 ***
has_solarbattery × faces_tou	-0.2474	0.0429	-5.76	0.000 ***
faces_tou	0.0696	0.0277	2.51	0.012 **
has_controlled_load × faces_tou	-0.0436	0.0562	-0.78	0.438
has_controlled_load	0.0442	0.0286	1.54	0.122
annual_bought	0.0001408	0.0000056	25.16	0.000 ***
max_import_capacity	-0.0010	0.0044	-0.22	0.826
has_ev	-0.1183	0.0452	-2.62	0.009 ***
Feb	0.0785	0.0039	19.91	0.000 ***
Mar	0.0244	0.0050	4.91	0.000 ***
Apr	-0.5550	0.0153	-36.38	0.000 ***
May	-0.4404	0.0168	-26.22	0.000 ***
Jun	0.3972	0.0105	37.88	0.000 ***
Jul	0.3892	0.0105	36.95	0.000 ***
Aug	0.1949	0.0095	20.52	0.000 ***
Sep	-0.5138	0.0160	-32.09	0.000 ***
Oct	-0.6227	0.0146	-42.62	0.000 ***
Nov	-0.1431	0.0054	-26.70	0.000 ***
Dec	-0.0525	0.0040	-13.15	0.000 ***

QLD	0.6495	0.0215	30.21	0.000 ***
SA	1.1092	0.0251	44.13	0.000 ***
VIC	0.6903	0.0198	34.91	0.000 ***
Constant	0.5033	0.0351	14.35	0.000 ***

Robust standard errors clustered by siteid. $N = 3,049,653$; 5,096 households; reference state = NSW. R^2 : within = 0.166, between = 0.572, overall = 0.340. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A3. Dependent variable: ihs_daily_peak_shadowB (full sample, random effects, state FE)

Variable	Coef.	Std. Err.	z	p-value
has solar	0.1978	0.0308	6.41	0.000 ***
has solarbattery	0.0510	0.0331	1.54	0.124
has solar $\times$ faces tou	-0.2179	0.0583	-3.74	0.000 ***
has solarbattery $\times$ faces tou	-0.1904	0.0529	-3.60	0.000 ***
faces tou	0.0882	0.0294	3.00	0.003 ***
has controlled load $\times$ faces tou	-0.0607	0.0727	-0.83	0.404
has controlled load	0.0562	0.0351	1.60	0.110
annual bought	0.0001429	0.0000055	25.75	0.000 ***
max import capacity	-0.0013	0.0044	-0.29	0.773
has ev	-0.1445	0.0451	-3.20	0.001 ***
Feb	0.0577	0.0036	15.84	0.000 ***
Mar	-0.0196	0.0047	-4.13	0.000 ***
Apr	-0.0661	0.0064	-10.37	0.000 ***
May	0.0840	0.0082	10.26	0.000 ***
Jun	0.3266	0.0100	32.54	0.000 ***
Jul	0.3568	0.0100	35.58	0.000 ***
Aug	0.1447	0.0090	16.07	0.000 ***
Sep	-0.0422	0.0083	-5.08	0.000 ***
Oct	-0.1523	0.0063	-24.30	0.000 ***
Nov	-0.1524	0.0050	-30.21	0.000 ***
Dec	-0.0808	0.0036	-22.26	0.000 ***
QLD	0.6233	0.0231	27.02	0.000 ***
SA	1.0827	0.0273	39.62	0.000 ***
VIC	0.6711	0.0208	32.27	0.000 ***
Constant	0.3631	0.0348	10.44	0.000 ***

Robust standard errors clustered by siteid. $N = 3,049,653$; 5,096 households; reference state = NSW. Reproduces Table 5 (Section 6.2) with fixed effects shown in full. R^2 : within = 0.040, between = 0.567, overall = 0.287. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A4. Dependent variable: ihs_daily_peak_user (5pm–9pm window; full sample, random effects, state FE)

Variable	Coef.	Std. Err.	z	p-value
has solar	0.3353	0.0401	8.36	0.000 ***
has solarbattery	0.1485	0.0405	3.67	0.000 ***
has solar $\times$ faces tou	-0.2543	0.0785	-3.24	0.001 ***
has solarbattery $\times$ faces tou	-0.1809	0.0668	-2.71	0.007 ***
faces tou	0.1344	0.0345	3.90	0.000 ***
has controlled load $\times$ faces tou	-0.1048	0.0991	-1.06	0.290
has controlled load	0.0359	0.0460	0.78	0.435
annual bought	0.0001426	0.0000055	26.06	0.000 ***
max import capacity	-0.0040	0.0043	-0.94	0.348
has ev	-0.1566	0.0478	-3.28	0.001 ***
Feb	0.0725	0.0035	20.43	0.000 ***
Mar	0.0172	0.0044	3.88	0.000 ***
Apr	-0.0617	0.0061	-10.07	0.000 ***
May	0.0926	0.0078	11.80	0.000 ***
Jun	0.3531	0.0093	37.80	0.000 ***
Jul	0.3526	0.0094	37.67	0.000 ***
Aug	0.1834	0.0085	21.58	0.000 ***

Sep	-0.0187	0.0079	-2.36	0.018 **
Oct	-0.1395	0.0061	-23.03	0.000 ***
Nov	-0.1328	0.0048	-27.61	0.000 ***
Dec	-0.0579	0.0037	-15.85	0.000 ***
QLD	0.0743	0.0260	2.86	0.004 ***
SA	-0.2122	0.0311	-6.83	0.000 ***
VIC	0.0368	0.0228	1.61	0.108
Constant	0.6960	0.0368	18.90	0.000 ***

*Robust standard errors clustered by siteid. N = 3,049,653; 5,096 households; reference state = NSW. R²: within = 0.061, between = 0.527, overall = 0.302. *p<0.10, **p<0.05, ***p<0.01.*

References

- Bellemare, M.F. & Wichman, C.J. (2020). Elasticities and the inverse hyperbolic sine transformation. *Oxford Bulletin of Economics and Statistics*, 82(1), 50–61.
- Burns, K. & Mountain, B. (2021). Do households respond to Time-Of-Use tariffs? Evidence from Australia. *Energy Economics*, 95, 105070.
- Cappers, P.A. & Todd-Blick, A. (2021). Heterogeneity in own-price residential customer demand elasticities for electricity under time-of-use rates: Evidence from a randomized-control trial in the United States. *Utilities Policy*, 73, 101314.
- Caves, D.W. & Christensen, L.R. (1980). Econometric analysis of residential time-of-use electricity pricing experiments. *Journal of Econometrics*, 14(3), 287–306.
- Choo, D., Jang, H. & Kim, J. (2025). Strategic and price-dependent responses to opt-in time-of-use tariffs: Evidence from Jeju Island, Korea. *Energy Policy*, 206, 114724.
- Filippini, M. (1995). Swiss residential demand for electricity by time-of-use: An application of the Almost Ideal Demand System. *The Energy Journal*, 16(1), 27–40.
- Herriges, J.A., Caves, D.W. & Kuester, K.A. (1989). Load shifting under voluntary residential time-of-use rates. *The Energy Journal*, 10(4), 83–99.
- Ito, K. (2014). Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. *American Economic Review*, 104(2), 537–563.
- Qiu, Y., Xing, B., Patwardhan, A., Hultman, N. & Zhang, H. (2022). Heterogeneous changes in electricity consumption patterns of residential distributed solar consumers due to battery storage adoption. *iScience*, 25, 104352.
- Restel, L. & Say, K. (2025). Counteracting the duck curve: Prosumage with time-varying import and export electricity tariffs. *Energy Policy*, 198, 114461.
- Sharma, V., Aziz, S.M., Haque, M.H. & Kauschke, T. (2022). Energy economy of households with photovoltaic system and battery storage under time of use tariff with demand charge. *IEEE Access*, 10, 33069–33082.