

# Online Appendix

## Start-up Costs and Market Power: Lessons from the Renewable Energy Transition

by Akshaya Jha and Gordon Leslie

<b>A</b>	<b>Data appendix</b>	<b>3</b>
<b>B</b>	<b>Capabilities, operating patterns, and spell profits of coal generating units</b>	<b>5</b>
B.1	Trends in the operations of coal, CCGT, and OCGT units . . . . .	5
B.2	Cost function for coal-fired units . . . . .	12
B.3	Fixed costs and the profitability of operating spells across generating units . . . . .	14
<b>C</b>	<b>Descriptive trends in market outcomes: Additional figures and tables</b>	<b>19</b>
C.1	Trends over time . . . . .	19
C.2	Descriptive trends aggregated across coal and gas units . . . . .	21
<b>D</b>	<b>Dynamic production function: Additional detail on methods and estimates</b>	<b>24</b>
<b>E</b>	<b>Additional details on constraints in our least-cost allocation problem</b>	<b>30</b>
E.1	Minimum and maximum feasible output levels and maximum ramping constraints . . . . .	30
E.2	Capacity adjustments and must-run status for units with ancillary service obligations . . . . .	32
E.3	Transmission congestion in WA's wholesale market . . . . .	36
E.4	Deadweight loss estimates under static and dynamic frameworks . . . .	38
<b>F</b>	<b>Linking market outcomes to solar penetration: Additional results</b>	<b>42</b>

F.1	Market outcomes aggregated across the coal- and gas-fired fleet . . .	42
F.2	Robustness to different sets of fixed effects . . . . .	44
F.3	Trimming the top 1% and bottom 1% of the outcome . . . . .	46
F.4	Predicted versus observed outcomes . . . . .	46
F.5	Effects of solar by season . . . . .	48
F.6	Estimation excluding lagged wholesale demand . . . . .	52
<b>G</b>	<b>Firm-level evidence linking rooftop solar penetration and competi- tion: Additional figures</b>	<b>54</b>
G.1	Inverse semi-elasticities and rooftop solar penetration: Descriptive trends . . . . .	54
G.2	Inverse semi-elasticities and market outcomes . . . . .	56
<b>H</b>	<b>Competitive impacts from large-scale solar penetration</b>	<b>59</b>
H.1	Methodology . . . . .	59
H.2	Results . . . . .	60
<b>I</b>	<b>The impacts of rooftop solar penetration on carbon emissions</b>	<b>64</b>

## A Data appendix

This section provides details on the data collection and construction process. Summarizing, our analysis uses publicly-available data on daily fuel usage by gas-fired electricity generation units, half-hourly electricity output from each unit, generation unit characteristics, wholesale electricity market outcomes, fuel prices, and rooftop solar installations and performance.

We obtained fuel consumption data for gas-fired units from August 1st 2013 through December 31st 2018 from the Gas Bulletin Board at the Australian Energy Market Operator (GGB-WA (2013-2018)). This database lists the daily total amount of gas consumed from each “feeder”. Feeders are matched to generating units in order to estimate each unit’s production function including the fuel used for starts; see Appendix D for more details.

Wholesale electricity market (WEM) data were also obtained from the Australian Energy Market Operator (AEMO (2014-2018)). These data are available for each 30 minute interval and include information on day-ahead and real-time prices, load, and the output produced by each generating unit. The supply curves submitted to the day-ahead and real-time markets for each half-hour are available either at the portfolio level or the generation unit level.

Unit characteristics were collected from annual reports compiled for the WEM market operator. These reports, such as SKM-MMA (2014), include engineering estimates of the heat rates at minimum and maximum output for each unit but do not include estimates of start-up costs. We document in Appendix D that the

heat rate estimates listed in these reports are similar to our econometric estimates for gas-fired units. We also use the engineering estimates of heat rates from SKM-MMA (2014) to calculate the fuel used by coal-fired units. Finally, each unit's fuel cost in each half-hour is calculated by multiplying fuel use by the relevant quarterly fuel price reported by the Western Australian Department of Mines and Petroleum (WA-DMP (2014-2018)).<sup>39</sup>

The Australian Photovoltaic Institute provides half-hourly data on the capacity of rooftop solar installed in Western Australia as well as an estimate of the aggregate output from rooftop solar panels in each half-hour-of-sample (APVI (2015-2018)). We validate this data source using data on rooftop solar installations from the Australian Government Clean Energy Regulator (AGCER (2012-2018)).

---

<sup>39</sup>The fuel prices listed in WA-DMP (2014-2018) are the prices paid to producers by all gas customers. Consequently, we scale the gas price time series by  $\frac{4.32}{4.96}$  to align this series with the prices paid by gas-fired generation units provided by AEMO (2017) for the year 2017.

## **B Capabilities, operating patterns, and spell profits of coal generating units**

This section is split into three parts. First, we present the operating patterns of coal and gas units, documenting that coal unit starts, stops, and propensity to run are largely unchanged coincident with the rooftop solar boom from 2014 to 2018. The next subsection discusses how we calculate operating costs for coal units inclusive of variable, start-up, and running costs; we note existing academic and regulatory evidence that ramping costs are not economically meaningful for coal units. Finally, we present the operating spell-level distribution of the profits earned by coal and gas units. In doing so, we provide evidence that redispatching coal units as part of our dynamic competitive benchmark would likely lead to lower daily total production costs and thus daily average competitive prices relative to our current approach of fixing coal output at observed levels.

### **B.1 Trends in the operations of coal, CCGT, and OCGT units**

**Starts, stops and propensity to run by year:** Figure B.1 documents the hourly average proportion of units running, proportion of units starting, and proportion of units stopping separately for 2014, 2015-2016, and 2017-2018 for coal, CCGT, and OCGT units. We focus on proportions rather than the number of units so that the y-axis is comparable across unit types (i.e., there are more OCGT units than either

coal or CCGT units).

Coincident with the rooftop solar boom from 2014-2018, Figure B.1a suggests a very slight increase in the proportion of coal units running. Importantly, however, this average increase does not vary across hours of the day in any year, as would be expected if coal unit starts or stops were a function of the level of rooftop solar output. Instead, the descriptive trends suggest that the increases in rooftop solar output during the daytime hours from 2014 to 2018 served primarily to reduce the proportion of CCGT units running between 5am-10pm in favor of OCGT units with lower start-up costs but higher variable costs.

Before interpreting the trends over time in starts and stops, note that the scale of the y-axis for the proportions of coal units starting and stopping is far smaller than the y-axis scale for the proportions of CCGT units starting and stopping. Put more simply, coal units start and stop far less frequently than gas units. In aggregate, there are 295 operating spells by coal-fired generating units from 2014-2018 (with a median length of 441 hours) compared to 9,830 operating spells by gas-fired generating units from 2014-2018 (with a median spell length of 5.5 hours for OCGT units and 17.5 hours for CCGT units).

From panels (d) and (g) of Figure B.1, we see no meaningful trends in the proportions of coal units starting or stopping across hours of the day.<sup>40</sup> Even more

---

<sup>40</sup>The slightly higher rate of coal units stopping at the end of the day and on Fridays (as we document below) reflects that—if a plant is to be shut down for annual maintenance, they will likely choose to shut down for maintenance over the weekend when average demand levels are lower. The key point is that coal units during our sample period do not change the frequency with which they start or stop as a function of rooftop solar output or intra-day variation in prices or demand.

importantly, we see no differences in trends in starts and stops by coal units across 2014 versus 2015-2016 versus 2017-2018, suggesting that the rooftop solar boom from 2014-2018 did not impact coal starts or stops.

The within-day timing of the starts and stops of CCGT units (panels e and h) demonstrate that—before 2017—these units frequently had spells that started just before the morning demand peak and stopped just after the evening demand peak. After 2017, we see that within-day spells covering both the morning and evening peaks become quite rare for CCGT units. This shortfall in CCGT output at the morning and evening peaks is replaced by OCGT units characterized by higher variable costs but lower start-up costs than CCGT units. Namely, panels c, f, and i document slight increases in proportions of OCGT units running as well as stopping and starting at the morning and evening demand peaks over time.

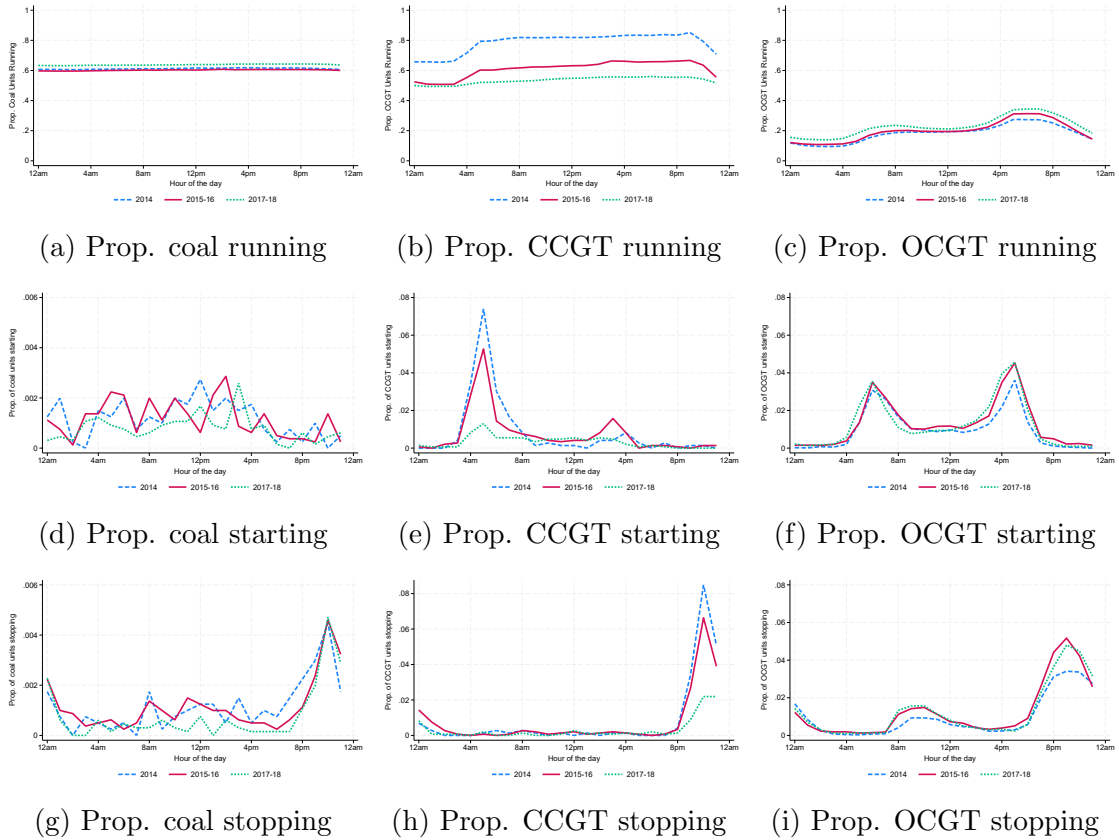
In sum, this suggests that the solar-driven changes in unit starts, stops, and propensity to run are overwhelmingly driven by changes in the operations of CCGT units—with OCGT units starting and stopping to fill in the shortfall in output during morning and evening demand peak hours.<sup>41</sup> In contrast, coal units did not noticeably change their starts, stops, or propensity to run coincident with rooftop solar boom from 2014 to 2018.

**Patterns in capacity running at minimum output:** We next discuss the proportion of generating capacity producing at minimum safe operating levels. Defining

---

<sup>41</sup>This is confirmed by Figure 9, which documents that hourly average total output from Synergy’s CCGT and OCGT units fell from 2014 to 2018, replaced by increases in output from 2014 to 2018 from competitors’ gas units.

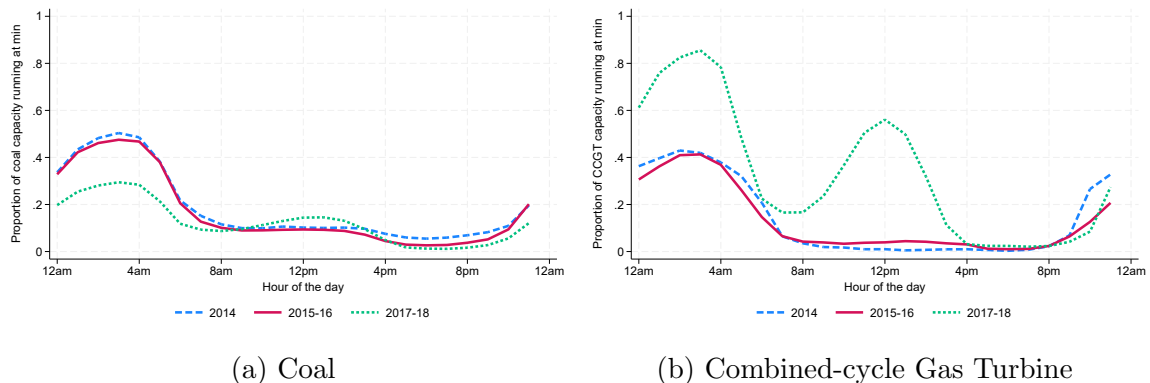
Figure B.1: Hourly average proportions of units running, starts, and stops by year



**Notes:** This figure plots hour-of-the-day specific averages of the proportion of units running in the hour (top row), hourly total number of starts divided by number of units (middle row), and hourly total number of stops divided by the number of units (bottom row). These averages are constructed separately for 2014, 2015-2016, and 2017-2018 for coal, combined-cycle gas turbine (CCGT), and open-cycle gas turbine (OCGT) units. Note that the scale of the y-axis for the proportions of coal units starting and stopping is far smaller than the y-axis scale for the proportions of CCGT and OCGT units starting and stopping.

$\underline{K}_i$  and  $\overline{K}_i$  as unit  $i$ 's minimum safe operating level and maximum production capacity respectively, we say that unit  $i$  is “running at minimum” if its output level is between  $\underline{K}_i$  and  $\underline{K}_i + 0.25 \times (\overline{K}_i - \underline{K}_i)$ . The qualitative conclusions drawn are similar for cut-offs other than 0.25.

Figure B.2: Hourly average proportion of capacity running at minimum by year



**Notes:** This figure plots hour-of-the-day specific averages of the hourly total generating capacity running at minimum divided by hourly total capacity running at any output level. We do so separately for the years 2014, 2015-2016, and 2017-2018 for coal units (left panel) and combined-cycle gas turbine (CCGT) units (right panel). Defining  $\underline{K}_i$  and  $\overline{K}_i$  as unit  $i$ 's minimum safe operating level and maximum production capacity respectively, unit  $i$  is “running at minimum” if its output level is between  $\underline{K}_i$  and  $\underline{K}_i + 0.25 \times (\overline{K}_i - \underline{K}_i)$ .

Figure B.2 plots hour-of-the-day averages of the proportion of capacity running at minimum: hourly total generating capacity running at minimum divided by hourly total capacity running at any output level. We plot separate hour-of-day averages for the years 2014, 2015-2016, and 2017-2018. Before considering the year-to-year differences, note that both coal and CCGT units are more likely to be running at minimum overnight relative to other times of day. This is due to the low levels of demand on average between 12am-6am. These trends highlight the importance of accounting for dynamics. Rather than operating relatively few units at higher output

levels overnight, suppliers produce using relatively more units producing at close to minimum safe operating levels in order to avoid shutting units down and thus having to incur start-up costs in subsequent hours.

Figure B.2b documents that CCGT units are more likely to be running at minimum during the day coincident with the rooftop solar boom. This is intuitive: during the day, output from rooftop solar panels displaces output from CCGT units, leading to a higher propensity to run at minimum. In contrast, Figure B.2a documents only at most slight changes over time in the propensity of coal units to run at minimum outside of overnight hours (i.e., 12am-6am).

Figure B.2b also indicates that the proportion of CCGT capacity running at minimum between 12am-6am increases coincident with the rooftop solar boom. This is because—as documented in Figure B.1b—CCGT units are less likely to be running at all. Namely, Figure 9a documents that Synergy’s CCGT unit—characterized by higher start-up costs but lower variable costs than other gas units—operates less overnight in 2018 than 2014.<sup>42</sup> This is partly because solar-induced decreases in demand during the day make it less likely that this unit can recover its start-up costs by operating throughout the day. However, conditional on running, Synergy’s CCGT unit is more likely to be operating above minimum than other units due to its low variable costs. Therefore, coincident with the rooftop solar boom, the proportion of CCGT capacity running at minimum across all operating CCGT units has increased overnight since the denominator—hourly total CCGT capacity running—has decreased, while the numerator—hourly total CCGT capacity running

---

<sup>42</sup>Table D.1 documents fuel start-up, running, and variable costs for the gas units in our sample.

at minimum—has remained similar.

Finally, Figure B.2a indicates that the proportion of coal generating capacity running at minimum decreases over time. This is because the output reduction from CCGT units between 12am-6am is replaced by increases in output from coal. It is important to emphasize that—despite this increase in output among operating coal units—coal units are unlikely to be setting market prices. Even overnight, market prices are likely set by CCGT units with higher average variable costs and with a shorter operating spell length—requiring higher average prices over the spell for CCGT units to recover their fixed and variable costs. We elaborate on this point in Section B.3.

Summarizing, the trends in coal units running at minimum overnight presented in this section make clear the importance of holding fixed aggregate coal output at *observed* levels when redispatching gas units within our framework to calculate competitive benchmark output levels and prices. Redispatching coal units without fully accounting for their minimum safe operating levels, minimum start times, maximum ramp rates, and other operating constraints would allow coal units to increase output to potentially infeasible levels—resulting in lower daily total production costs and thus underestimates of daily average competitive benchmark prices. By subtracting observed aggregate output from coal units from the net demand to be served by the gas-fired fleet, we ensure that all relevant operating constraints for coal units are respected and remain conservative when setting competitive benchmark prices.

The trends in propensity to run at minimum presented in this section also under-

score the importance of including running costs as part of unit-level cost functions. Incorporating running costs—the costs of producing positive output that are not tied to the level of output—allows us to capture a key feature of power plant production technology: operating costs per MWh are higher at lower output levels. The evidence in this section suggests that accounting for running costs will become increasingly important as production from renewables displaces fossil-fuel output.

## B.2 Cost function for coal-fired units

Unlike the gas units in our sample, we do not directly observe fuel use by coal-fired generators. We extend the approach of Borenstein et al. (2002) to estimate coal use from engineering estimates of heat rates. That is, we calculate the fuel use associated with a 1 MWh increase in output (i.e.,  $\alpha_i^V$ ) and the fuel use associated with running not tied to output levels (i.e.,  $\alpha_i^R$ ) using reported heat rates at minimum safe operating level and maximum capacity from SKM-MMA (2014).<sup>43</sup> With this in hand, total coal use  $F_{i,t}$  is calculated as:

$$F_{i,t} = \alpha_i^V O_{i,t} + \alpha_i^R 1[O_{i,t} > 0] \quad (\text{B.1})$$

We calculate the operating costs  $TC_{i,t}$  incurred by coal-fired unit  $i$  that produces

---

<sup>43</sup>Define minimum safe operating level and maximum operating capacity as  $\underline{K}$  and  $\overline{K}$ , and the heat rate at minimum and maximum as  $HR_{min}$  and  $HR_{max}$ . Then a linear transformation is used to calculate the fuel use from a 1 MWh increase in output,  $\alpha^V = (HR_{max} * \overline{K} - HR_{min} * \underline{K}) / (\overline{K} - \underline{K})$ , and the fuel use conditional on producing non-zero output,  $\alpha^R = HR_{min} * \underline{K} - \alpha^V * \underline{K}$ .

output  $O_{i,t}$  in half-hour-of-sample  $t$  as:

$$TC_{i,t} = \underbrace{P_t^{\text{Coal}} F_{i,t}}_{\text{Fuel Cost}} + \underbrace{\text{VOM}_i O_{i,t}}_{\text{Nonfuel Variable Cost}} + \underbrace{\text{SUC}_i 1[O_{i,t} > 0, O_{i,t-1} = 0]}_{\text{Start-up Cost}}$$

where  $P_t^{\text{Coal}}$  is the price of coal in AUD per GJ and  $F_{i,t}$  is the coal use (in GJ) predicted by Equation (B.1).  $\text{VOM}_i$  is unit  $i$ 's non-fuel variable operating and maintenance costs (in AUD per MWh). Appendix A provides further information on the data used for coal prices and  $\text{VOM}_i$ . Remaining conservative, we apply start-up cost  $\text{SUC}_i$  equal to the highest start-up cost estimate for coal units estimated in Reguant (2014).<sup>44</sup> Since she estimates  $\text{SUC}_i$  based on suppliers' observed bidding behavior, both fuel and nonfuel components of start-up costs are reflected in her estimates.

Our findings are not at all sensitive to the assumptions regarding the start-up or ramping costs of coal units because these costs make up only 1% of the total operating costs incurred by coal units over their operating spells. Note that Reguant's preferred specification does not include ramping costs since her estimates of the ramping costs incurred by coal units are economically small (at most 0.1% of variable costs) and not statistically different from zero.<sup>45,46</sup>

---

<sup>44</sup>Reguant (2014) has a maximum per MW start-up cost of 109.72 Euros (177.75 AUD) across her reported model specifications.

<sup>45</sup>To calculate the dynamic costs incurred by coal units, we utilize the estimates from specification (4) in Table 3 of Reguant (2014). This specification sets ramping costs equal to zero and has the greatest start-up cost per MW of capacity. The quantitative and qualitative implications of our analysis remain the same if we instead consider estimates from specification (5) in Table 3, which includes both start and ramping costs.

<sup>46</sup>The largest point estimate of ramping costs in Reguant (2014) is in specification (5) in Table 3 of her paper. We input the maximum ramp rate observed across coal units in our data (30MW/half-hour; or 60MW/hour), which if applied linearly is 30MWh of energy over an hour. The ramping costs under Reguant's model estimates from increasing output from 120MWh to 150MWh over an hour is only 0.1% of the total variable costs of producing 150 MWh..

### **B.3 Fixed costs and the profitability of operating spells across generating units**

In this section, we calculate and report the distribution of profits earned over each operating spell (i.e., from start to stop) for coal and gas units. For coal units, we document that operating profits are far greater than zero for the vast majority of spells: (1) even after incorporating start-up and ramping cost estimates from Reguant (2014), and (2) regardless of whether operating revenues are constructed based on observed wholesale prices or our dynamic competitive benchmark prices. In contrast, many spells for combined-cycle gas units and especially open-cycle gas units have operating profits close to zero. This provides evidence that no additional constraints are necessary in the price-setting component of our dynamic competitive benchmark to ensure cost recovery for coal units. Namely, our competitive benchmark prices set in order to ensure that gas units recover their fixed and variable costs are high enough that coal units producing at observed output levels would earn substantial profits at these prices—these profits are far more than sufficient to ensure that coal units would recover their fixed and variable costs as well.

Figure B.3 documents the cumulative distribution function (CDF) of spell-level profits when calculating operating revenues using observed wholesale prices; we document separate CDFs for coal, combined-cycle gas turbine (CCGT), and open-cycle gas turbine (OCGT) units. Focusing on the left panel, we see that the distribution of spell-level profits is shifted far to the right for coal units relative to other types of units. Put another way, spells by coal units typically earn substantially larger

operating profits than spells by gas units. Indeed, the operating profits earned by coal units across most spells are large enough to dwarf any reasonable estimate of the start-up costs incurred by coal units, noting again that we currently calculate operating profits utilizing start-up cost estimates from Reguant (2014).

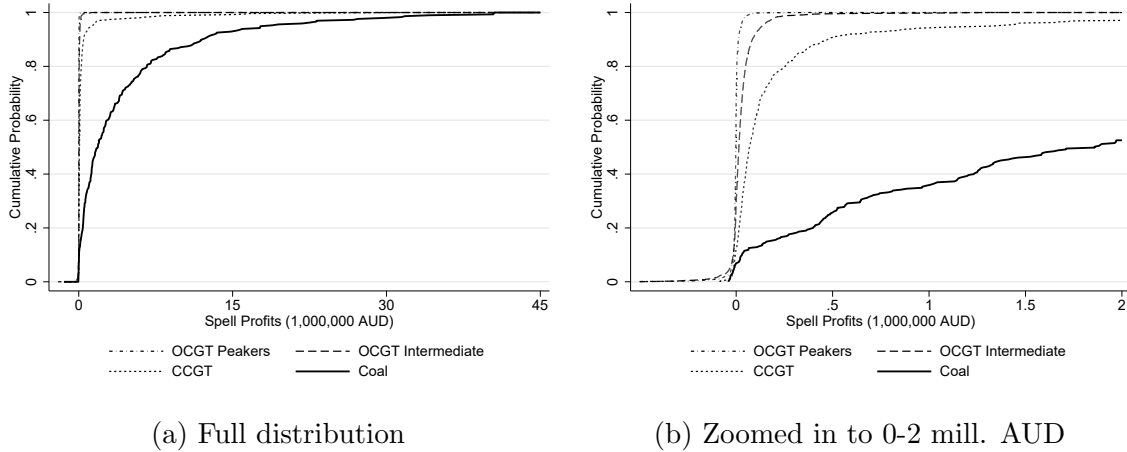
Since spell-level operating profits are so much larger for coal, we zoom in to 0-2 million AUD in profits in the right panel of Figure B.3. Again, we see that the distribution of spell profits is shifted to the right for coal relative to the other types of units. Indeed, the ordering of the empirical distributions of spell profits by unit type make clear that suppliers choose output levels such that units with higher start-up costs but lower variable costs (e.g., coal) earn more per operating spell than units with lower start-up costs but higher variable costs (e.g., peaker OCGT units).<sup>47</sup> Much of the distribution of spell profits for CCGT and OCGT units is close to zero, providing evidence that our estimates of operating costs for these units are not substantially under-estimated. If, for example, we set nonfuel start-up costs for OCGT units to be much larger, this would result in an unreasonably large mass of OCGT unit spells with negative spell profits.

A key constraint in our dynamic framework is that competitive prices across the day must be sufficient for all units to recover their fixed and variable costs. Figure B.4 indicates that the “competitive rents” earned by coal units—that is, the profits they would earn if operating revenues were calculated using our competitive price series—are substantial, far more than sufficient for coal units to recover their fixed and variable costs for the vast majority of their operating spells.

---

<sup>47</sup>We split OCGT units into peakers and intermediate classes to highlight this feature of the data.

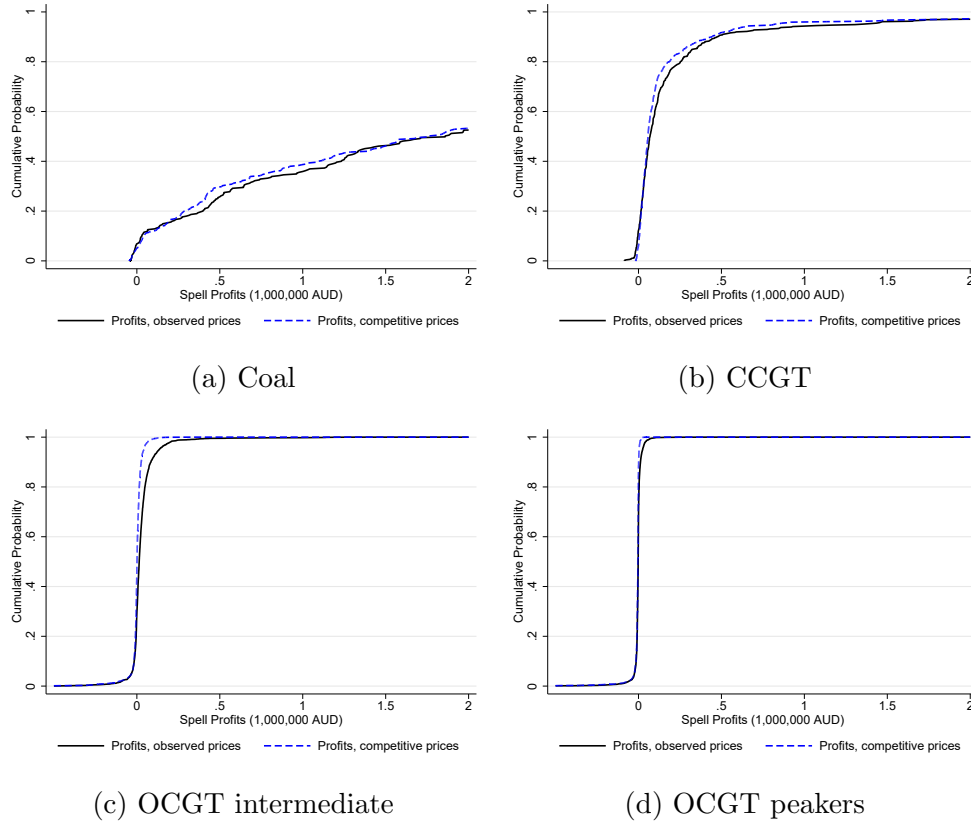
Figure B.3: Spell-level operating profits by unit type



**Notes:** This figure plots the distribution of operating profits earned across each unit’s “operating spell”: the profits earned from the unit starting up to shutting down. We plot this distribution separately for four types of units: coal, combined-cycle gas turbine (CCGT), open-cycle gas turbine (OCGT) intermediate, and OCGT peakers. The left panel of this figure plots the full distribution of spell-level profits while the right panel zooms in to between 0-2 million AUD in profits. The variable component of operating costs for coal units is calculated using a combination of engineering heat rate estimates, coal prices, and nonfuel variable operating and maintenance costs; we also include start-up costs based on the start-up cost estimate from specification (4) in Table 3 of Reguant (2014). Operating costs for gas units are calculated as described in Section A.

**Category definitions:** Gas units are categorized first by their technology: combined-cycle gas turbine (CCGT) versus open-cycle gas turbine (OCGT). Within OCGT, we categorize units into “intermediate” versus “peaker” categories based on heat rates at maximum output levels. The ALINTA WGP, KWINANA and N’GEN N’BUP open-cycle gas turbine (OCGT) units are in the intermediate group (i.e., heat rates between 9.4-11.1 GJ/MWh), and the remaining OCGT peaker units are in the high heat rate group (i.e., heat rates greater than 12.1 GJ/MWh)

Figure B.4: Spell-level operating profits and competitive rents



**Notes:** This figure plots the distribution of operating profits earned across each unit’s “operating spell”: the profits earned from the unit starting up to shutting down. We plot this distribution separately for four types of units: coal (top left panel), combined-cycle gas turbine (top right panel), open-cycle gas turbine (OCGT) intermediate (bottom left panel), and OCGT peakers (bottom right panel). For all four panels, we zoom in to between 0-2 million AUD in profits rather than plot the full distribution of profits. We plot the distribution of profits calculated based on observed wholesale prices (solid black line) and the competitive benchmark prices from our framework (termed competitive rents, dashed blue line). The variable component of operating costs for coal units is calculated using a combination of engineering heat rate estimates, coal prices, and nonfuel variable operating and maintenance costs; we also include start-up costs based on the start-up cost estimate from specification (4) in Table 3 of Reguant (2014). Operating costs for gas units are calculated as described in Section A.

**Category definitions:** Gas units are categorized first by their technology: combined-cycle gas turbine (CCGT) versus open-cycle gas turbine (OCGT). Within OCGT, we categorize units into “intermediate” versus “peaker” categories based on heat rates at maximum output levels. The ALINTA WGP, KWINANA and N’GEN N’BUP open-cycle gas turbine (OCGT) units are in the intermediate group (i.e., heat rates between 9.4-11.1 GJ/MWh), and the remaining OCGT peaker units are in the high heat rate group (i.e., heat rates greater than 12.1 GJ/MWh)

Summarizing, fixing coal units' output at observed levels in the least-cost output allocation component of our dynamic competitive benchmark is conservative, since allowing the redispach of coal units would result in increases in coal unit output and thus decreases in daily total costs and daily average competitive prices. At observed output levels, coal units would earn substantial profits under our dynamic competitive price series. This implies that imposing an additional constraint in the price-setting component of the dynamic framework that prices must be sufficient for coal units to earn revenues in excess of their fixed and variable costs would have very little effect on our dynamic competitive benchmark prices.

## C Descriptive trends in market outcomes: Additional figures and tables

### C.1 Trends over time


Figures C.1a-C.1d present trends in real-time wholesale prices, aggregate heat rate for the gas-fired fleet, and fuel prices over 2015-2018. Each series is relatively flat, with only slight upward trends in electricity prices, heat rates and coal prices over time. Gas prices also remained relatively stable with the exception of a 15% reduction in 2017.

Installed rooftop solar capacity increased by 135% from 2015 to 2018 (see Figure C.2). Coincident with this solar boom, wholesale electricity prices in the day-ahead market increased by 17% while real-time prices increased by 6%.<sup>48</sup> However, the annual aggregate heat rate for the gas-fired fleet only changed by less than 1% from 2015 to 2018. The increase in wholesale prices also cannot be explained by changes in fuel prices: natural gas prices decreased by 16% and coal prices decreased by 3% from 2015 to 2018. This motivates an analysis of start-up costs and market power: rooftop solar production displaces fossil fuel units during the day, forcing these units to start up in order to effectively compete at sunset. Both increases in the competitive benchmark prices necessary to recover start-up, variable, and running costs and increases in the exercise of market power can result in the increases in

---


<sup>48</sup>Day-ahead prices increased from 40.84 AUD/MWh in 2015 to 47.90 AUD/MWh in 2018 while real-time prices increased from 47.39 AUD/MWh to 50.20 AUD/MWh. Both sets of prices are load-weighted annual averages for the years 2015 versus 2018.

Figure C.1: Electricity prices, heat rate and fuel prices: 2015-2018



(a) Load-Weighted Average Real-Time Electricity Price (AUD per MWh)

(b) Aggregate Heat Rate for the Gas-Fired Fleet (GJ/MWh)



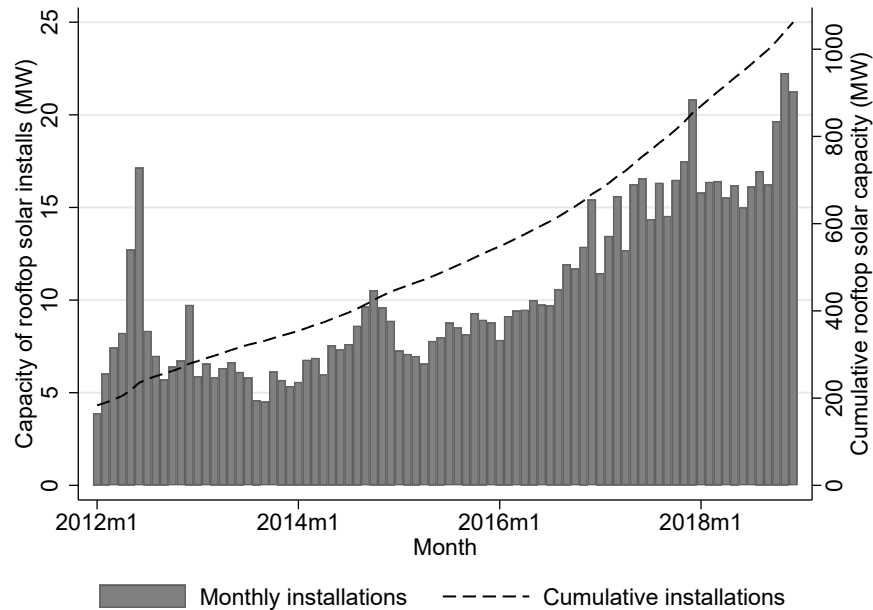
(c) Natural Gas Price (AUD per GJ)

(d) Coal Price (AUD per MT)

**Notes:** Panel (a) of this figure plots the monthly load-weighted average price in the balancing market (i.e., the “real-time” price); note that wholesale electricity demand is called “load”. Panel (b) of this figure plots the monthly total fuel consumed by the gas-fired fleet (in GJ) divided by the monthly total output produced by the gas-fired fleet (in MWh). Panels (c) and (d) report quarterly average natural gas prices and coal prices in AUD per GJ and AUD per MT respectively. The sample period considered for all panels is 2015-2018.

wholesale prices observed at sunset.

Figure C.2: Trends in the installation of rooftop solar capacity



**Notes:** This figure displays monthly total capacity of new installations of rooftop solar panels (the histogram) and the total cumulative amount of rooftop solar capacity installed up to that month-of-sample (the dashed line).

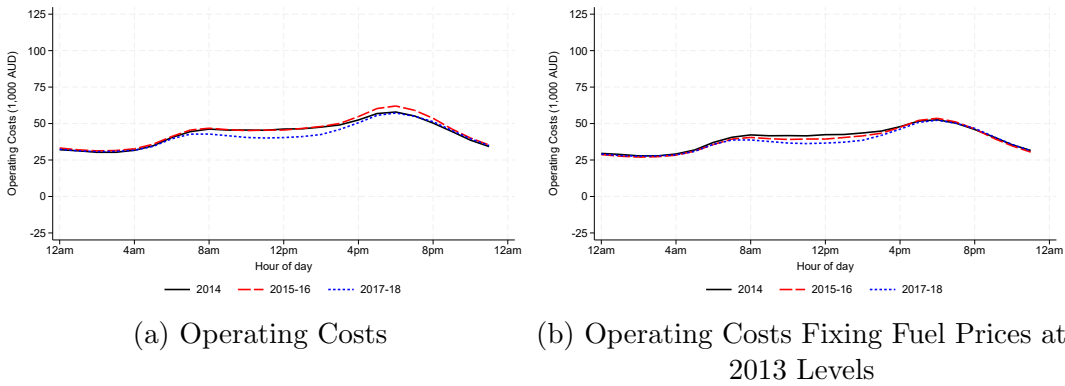
## C.2 Descriptive trends aggregated across coal and gas units

Figure C.3 plots hour-of-day specific averages of hourly aggregate operating costs across both coal and gas units for the year 2014, the years 2015-2016, and the years 2017-2018.<sup>49</sup> The trends in Figure C.3 are quite similar to those presented in Figure 3 focusing on the aggregate costs incurred only by the gas fleet.

<sup>49</sup>Note that a carbon tax was repealed on July 1 2014 (Leslie, 2018). For the purpose of displaying changes in hourly average operating costs over time, the descriptive figures in this section exclude the carbon tax costs incurred by coal- and gas-fired units. All other analyses include carbon tax costs. Namely, the carbon tax is accounted for when calculating dynamic competitive prices and output levels, when presenting descriptive trends in rents, and in the regression analyses that assess how small increases in rooftop solar penetration impact market outcomes.

Namely, after fixing fuel prices at 2013 levels, the aggregate operating costs incurred by the coal and gas fleet are lower during the day in 2018 relative to 2014, consistent with output from rooftop solar increasingly displacing production from fossil-fuel units. However, solar resources do not produce in the evening. For this reason, hourly average operating costs in the evening remain largely unchanged from 2014 to 2018 after holding fuel prices fixed.

Figure C.3: Average operating costs across coal and gas units by hour of the day: 2014-2018




**Notes:** This figure plots hour-of-day specific averages of hourly aggregate operating costs across both coal and gas units. The left panel focuses on aggregate operating costs while the right panel focuses on aggregate operating costs calculated after fixing fuel prices at 2013 levels. This figure excludes the costs associated with the carbon tax repealed on July 1 2014 (Leslie, 2018).

Figure C.4 presents hour-of-the-day specific averages of hourly aggregate operating profits, competitive rents, and market power rents for the year 2014, the years 2015-2016, and the years 2017-2018. In contrast to Figure 5 in the main text, we aggregate across both coal and gas units when calculating the aggregate magnitudes in Figure C.4. The trends in Figure C.4 are quite similar to those presented in Figure 5 focusing solely on the gas-fired fleet.

Namely, operating profits, competitive rents, and market rents fall during the day coincident with the rooftop solar boom from 2014-2018. However, we see increases in operating profits in the evening from 2014-2018—with these increases in profits driven primarily by increases in market power rents rather than competitive rents. Consistent with our central hypothesis, these trends suggest that the dramatic increase in rooftop solar penetration from 2014-2018 led to increases in the exercise of market power in the evening.

Figure C.4: Hourly average competitive rents and market power rents across coal and gas Units: 2014-2018



**Notes:** This figure plots hour-of-day specific averages of hourly aggregate operating profits, competitive rents, and market power rents across coal-fired and gas-fired units. The left panel reports operating profits: revenues less operating costs incurred in that interval. The middle panel reports competitive rents: the operating profits that would be earned if the wholesale price was set at the competitive price described in Section IV. The right panel reports market power rents: operating profits earned in excess of competitive rents.

## D Dynamic production function: Additional detail on methods and estimates

We assume that each unit’s production function takes the form specified in Equation (1) in Section A. The sample period used to estimate each unit’s production function spans from January 1st 2013 to December 31st 2018. Only days with positive gas use are included when estimating each unit’s production function. We estimate each unit’s production function using ordinary least squares and report heteroskedasticity-consistent standard errors.

**Why not include ramping costs?** We consider a production function similar to the ones used in Wolak (2007) and Reguant (2014). Specifically, both of these papers specify start-up costs as in Equation (1). However, in contrast with these papers, we include the fuel use associated with producing positive output not tied to output levels (i.e., “running fuel”) but not ramping fuel use. We do so for three reasons.

First, our dynamic competitive benchmark includes maximum ramping constraints that we have ensured are relatively conservative, specifically to eliminate the possibility of gas units incurring additional costs such as equipment depreciation from ramping rapidly (i.e., we only allow ramping at “comfortable” rates, not all the way up to the technical limits specified by manufacturers). Second, we follow previous literature in electricity economics that focuses on dynamic costs but still assume ramping costs are zero for gas units (e.g., Mansur (2008); Cullen (2015)). Reguant (2014) estimates an economically small and not statistically significant ramping cost

for both coal and CCGT units in Spain, leading her to exclude ramping costs from her primary specification. Following this literature, we consider maximum ramping constraints to be economically relevant, but not ramping costs.

Third, ramping costs are not typically referenced in wholesale market rules or regulatory documentation. For example, WA’s market operator requires gas generating units to report maximum ramping rates but does not require information on ramping costs. As another example, in a recent regulatory ruling on the legality of Synergy’s bidding behavior in the real-time market, the cost structures of Synergy’s units were characterized by start-up costs and times, ramp rates, as well as capital and variable costs, but none of the documentation referenced ramping costs (Economic Regulation Authority vs. Synergy (2019-2022)).

**Feeders linked to multiple generating units:** For the ALINTA\_WGP, KEMERTON, KWINANA and MUNGARRA facilities, there are multiple generating units at each plant site. For each of these facilities, generating units at the same site have the same production technology and are thus likely to have similar production functions (SKM-MMA (2014)). Therefore, we estimate the following equation for each feeder  $i$  providing gas to  $J$  identical generating units:

$$\sum_{t=1}^{48} G_{i,t,d} = \alpha_i^V \sum_{j=1}^J \sum_{t=1}^{48} O_{i,j,t,d} + \alpha_i^S \sum_{j=1}^J \sum_{t=1}^{48} S_{i,j,t,d} + \alpha_i^R \sum_{j=1}^J \sum_{t=1}^{48} R_{i,j,t,d} + \sum_{j=1}^J \sum_{t=1}^{48} \epsilon_{i,j,t,d}$$

where  $O_{i,j,t,d}$  is unit  $j$ ’s output in half-hour-of-day  $t$  of day  $d$ . As before, start-up term  $S_{i,j,t,d} \equiv 1[O_{i,j,t,d} > 0] \times 1[O_{i,j,t-1,d} = 0]$  is equal to one if and only if unit  $j$

started up in half-hour  $t$ , and the running term  $R_{i,j,t,d} \equiv 1[O_{i,j,t,d} > 0]$  is equal to one if and only if unit  $j$  is running in half-hour  $t$ .

For the PINJAR facility, there are six generating units that use the same technology (PINJAR1-PINJAR7) while a different technology is used for three other generating units (PINJAR9-PINJAR11). For this facility, we estimate the following equation for feeder  $i$  providing gas to  $J$  generating units of type 1 and  $K$  generating units of type 2:

$$\begin{aligned} \sum_{t=1}^{48} G_{i,t,d} = & \alpha_1^V \sum_{j=1}^J \sum_{t=1}^{48} O_{i,j,t,d} + \alpha_1^S \sum_{j=1}^J \sum_{t=1}^{48} S_{i,j,t,d} + \alpha_1^R \sum_{j=1}^J \sum_{t=1}^{48} R_{i,j,t,d} + \sum_{j=1}^J \sum_{t=1}^{48} \epsilon_{i,j,t,d} \\ & + \alpha_2^V \sum_{k=1}^K \sum_{t=1}^{48} O_{i,k,t,d} + \alpha_2^S \sum_{k=1}^K \sum_{t=1}^{48} S_{i,k,t,d} + \alpha_2^R \sum_{k=1}^K \sum_{t=1}^{48} R_{i,k,t,d} + \sum_{k=1}^K \sum_{t=1}^{48} \epsilon_{i,k,t,d} \end{aligned}$$

**Parameter estimates and comparisons to engineering heat rates:** Table D.1 reports each unit's parameter estimates. Note that each facility can be home to multiple units; for example, a facility with 6 units, each of which has 38MW of capacity, would be listed as  $38 \times 6$ . The last three columns of Table D.1 list the parameter estimates corresponding to each unit's production function: variable fuel component ( $\hat{\alpha}_i^V$ ), start-up fuel requirement ( $\hat{\alpha}_i^S$ ), and running fuel requirement ( $\hat{\alpha}_i^R$ ). We report heteroskedasticity-consistent standard errors for each parameter estimate in parentheses.

Table D.2 compares the heat rates at minimum safe operating levels and maximum operating capacities implied by our estimates to the engineering estimates reported in SKM-MMA (2014). Specifically, this report lists each unit's capacity as

well as engineering estimates of each unit’s heat rates when producing at its minimum safe operating level and at capacity. The last two columns of Table D.2 list the heat rates implied by our parameter estimates at minimum and maximum output,<sup>50</sup> with the preceding two columns reporting the corresponding engineering estimates.

Our econometric estimates in Table D.1 demonstrate that gas-fired units must use non-trivial quantities of fuel in order to start up. Further, there is significant heterogeneity in heat rates and start-up costs across gas-fired units. Combined Cycle Gas Turbine (CCGT) units are considered “baseload” because they have low heat rates but require larger amounts of fuel to start up. In contrast, Open Cycle Gas Turbine (CCGT) units are considered “peakers” because they are characterized by higher heat rates but smaller amounts of fuel required to start up.

From Table D.2, we see that the heat rates at maximum output are very similar across econometric and engineering estimates for all units listed in SKM-MMA (2014). Our econometric estimates of heat rates at minimum safe operating levels largely align with the engineering estimates. In the case of Cockburn, the difference may be because the minimum safe operating level used in the engineering report is significantly higher than the minimum safe operating level we calculate from the data.

---

<sup>50</sup>These heat rates are  $\frac{\hat{\alpha}_i^V X + \hat{\alpha}_i^R}{X}$ , where  $X$  either takes the value of the unit’s empirically-defined minimum safe operating level or maximum output.

Table D.1: Econometric estimates of the parameters of each gas-fired unit’s production function

Unit	Owner	Capacity (MW)	HR ( $\hat{\alpha}_i^V$ ) ( $\frac{\text{GJ}}{\text{MWh}}$ )	Start Up ( $\hat{\alpha}_i^S$ ) (GJ)	Running ( $\hat{\alpha}_i^R$ ) (GJ)
<i>Combined Cycle Gas Turbine (CCGT) Units</i>					
COCKBURN	Synergy	237 × 1	6.0 (0.111)	748.6 (27.5)	183.5 (11.5)
N’GEN_K’ANA	NewGen	324 × 1	7.0 (0.228)	3014.8 (623.3)	119.7 (29.0)
<i>Open Cycle Gas Turbine (OCGT) Units</i>					
ALINTA_WGP	Alinta	190 × 2	8.1 (0.05)	780.6 (19.6)	257.9 (3.9)
KEMERTON	Synergy	154 × 2	9.7 (0.169)	520.8 (29.0)	196.7 (10.1)
KWINANA	Synergy	100 × 2	7.8 (0.019)	181.4 (10.0)	79.5 (0.6)
MUNGARRA	Synergy	38 × 3	9.8 (0.457)	195.3 (36.6)	49.0 (5.4)
N’GEN_N’BUP	NewGen	342 × 1	7.9 (1.023)	786.0 (371.4)	295.6 (88.9)
PINJAR1-7	Synergy	38 × 6	12.6 (0.741)	143.4 (14.8)	26.0 (7.9)
PINJAR9-11	Synergy	123 × 3	8.8 (0.051)	485.8 (5.9)	173.3 (1.4)

**Notes:** This table lists the owner of each gas generating unit, the unit’s capacity (in MW), and our parameter estimates from the production function specified in Equation (1). Each facility can be home to multiple units; for example, a facility with 3 units, each of which has 38MW of capacity, would be listed as 38 × 3. Each unit’s capacity is listed in SKM-MMA (2014). We estimate each unit’s production function using the methodology discussed in Section A. Heteroskedasticity-consistent standard errors are reported in parentheses.

Table D.2: Engineering versus econometric estimates of each gas-fired unit’s heat rates at minimum and maximum feasible output levels

Facility	Owner	Capacity (MW)	Engineering Estimates		Econometric Estimates	
			HR@min (GJ/MWh)	HR@max (GJ/MWh)	HR@min (GJ/MWh)	HR@max (GJ/MWh)
<i>Combined Cycle Gas Turbine (CCGT) Units</i>						
COCKBURN	Synergy	237 × 1	9.4	9	13.3	7.4
N’GEN_K’ANA	NewGen	324 × 1	.	.	8.5	7.7
<i>Open Cycle Gas Turbine (OCGT) Units</i>						
ALINTA_WGP	Alinta	190 × 2	16.2	11.5	14.9	11.1
KEMERTON	Synergy	154 × 2	13.3	12.2	15.4	12.1
KWINANA	Synergy	100 × 2	15.2	9.4	13.9	9.4
MUNGARRA	Synergy	38 × 3	21.6	13.2	19.9	13.4
N’GEN_N’BUP	NewGen	342 × 1	.	.	15.7	9.7
PINJAR1-7	Synergy	38 × 6	22.5	13.2	17.9	14.2
PINJAR9-11	Synergy	123 × 3	21.7	12	22.6	12.8

**Notes:** This table lists the owner of each unit, each unit’s capacity (in MW), engineering estimates of heat rates when producing at minimum and maximum feasible output levels (in GJ/MWh), and our econometric estimates of these heat rates from the production function specified in Equation (1). Appendix D details the conversion of the parameters in Table D.1 to heat rates at minimum and maximum feasible output levels. Each facility can be home to multiple units; for example, a facility with 3 units, each of which has 38MW of capacity, would be listed as 38 × 3. Each unit’s capacity and engineering estimates of heat rates at minimum and maximum feasible output levels are listed in SKM-MMA (2014). Engineering estimates for N’GEN\_K’ANA and N’GEN\_N’BUP are not available.

## **E Additional details on constraints in our least-cost allocation problem**

The first subsection discusses how we calculate each gas unit’s minimum safe operating level, production capacity, and maximum ramping constraints. Next, we discuss how we adjust minimum safe operating levels and maximum operating capacities to account for the provision of ancillary services. Then, Section E.3 discusses how the dynamic competitive benchmark accounts for the two brief transmission congestion events during our sample period that affected gas units, along with details on the institutional features surrounding why transmission congestion is exceedingly rare in Western Australia. Finally, Section E.4 documents deadweight loss estimates under four permutations defined by: (1) statically versus dynamically determined output allocations, and (2) production costs calculated only based on variable costs versus start, running, and variable costs.

### **E.1 Minimum and maximum feasible output levels and maximum ramping constraints**

For each unit  $i$ , we calculate minimum safe operating level (i.e.,  $\underline{K}_i$ ), production capacity (i.e.,  $\overline{K}_i$ ), maximum ramping constraints conditional on the unit starting in half-hour  $t$  (i.e.,  $\overline{M}_i(1)$ ), and maximum ramping constraints conditional on the unit operating in half-hour  $t - 1$  (i.e.,  $\overline{M}_i(0)$ ) using half-hourly unit-level data on output. We impose that units of the same generating technology at the same facility

(e.g., PINJAR units 1-7) have the same parameters ( $\underline{K}_i, \overline{K}_i, \overline{M}_i(0), \overline{M}_i(1)$ ). Facility/technology groupings for gas units are listed in Table D.1.

**Minimum safe operating levels:** For each facility/technology group, we calculate  $\underline{K}_i$  as the 2nd percentile of half-hourly unit-level output—excluding observations if the unit started up in half-hours  $t - 1$ ,  $t$ , or  $t + 1$ . This ensures we are not considering observations where output is artificially low because the unit is starting up or ramping up or down.

**Production capacity:** For each facility/technology group, we calculate  $\overline{K}_i$  as the 98th percentile of half-hourly unit-level output—excluding observations if the unit is not operating or is observed to be operating at less than the minimum safe operating level  $\underline{K}_i$  in half-hour  $t$ .

**Maximum ramping constraints:** For each facility/technology group, we calculate the difference between output in half-hours  $t$  and  $t - 1$  (i.e., unit  $i$ 's ramp in half-hour  $t$  is  $O_{i,t} - O_{i,t-1}$ ). Maximum ramping capability if the unit started up in half-hour  $t$  is denoted  $\overline{M}_i(1)$  and maximum ramping capability if the unit did not start up is denoted  $\overline{M}_i(0)$ .

When calculating both  $\overline{M}_i(0)$  and  $\overline{M}_i(1)$ , we exclude observations where the ramping term is negative (i.e., the unit is ramping down rather than ramping up). We calculate  $\overline{M}_i(0)$  as the 98th percentile of ramp-ups—excluding observations where the unit started up in half-hour  $t$  and excluding observations where  $O_{i,t}$  is less than

$0.9 \times \overline{K}_i$ . The latter restriction ensures we are focusing on observations where the unit is plausibly ramping up as much as it can; if output increases by a small amount between half-hours  $t-1$  and  $t$ , this is unlikely to be informative of the unit’s maximum capability to ramp.

We calculate  $\overline{M}_i(1)$  as the 98th percentile of ramp-ups—excluding observations where the unit did *not* start in half-hour  $t$  and excluding observations where  $O_{i,t}$  is less than minimum safe operating levels. The latter restriction again ensures that we are focusing on observations where the unit is plausibly ramping up as much as it can in the half-hour; if the unit starts up but only produces a small amount in half-hour  $t$ , this may correspond to either an aborted start or a start late into the half-hour interval.

## **E.2 Capacity adjustments and must-run status for units with ancillary service obligations**

Most wholesale electricity markets specify either regulatory or market-based instruments to ensure that some units reserve a portion of their capacity to quickly increase or decrease their output over varying short time horizons. In WA, a generator may be paid to be available at the system operator’s discretion to rapidly increase or decrease their output to maintain system frequency and respond to other short-horizon contingencies that require an output adjustment to balance supply and demand.<sup>51</sup> We

---

<sup>51</sup>This is typically implemented via an automatic generation control system employed by the market operator, which allows the operator to rapidly change output levels at multiple generating units to maintain system frequency and ensure supply and demand are balanced.

will generically refer to requirements to increase output as “AS up” and requirements to decrease output as “AS down”. Generating units providing “AS up” must ensure they have sufficient generating capacity to quickly increase output levels if obligated to do so as part of their ancillary services requirements. Generating units providing “AS down” must ensure that their output levels are sufficiently above their minimum safe operating levels, so that they can quickly decrease output levels if required.


Figure E.1 graphically displays how ancillary service requirements affect the production capacity available in the energy market. If unit  $i$  is not providing ancillary services in half-hour  $t$ , it can either produce no output (the dot at zero in the top row of the schematic) or output levels between minimum safe operating level  $\underline{K}_i$  and capacity  $\overline{K}_i$ . In contrast, a unit providing ancillary services must be available to adjust output as needed (i.e., its output cannot be zero), and is restricted to produce within the range  $\underline{K}_{i,t}^{AS}$  and  $\overline{K}_{i,t}^{AS}$  determined by its “AS down” and “AS up” requirements.

To account for the provision of ancillary services in our dynamic competitive benchmark, we first identify the gas units that offer ancillary services in each half-hour. We then tighten the constraints on their feasible ranges of output based on the quantities of AS up and AS down they provide each half-hour. Tightening these constraints weakly increases the optimized daily total cost from our least-cost output allocation problem. This in turn serves to weakly increase the average of competitive prices across the day.<sup>52</sup>

---

<sup>52</sup>Though daily average competitive prices must weakly increase with tighter constraints on output levels, it is possible for competitive prices in some half-hours to fall given that we solve this problem over a 48-half-hour window.

Figure E.1: Feasible output ranges for generating units providing versus not providing ancillary services



**Notes:** Unit  $i$  in half-hour  $t$  either provides ancillary services ( $AS_{i,t} = 1$ ) or does not ( $AS_{i,t} = 0$ ). If unit  $i$  does not provide ancillary services in half-hour  $t$ , it can either produce nothing (i.e., the dot at zero in the top line) or produce any level of output between its minimum safe operating level  $\underline{K}_i$  and maximum operating capacity  $\overline{K}_i$ . However, if the unit is providing ancillary services, it must be available to increase or decrease output as needed; its output cannot be zero. The unit's minimum safe operating level is adjusted to be  $\underline{K}_{i,t}^{AS} = \underline{K}_i + AS_{i,t}^{DOWN}$  and its maximum possible production level becomes  $\overline{K}_{i,t}^{AS} = \overline{K}_i + AS_{i,t}^{UP}$ .  $AS_{i,t}^{UP}$  ( $AS_{i,t}^{Down}$ ) is the maximum amount by which the unit can be asked to increase (decrease) its output level as part of its ancillary service obligations.

Due to regulatory precedent, one firm, Synergy, submits a firm-level offer curve into the real-time market. This offer curve specifies an amount of energy offered at the minimum price—labeled in the data as being due to the provision of ancillary services. This ensures that Synergy’s generating units collectively are operating at levels sufficient to allow them to decrease output if required to do so as part of their ancillary service obligations (i.e., provide “AS down” services). Synergy’s offer curve also specifies an amount of energy offered at the maximum price—again labeled in the data as being due to the provision of ancillary services. Energy offered at the maximum price ensures that enough generating capacity is withheld to allow units to increase output if required to do so as part of their ancillary service provisions (i.e., provide “AS up” services).

Since Synergy submits a firm-level offer curve, it is not specified which of its units are providing ancillary services on a given day-of-sample. However, we can use unit-level output profiles to ascertain which units are providing ancillary services. For example, if we observe an open-cycle gas-turbine (OCGT) unit producing for all 24 hours of the day, it is almost surely providing ancillary services given that the typical OCGT output profile is to produce only during peak demand hours. Formally, in each day, we focus on units owned by Synergy that operate for at least 20 hours and assign units to the provision of ancillary services until both AS-down and AS-up requirements would be satisfied at the units’ observed levels of output.

After assigning units to provide ancillary services, we fix their output at observed levels for the least-cost allocation problem, or in the framing of Figure E.1, their feasible output range is reduced to a single point. By setting these units’ output

levels to their observed level of output to ensure that ancillary service requirements are met, we are effectively no longer allowing them to be redispatched as part of our dynamic framework. Holding the output of Synergy’s units assigned to provide ancillary services fixed at observed levels rather than allowing these units to be redispatched (with bounds on their output levels) is conservative. Namely, reducing the set of units that can be redispatched can only result in increases in the optimised daily total costs from our least-cost allocation, pushing daily average competitive prices up.

### **E.3 Transmission congestion in WA’s wholesale market**

Being an isolated and islanded grid, Western Australia (WA) has historically had an overbuilt transmission network. Due to this, WA operated during our sample period under a principle they term “firm-access rights”: each dispatchable power station should have the ability to generate at its maximum operating capacity without being curtailed by transmission constraints at any time. Indeed, day-ahead and real-time markets in WA clear at a uniform price that does not vary across locations to reflect transmission congestion. Consistent with WA’s market clearing processes, our dynamic competitive framework sets one uniform price that does not vary across locations to reflect the potential for transmission congestion.

The market regulator provides quarterly reports on transmission congestion events: events where “the configuration of the transmission network ... has an effect or potential effect of constraining or otherwise varying the output of a generator”

(AEMO (2014-2016); ERA (2016-2018)). These reports are often less than a page, highlighting how rare transmission congestion events are in WA. Indeed, there are only two transmission congestion events listed for gas-fired units between 2014-2018, with the stated reason in each case being bushfires threatening other lines in the network. This confirms that transmission congestion is exceedingly rare and that transmission constraints are no more likely to bind in the evening or in a way that is related to solar generation.

Our least-cost output allocation problem, discussed in Section B, accounts for the two transmission congestion events affecting gas units. Briefly, each of the two transmission congestion events reports a single gas unit that the market operator deemed must be online (i.e., producing positive output) as a contingency during brief bushfire threats to parts of the network. During the time period of the congestion event, our competitive benchmark forces the relevant gas unit to produce at observed levels – their production cannot be reallocated to a lower-cost source.

We note that our competitive benchmark allocates output across gas units in a dynamic least-cost manner holding output from non-gas units fixed at observed levels. Holding output from non-gas units fixed at observed levels ensures that any constraints on output imposed by transmission congestion events are respected for these units as well. Consequently, the revised framework respects all of the constraints on output levels imposed by transmission congestion events, ensuring that transmission congestion does not contribute to the gap between market outcomes and the dynamic competitive benchmark in any half-hour-of-sample.

## E.4 Deadweight loss estimates under static and dynamic frameworks

Table E.1 below summarizes the methodological differences across our dynamic competitive benchmark, traditional static formulations (e.g., Borenstein et al. (2002); Joskow and Kahn (2002)), and Mansur (2008). This table illustrates the two ways in which our dynamic competitive benchmark differs from traditional static formulations: (1) output allocation, and (2) calculation of costs. Namely, static benchmarks dispatch units in each half-hour from lowest to highest average variable cost (i.e., a static output allocation method). Static benchmarks also calculate the costs associated with observed and simulated output levels only based on average variable costs, without accounting for start-up and running costs (i.e., a static cost calculation method).

Mansur (2008) improved upon static benchmarks by calculating benchmark output levels using a dynamic regression model estimated on a sample period when output levels were plausibly allocated amongst generating units in a dynamic least-cost manner (i.e., a dynamic output allocation method). However, costs in Mansur (2008) were still calculated only based on average variable costs, omitting start-up and running costs. Finally, our dynamic competitive benchmark both dynamically allocates output levels for each unit across half-hours of the day and accounts for start-up and running costs when calculating costs.

Mansur (2008) (hereafter “Mansur”) makes the point that deadweight loss calculated based on total variable costs must necessarily be smaller when output is

Table E.1: Comparison across methods to determine competitive benchmark

Cost Calculation	Output Allocation	
	Static	Dynamic
Static	BBW	Mansur
Dynamic	Jha and Leslie	

**Notes:** This table illustrates our contribution to existing literature on competitive benchmarks applicable to wholesale electricity markets. The rows correspond to whether costs are calculated only based on average variable costs (a “static cost calculation method”) or incorporating variable, start-up, and running costs (a “dynamic cost calculation method”). The columns correspond to whether dynamics are accounted for when calculating competitive benchmark output levels (static versus dynamic “output allocation” methods). “BBW”, “Mansur”, and “Jha and Leslie” refer to Borenstein et al. (2002), Mansur (2008), and this paper respectively.

allocated according to dynamic competitive benchmarks rather than static formulations such as Borenstein et al. (2002) (hereafter “BBW”). This is because BBW minimizes fleet-wide total variable costs in each half-hour by allocating output from lowest to highest average variable cost unit at capacity until demand is met. Consequently, *any* other output allocation method that adds operating constraints will result in higher total variable costs than BBW’s benchmark, and therefore deadweight loss will be estimated to be lower. Namely, both our output allocation and the allocation from Mansur will result in lower deadweight loss than BBW provided deadweight loss is calculated only based on total variable costs.

However, Mansur omits start-up and running costs from his calculation of production costs. If these are non-zero, the sign of the bias in DWL from static benchmarks relative to a correctly specified dynamic benchmark is theoretically ambiguous. On the one hand, DWL may be biased upward under static frameworks because static frameworks do not account for relevant operating constraints when setting

output levels—allowing technically infeasible production schedules that drive down benchmark-simulated costs. On the other hand, DWL may be biased downward under static frameworks because static frameworks do not account for start-up and running costs both when allocating output and when calculating the costs associated with observed and benchmark-simulated output levels. For example, starting up a high start cost, low variable cost unit to meet evening peak demand — instead of starting up a low start cost, high variable cost unit — could be misdiagnosed as efficient under the static framework.

Table E.2 presents deadweight loss estimates under each permutation of (1) static versus dynamic cost functions, and (2) static versus dynamic output allocation. First, we find a similar result to Mansur that moving from statically determined output to dynamically determined output significantly increases total variable costs under the least cost benchmark and thus reduces the estimate of DWL based solely on differences in total variable costs (i.e., compare Column 4 to Column 3).

However, moving from traditional static frameworks to our dynamic framework increases total production costs under both the observed output allocation and the least-cost allocation (i.e., compare Column 4 to Column 1). This is because static frameworks ignore start and running costs both when setting output levels and when determining costs. Consequently, the sign of the bias in DWL estimates from static frameworks is theoretically ambiguous. Empirically, we find that the static framework substantially over-estimates DWL relative to our dynamic framework. This bias is driven largely by changes in output allocation under the dynamic framework rather than changes in the cost calculation method (i.e., including start and running

costs).

Table E.2: Hourly average observed production costs, simulated production costs, and DWL from static versus dynamic methods to allocate output and calculate costs

	(1)	(2)	(3)	(4)
<b>Output allocation method</b>	Dynamic	Static	Dynamic	Static
<b>Cost calculation method</b>	Dynamic	Dynamic	Static	Static
<b>Costs of observed output levels</b>	16,656	16,656	15,174	15,174
<b>Costs of simulated output levels</b>	15,177	10,873	14,606	9,282
<b>DWL (Observed Cost-Simulated Cost)</b>	1,480	5,783	568	5,892

**Notes:** This table presents the hour-of-day specific averages of observed production costs, the production costs simulated from the relevant competitive benchmark, and deadweight loss (i.e., the difference between observed production costs and simulated production costs). To set output levels for each unit, we consider the dynamic competitive benchmark specified in Section IV (labeled “Dynamic” output allocation method in the table) versus a traditional static competitive benchmark akin to Borenstein et al. (2002) (labeled “Static” output allocation method in the table). When calculating the costs associated with observed and competitive-benchmark-simulated output levels, we can use either our dynamic cost functions that account for variable, start-up, and running costs (labeled “Dynamic” cost calculation method in the table) or we can use traditional static formulations that only account for variable costs (labeled “Static” cost calculation method).

## **F Linking market outcomes to solar penetration: Additional results**

This section presents additional results pertaining to the regression analysis in Section VI.

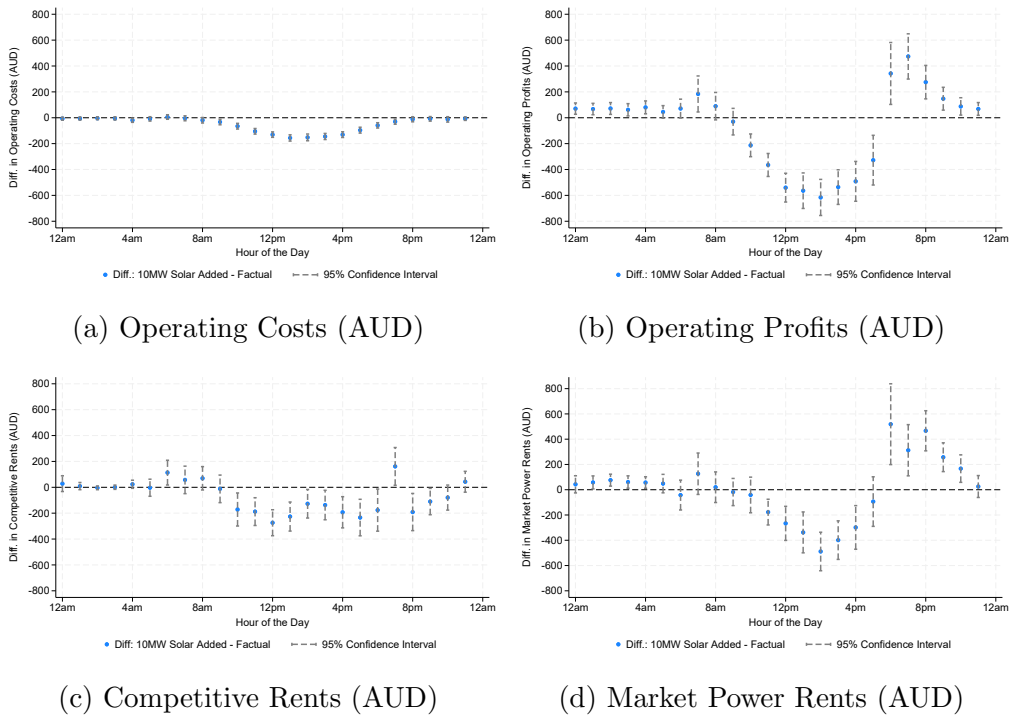
### **F.1 Market outcomes aggregated across the coal- and gas-fired fleet**

Figure F.1 plots the hourly average change in model predictions from adding 10MW of rooftop solar capacity to the system. The bars presented are 95% confidence intervals based on standard errors clustered by month-of-sample. In contrast with the figures in Section B, the market outcomes considered in this section are aggregated across both the coal- and gas- fleet. The conclusions drawn remain exactly the same as those in the main text figures focusing exclusively on the gas-fired fleet.

Namely, we see from Figure F.1 that a 10 MW increase in rooftop solar capacity corresponds to substantial reduction in the aggregate operating costs, operating profits, competitive rents, and market power rents earned by the coal- and gas- fired fleet during the day. This is intuitive: output from rooftop solar panels displaces fossil-fuel production during the day, lowering prices and potentially increasing competitiveness.

However, Figure F.1b indicates that a 10MW increase in solar capacity corresponds to *increases* in operating profits in the evening and early morning. Compar-

Figure F.1: Hourly average changes in predicted market outcomes from adding 10 MW of rooftop solar capacity: Coal and gas units



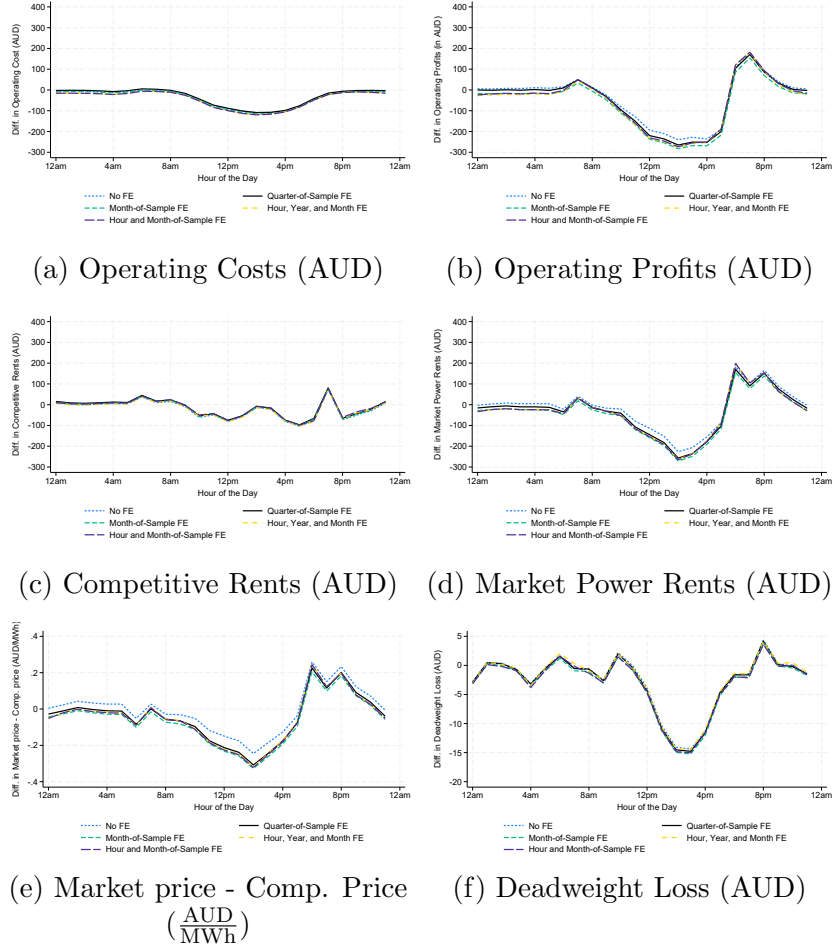
**Notes:** This figure plots the hour-of-day specific averages of hourly aggregate differences in predictions of market outcomes with versus without adding 10MW of rooftop solar capacity to the system. The bars presented are 95% confidence intervals based on standard errors clustered by month-of-sample. Panel (a) considers the total operating costs incurred by coal- and gas-fired units. Panels (b), (c), and (d) report averages for operating profits, competitive rents and market power rents, again aggregated over both coal- and gas- fired units. Competitive rents, market power rents, competitive benchmark prices, and deadweight loss are calculated using the dynamic framework discussed in Section IV.

ing Figures F.1c and F.1d, we see that increases in predicted operating profits in the evening are due primarily to increases in market power rents rather than competitive rents. Namely, competitive rents—the aggregate profits earned by the fossil-fuel-fired fleet if revenues were determined based on our competitive price series—fall with rooftop solar penetration from 10am-8pm, with the exception of a small average increase at 7pm. In contrast, we see large increases in market power rents—operating profits minus competitive rents—in the evening with a 10MW increase in rooftop solar penetration. This provides evidence that the bulk of the increases in operating profits in the evening are driven by increases in the exercise of market power.

## **F.2 Robustness to different sets of fixed effects**

The differences in average predictions plotted in Figure 7 were calculated by estimating Equation (12) in Section A. In this subsection, we show that the patterns in predictions remain similar if we predict outcomes after adding different sets of fixed effects to Equation (12). Specifically, Figure F.2 plots the hourly average differences in predictions with versus without adding 10 MW of solar capacity from 5 different regression models including: (1) no fixed effects (our primary specification), (2) quarter-of-sample fixed effects, (3) month-of-sample fixed effects, (4) year fixed effects, hour-of-day fixed effects, and month-of-year fixed effects, and (5) month-of-sample fixed effects and hour-of-day fixed effects.

Figure F.2: Hourly average differences in predicted outcomes: Different sets of fixed effects



**Notes:** This figure plots the hour-of-day specific averages of hourly aggregate differences in predictions of market outcomes with versus without adding 10MW of rooftop solar capacity to the system. In each panel, predictions are calculated using the regression specification in Section A, but including (1) no fixed effects, (2) quarter-of-sample fixed effects, (3) month-of-sample fixed effects, (4) year fixed effects, hour-of-day fixed effects, and month-of-year fixed effects, and (5) month-of-sample fixed effects and hour-of-day fixed effects. Panel (a) considers the total operating costs incurred by gas-fired units. Panels (b), (c), and (d) report average predicted differences for operating profits, competitive rents and market power rents, again aggregated across the gas-fired fleet. Panels (e) and (f) focus on the markup in market prices above competitive prices and deadweight loss (i.e., the difference between observed production costs and those from our dynamic competitive benchmark) respectively. Competitive rents, market power rents, competitive benchmark prices, and deadweight loss are calculated using the dynamic framework discussed in Section IV.

### **F.3 Trimming the top 1% and bottom 1% of the outcome**

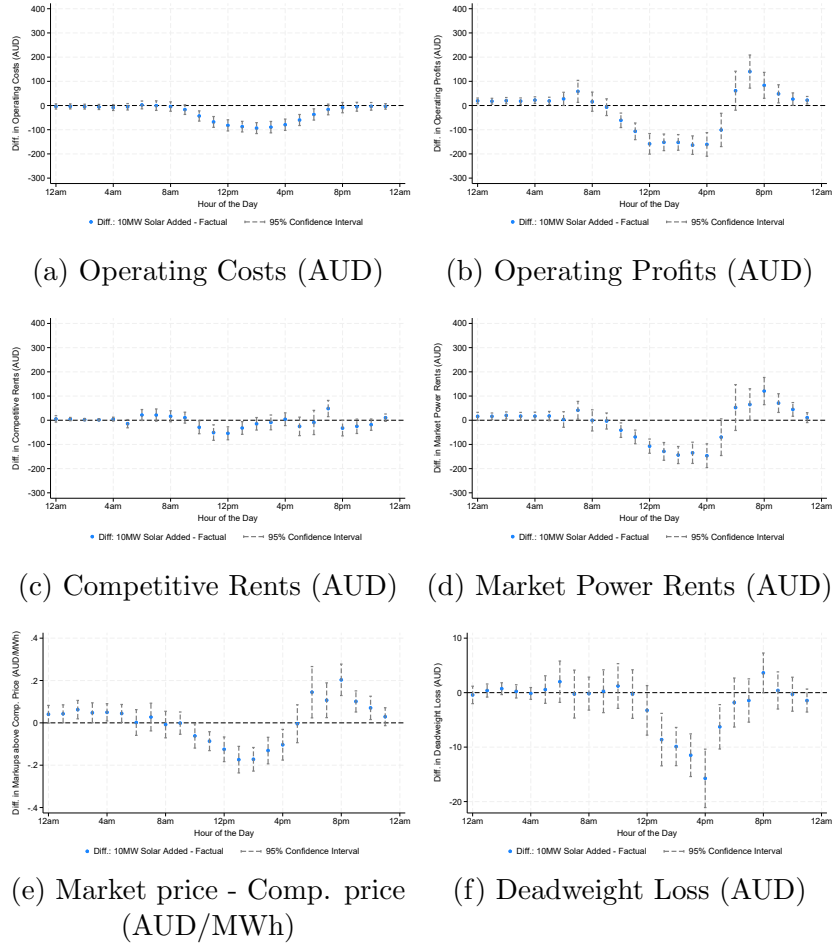
In this subsection, we trim the top 1% and bottom 1% of the relevant market outcome before estimating Equation (12) from Section A. We predict the market outcome with versus without adding 10 MW of solar capacity to the system for each half-hour-of-sample for the observations remaining after trimming. The hour-of-day specific averages of hourly aggregate differences in predictions are presented in Figure F.3.

The story remains the same regardless of whether the outcomes are trimmed or not. Namely, Figure F.3 indicates that adding 10 MW of solar capacity to the system corresponds to reductions in all of our market outcomes during the day. However, operating profits and markups in wholesale prices above competitive benchmark prices are larger in the evening when adding 10 MW of solar capacity. As with Figure 7 in the main text, increases in operating profits and markups in the evening are driven primarily by increases in market power rents rather than increases in competitive rents (see Figures F.3c-F.3e).

### **F.4 Predicted versus observed outcomes**

Figure F.4 plots hour-of-day specific averages of observed market outcomes versus market outcomes predicted using Equation (12) from Section A. As in the machine learning literature, we focus on “out-of-sample” prediction in this figure (e.g., Burlig et al. (2020); Jarvis et al. (2022)). Specifically, we estimate Equation (12) on a randomly chosen 80% of the data-set and assess observed versus predicted outcomes for the other 20% of the data-set. The out-of-sample  $R^2$  is reported underneath each

Figure F.3: Hourly average predicted changes in trimmed market outcomes when adding 10MW of rooftop solar capacity



**Notes:** This figure plots the hour-of-day specific averages of hourly aggregate differences in predictions of market outcomes with versus without adding 10MW of rooftop solar capacity to the system. We trim the top 1% and bottom 1% of the market outcome prior to estimating the regression model. The bars presented are 95% confidence intervals based on standard errors clustered by month-of-sample. Panel (a) considers the hourly total operating costs incurred by gas-fired units. Panels (b), (c), and (d) report average predicted differences for hourly total operating profits, competitive rents, and market power rents, again aggregated across the gas-fired fleet. Panels (e) and (f) focus on the hourly average markup in market prices above competitive benchmark prices and hourly aggregate deadweight loss (i.e., the difference between observed production costs and those simulated from our dynamic competitive benchmark) respectively. Competitive rents, market power rents, competitive benchmark prices, and deadweight loss are calculated using the dynamic framework discussed in Section IV.

panel.

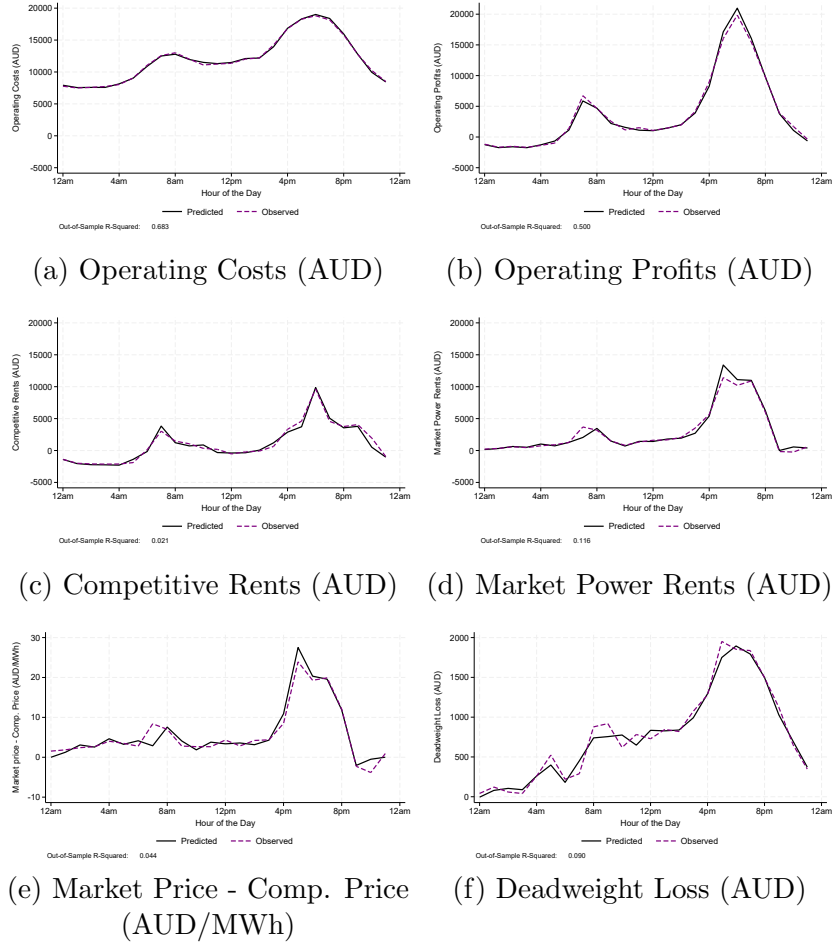
Panel (a) of Figure F.4 focuses on the hourly total operating costs incurred by gas-fired units while Panel (b) focuses on the hourly total operating profits earned by these units. Competitive rents and market power rents are plotted in Panels (c) and (d) respectively. Finally, Panel (e) focuses on the hourly average markup in wholesale prices above our competitive benchmark price while Panel (f) focuses on hourly aggregate deadweight loss. For Panels (c)-(f), the relevant magnitude is constructed using the prices and output levels from the dynamic competitive benchmark presented in Section IV.

Figure F.4 demonstrates that our linear regression model predicts all six market outcomes well on average even out-of-sample. Even for panels with relatively low out-of-sample  $R^2$ , the hourly averages of observed versus predicted outcomes track each other quite closely. This is especially comforting because we compare the hourly averages of predictions from Equation (12) with versus without adding 10MW of rooftop solar capacity to the system.

## F.5 Effects of solar by season

This subsection explores how hour-of-day specific average differences in predicted outcomes with versus without adding 10MW of solar capacity to the system vary by season. We see from Figure F.5 that a 10MW increase in solar capacity corresponds to the largest reductions in operating costs, operating profits, and market power rents during daylight hours in the summer relative to the other seasons. This is intuitive:

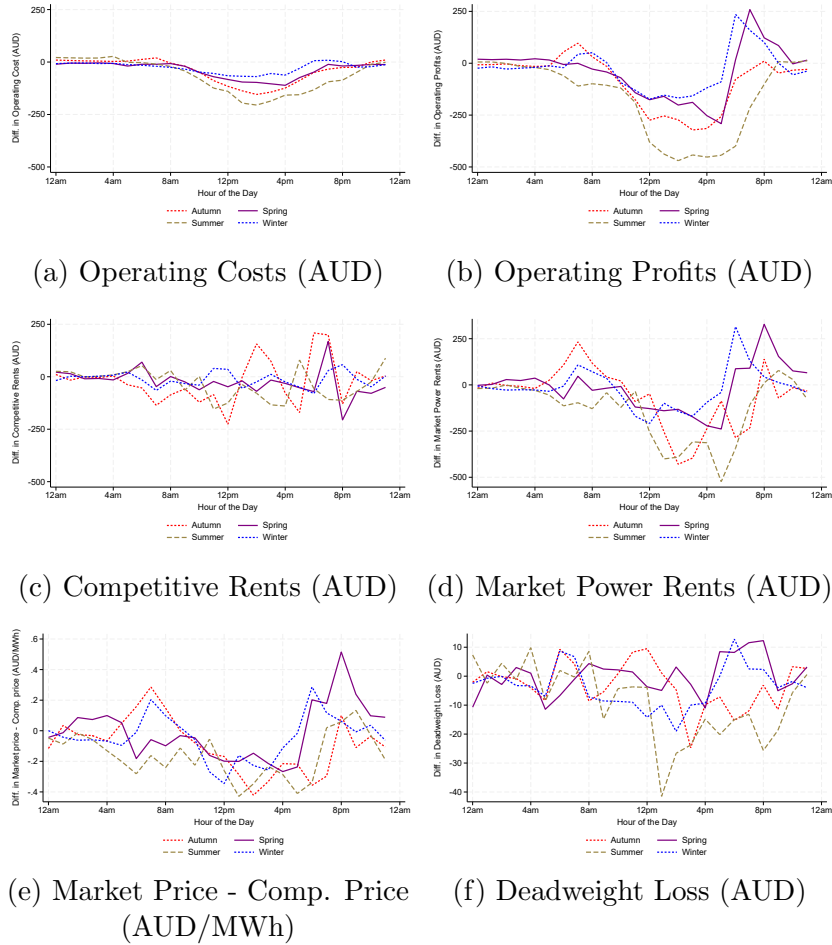
Figure F.4: Hourly average predicted and observed market outcomes



**Notes:** This figure plots the hour-of-day specific averages of observed and predicted outcomes. We estimate Equation (12) on a randomly chosen 80% of the data-set and compare observed versus predicted outcomes for the other 20% of the data-set. Panel (a) considers the hourly total operating costs incurred by gas-fired units. Panels (b), (c), and (d) report averages for hourly total operating profits, competitive rents and market power rents, again aggregated across the gas-fired fleet. Panels (e) and (f) focus on the hourly average markup in market prices above competitive benchmark prices and hourly aggregate deadweight loss (i.e., the difference between observed production costs and those simulated from our dynamic competitive benchmark) respectively. Competitive rents, market power rents, competitive benchmark prices, and deadweight loss are calculated using the dynamic framework discussed in Section IV.

production from solar is tied both to how long the sun is out and how brightly the sun shines; output from a given amount of solar capacity will thus be largest in the summer. Of course, increases in rooftop solar output directly correspond to decreases in the residual demand to be served by the gas-fired fleet.

Figure F.5: Hourly average changes in predicted market outcomes from adding 10MW of rooftop solar capacity by season



**Notes:** This figure plots hour-of-day specific averages of differences in predictions of market outcomes with versus without adding 10MW of rooftop solar capacity; these averages are calculated separately for each season. Panel (a) considers the hourly total operating costs incurred by gas-fired units. Panels (b), (c), and (d) report averages for hourly total operating profits, competitive rents and market power rents, again aggregated across the gas-fired fleet. Panels (e) and (f) focus on markups in market prices above competitive benchmark prices and deadweight loss—the difference between hourly total observed operating costs and those from the dynamic competitive benchmark. Competitive rents, market power rents, competitive benchmark prices, and deadweight loss are calculated using the dynamic framework discussed in Section IV.

Interestingly, we see that a 10 MW increase in rooftop solar capacity corresponds to relatively small increases at sunset in operating profits, market power rents, and markups in the summer. This suggests that the market impacts of a given increase in solar capacity are tied not only to the level of output from this capacity, but also the pattern of solar output across hours of the day. Namely, 10 MW of solar capacity produces the most during the day in the summer. However, the sun also sets at a later hour in the summer, leading to a more gradual increase in wholesale demand from day to evening. This in turn leads to a more gradual increase in the number of gas-fired units starting up in the hours before peak evening demand. A supplier deciding whether to start their unit up to meet evening demand has a good sense of the level of competition they will face given this gradual increase in the number of units operating as the sun sets.

In contrast, in the other seasons, the sun sets relatively earlier and there is a steeper system ramp. Therefore, units that do turn on during this system ramp will operate for shorter spells than those that turn on during flatter system ramps. If many units are to start up, competition will be strong, leading to prices that may be insufficient to cover start-up costs. For this reason, some units will choose not to start up, leading to less effective competition in the evening and thus more pronounced increases in the ability to exercise market power for those that are operating, and subsequently, an increase in market power rents.

## F.6 Estimation excluding lagged wholesale demand

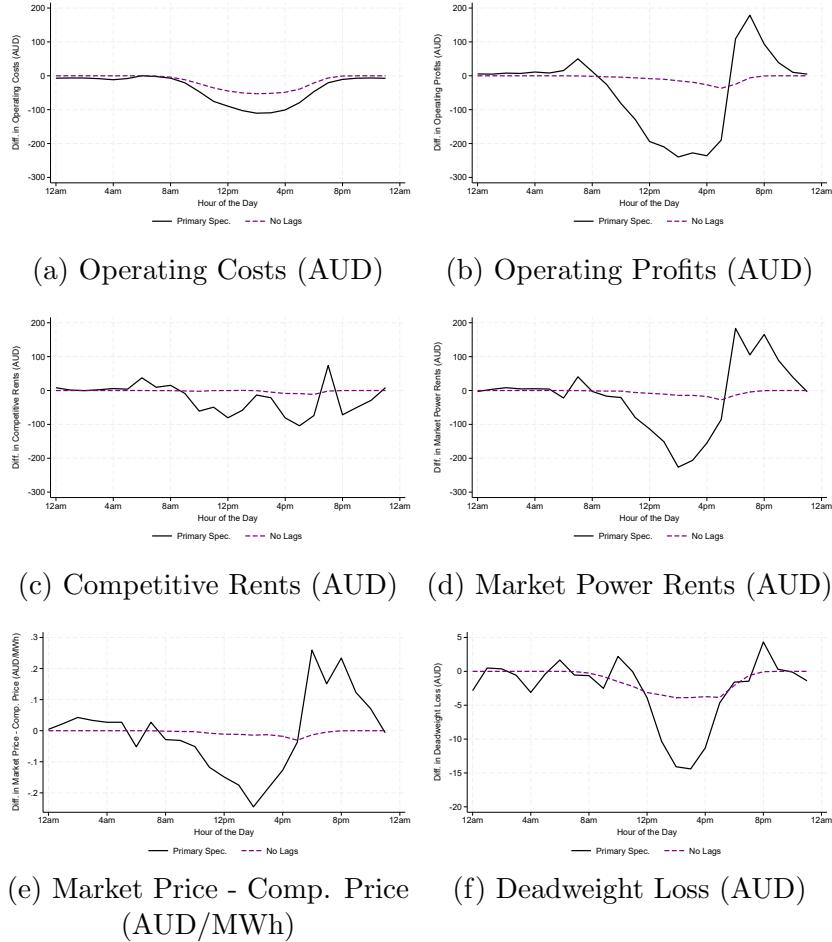
In this subsection, we estimate the relationship between current wholesale demand and market outcomes excluding the terms corresponding to lagged demand. Specifically, we estimate:

$$Y_t = \gamma_h D_t + \epsilon_t \tag{F.1}$$

where  $t$  indexes the half-hour-of-sample corresponding to half-hour of the day  $h$ . Figure F.6 plots the hour-of-day specific average differences in predictions with versus without adding 10 MW of solar capacity for predictions calculated in two different ways: (1) using Equation (12) which includes the 48 lags of half-hourly wholesale demand versus (2) using Equation (F.1) without lagged wholesale demand.

Including versus excluding lagged wholesale demand in the regression framework substantially changes the predicted impacts of solar capacity on all of our market outcomes. For example, relative to the regression framework without lags of demand, the specification including lags of demand predicts a much larger drop in operating profits, market power rents, and markups in wholesale prices above competitive benchmark prices during the day from a 10MW increase in solar capacity. Moreover, a 10 MW increase in solar capacity results in substantial increases in predicted operating profits, market power rents, and markups in the evening when using the regression model with lags but not the regression model without lags. This further highlights the importance of accounting for the dynamic impacts of solar penetration on market outcomes generated by start-up costs and operating constraints.

Figure F.6: Differences in hourly average outcomes predicted using regression models with versus without lagged wholesale demand



**Notes:** This figure plots the hour-of-day specific averages of differences in predictions of market outcomes with versus without adding 10MW of rooftop solar capacity. In each panel, predictions are calculated using two different regression specifications: (1) specifications including the 48 lags of half-hourly wholesale demand (see Equation 12) versus (2) specifications excluding lagged wholesale demand (see Appendix Equation F.1). Panel (a) considers the hourly total operating costs incurred by gas-fired units. Panels (b), (c), and (d) report averages for hourly total operating profits, competitive rents and market power rents, again aggregated across the gas-fired fleet. Panels (e) and (f) focus on the markups in market prices above our competitive benchmark prices and deadweight loss (i.e., observed production costs minus those from our dynamic competitive benchmark) respectively. Competitive rents, market power rents, competitive benchmark prices, and deadweight loss are calculated using the dynamic framework discussed in Section IV.

## **G Firm-level evidence linking rooftop solar penetration and competition: Additional figures**


This section presents additional figures pertaining to the firm-level analyses conducted in Section VII.

### **G.1 Inverse semi-elasticities and rooftop solar penetration: Descriptive trends**


Figure G.1 plots each firm's hour-of-day specific average inverse semi-elasticities for the sample periods 7/2014-6/2015, 7/2015-6/2016, and 7/2016-6/2018. We exclude the sample period 1/2014-6/2014 from this figure because Leslie (2018) demonstrates that inverse semi-elasticities (ISEs) were higher on average as a consequence of the carbon tax in place until 6/2014. We construct ISEs for the day-ahead market (left panels) and the real-time market (right panels).

For all three firms for both the day-ahead and real-time markets, we see that ISEs become higher on average in the evening coincident with the rooftop solar boom from 2014-2018. This is consistent with our central hypothesis: rooftop solar output displaces fossil-fuel units during the day, leading to less competition in the evening as some of the displaced units choose not to start or ramp back up. There are relatively small changes in ISEs from 2014-2018 outside of evening peak hours and the patterns are less consistent across firms and markets. This potentially points to important differences in each major supplier's ability and incentive to exercise market


Figure G.1: Hourly average inverse semi-elasticities by firm and market



(a) Synergy, Day-ahead market (b) Synergy, Real-time market



(c) Alinta, Day-ahead market (d) Alinta, Real-time market



(e) NewGen, Day-ahead market (f) NewGen, Real-time market

**Notes:** This figure plots hour-of-the-day specific average inverse semi-elasticities calculated for each firm for the day-ahead market (left panels) and real-time market (right panels) for the sample periods 7/2014-6/2015, 7/2015-6/2016, and 7/2016-6/2018. A firm's inverse semi-elasticity—based on the slope of its residual demand curve at the market clearing price—measures how much a firm can raise prices (in AUD per MWh) by decreasing its output from 1%.

power, tied to differences in capacity shares and generating mix. We speak more on these differences—and the resulting differences in equilibrium quantity supplied—in Section A.

A potential concern is that trends in ISEs over time are driven by secular changes in economic conditions rather than the rooftop solar boom. For this reason, Section B focuses on the regression estimates of solar-induced changes in ISEs that account for these secular trends.

## G.2 Inverse semi-elasticities and market outcomes

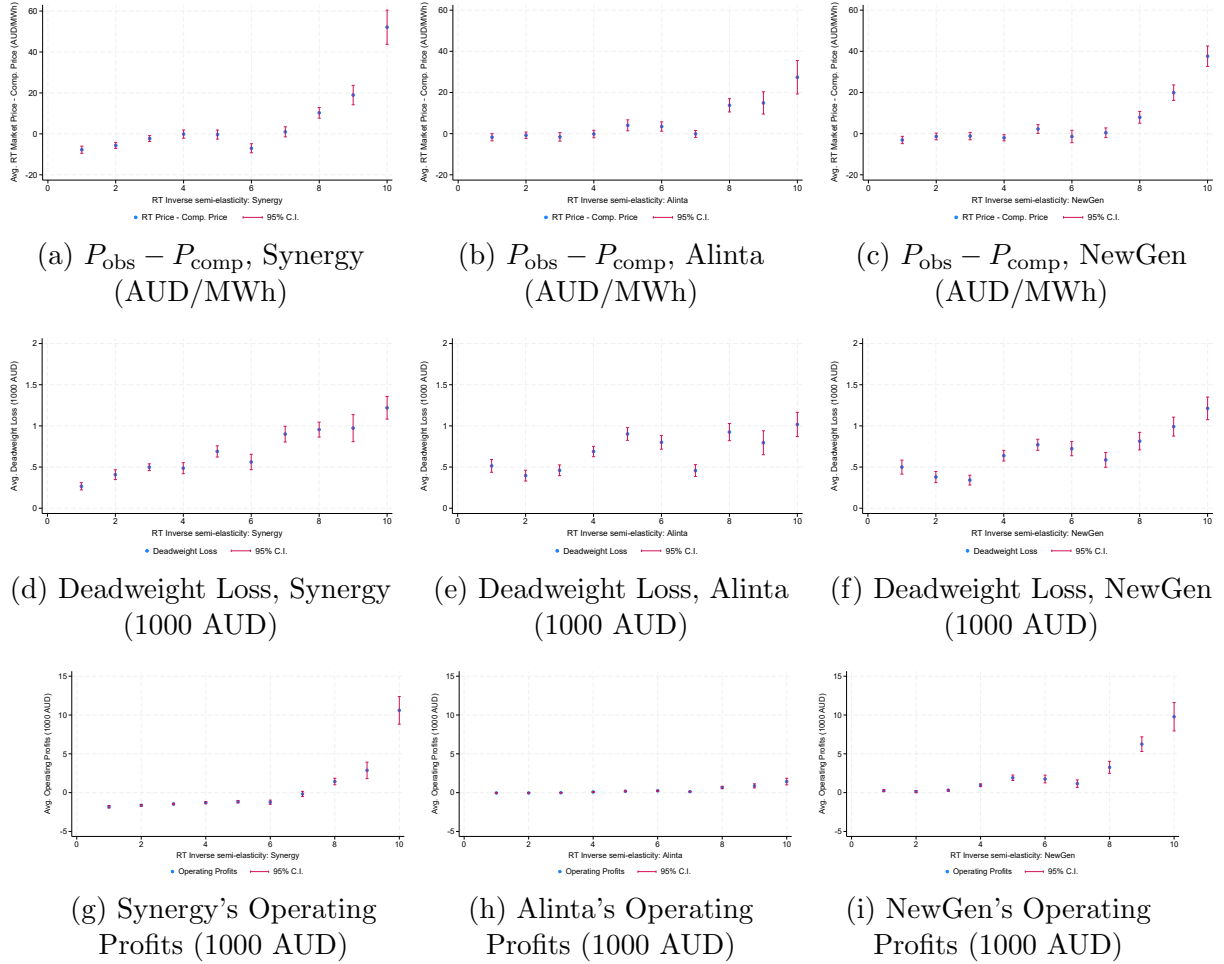
Figure G.2 plots the averages of each outcome as a function of binned firm-level inverse semi-elasticities (ISEs) in the real-time market.<sup>53</sup> The top row focuses on market prices minus competitive benchmark prices (labeled  $P_{\text{obs}} - P_{\text{comp}}$ ), the second row focuses on deadweight loss (i.e., observed operating costs minus operating costs from our competitive benchmark), while the bottom row focuses on the firm’s operating profits. To create these figures, we find the deciles of each firm’s inverse semi-elasticity over the sample period and categorize each half-hour-of-sample based on which decile it falls into.

For all three firms, price margins relative to our benchmark, deadweight loss, and the firm’s operating profits increase with the level of the firm’s inverse semi-elasticity. This evidence buttresses our assertion that deviations between observed outcomes and those from our competitive benchmark are indicative of the exercise

---

<sup>53</sup>The figures based on ISEs in the day-ahead market are similar, and are available upon request.

Figure G.2: Average outcomes by binned real-time inverse semi-elasticity



**Notes:** This figure plots the averages of market outcomes by binned inverse semi-elasticities (ISEs) in the real-time market for each major supplier. The top row focuses on market prices minus competitive benchmark prices (labeled  $P_{\text{obs}} - P_{\text{comp}}$ ), the second row focuses on deadweight loss (i.e., observed operating costs minus operating costs from our competitive benchmark), while the bottom row focuses on the firm's operating profits. To construct binned averages, we find the deciles of each firm's real-time ISE over the sample period and categorize each half-hour-of-sample based on which decile it falls into.

of market power. This evidence also buttresses our ISE measure: suppliers facing a steeper residual demand curve seem to take actions to raise market prices relative to our competitive benchmark prices, which serves to increase their operating profits.

# H Competitive impacts from large-scale solar penetration

In this section, we simulate least-cost market outcomes at counterfactual levels of rooftop solar penetration using the dynamic competitive benchmark outlined in Section IV. The simulation results inform the extent to which the large increases in operating profits coincident with the solar boom from 2015-2018 could have occurred in a competitive market.

## H.1 Methodology

We simulate gas unit output levels and prices for each day  $t$  in the year 2018 using the dynamic competitive framework outlined in Section IV. Namely, for each day  $t$ , we solve the dynamic competitive benchmark for the following three  $48 \times 1$  vectors of half-hourly demand: (1) the observed half-hourly pattern of wholesale demand; (2) the counterfactual half-hourly pattern of wholesale demand implied by 469MW of rooftop solar capacity (the capacity on April 1st 2015, the first day for which we have data on rooftop solar capacity); and (3) the counterfactual half-hourly pattern of wholesale demand implied by zero rooftop solar capacity.

We calculate the wholesale demand  $D_{t,h}^{sim}$  implied by each of these levels of rooftop solar capacity as:

1.  $D_{t,h}^{sim} = D_{t,h}^{obs}$ ,

$$2. D_{t,h}^{sim} = D_{t,h}^{obs} - 0.5(475 - K_t^{obs}) \times CF_{t,h}$$

$$3. D_{t,h}^{sim} = D_{t,h}^{obs} - 0.5(0 - K_t^{obs}) \times CF_{t,h}$$

where  $D_{t,h}^{obs}$  is the observed wholesale demand in half-hour  $h$  of day  $t$ , and  $K_t^{obs}$  is the observed level of rooftop solar capacity on day  $t$ . We multiply by 0.5 to reflect the total energy that would be generated by rooftop solar panels producing at nameplate capacity in a half-hour rather than a hour. Finally,  $CF_{t,h}$  is the observed capacity factor of rooftop solar: the amount of energy that 1MW of rooftop solar capacity would produce in half-hour  $h$  of day  $t$ . This capacity factor is calculated by dividing the aggregate observed production from rooftop solar panels in half-hour  $h$  of day  $t$  by the rooftop solar output implied by all panels producing at nameplate capacity in the half-hour.

In words, we calculate the wholesale demand at time  $t$  under the second counterfactual to be equal to observed total energy use at time  $t$  less the energy that would be generated if 469MW of rooftop solar was installed. This requires first adding the observed production by rooftop solar panels to observed wholesale demand to get total energy use, before subtracting the output that would be generated by 469MW of rooftop solar capacity.

## H.2 Results

Figure H.1 plots the hour-of-day specific averages of hourly aggregate market outcomes simulated under the dynamic competitive benchmark for each of three rooftop solar scenarios. The three solar scenarios considered are: (1) keeping the observed


level of solar capacity, (2) decreasing solar capacity to the levels on 4/1/2015, and (3) removing all solar capacity from the system. As a benchmark for interpreting the magnitudes from this counterfactual analysis, the amount of rooftop solar capacity in WA increased from 469MW to 1,092MW from April 1st 2015 to December 31st 2018.

Figure H.1a plots hour-of-day specific averages of the hourly aggregate operating costs incurred by the gas-fired fleet. This figure documents that operating costs would have increased substantially between 9am-8pm if we counterfactually reduced rooftop solar capacity, either to 4/1/2015 levels or removing rooftop solar entirely. We see very little changes overnight, when solar capacity does not produce.

Figure H.1b plots the hour-of-day specific averages of the hourly aggregate competitive rents earned across the gas-fired fleet. In this case, competitive rents are defined to be the profits that would be earned by the gas-fleet based on benchmark-simulated prices and output levels. Figure H.1b documents that competitive rents decrease with solar penetration during the day, consistent with increases in output from solar displacing fossil-fuel production.


We see slight increases in competitive rents in the evening with increases in rooftop solar capacity. This is attributable to increases in the competitive prices necessary for units displaced by solar output during the day to recover their start-up costs. However, these small increases in competitive rents in the evening are insufficient to explain the more than doubling of annual aggregate operating profits from 2014 to 2018 (a 28 million AUD increase between 2014 and 2018). These

Figure H.1: Hourly averages of simulated competitive market outcomes for different changes in solar capacity



(a) Operating Costs (1,000 AUD)

(b) Competitive Rents (1,000 AUD)



(c) Aggregate Incurred Costs from Starts (1,000 AUD)

(d) Number of Starts By Gas-Fired Units in the Hour

**Notes:** This figure plots the hour-of-day specific averages of hourly aggregate market outcomes simulated using the dynamic competitive benchmark outlined in Section IV. Namely, for each day  $t$ , we solve the dynamic competitive benchmark for the following three  $48 \times 1$  vectors of half-hourly demand: (1) the observed half-hourly pattern of wholesale demand; (2) the counterfactual half-hourly pattern of wholesale demand implied by 469MW of rooftop solar capacity (the capacity on April 1st 2015, the first day for which we have data on rooftop solar capacity); and (3) the counterfactual half-hourly pattern of wholesale demand implied by zero rooftop solar capacity. The top left panel plots the hour-of-day specific averages of hourly total operating costs incurred by the gas-fired fleet, while the top right panel considers competitive rents: the aggregate profits that would be earned by the gas-fleet based on benchmark-simulated prices and output levels. The bottom left panel focuses on the aggregate costs incurred due to starts by gas units, while the bottom right panel plots hour-of-day specific averages of hourly total number of starts by gas units.

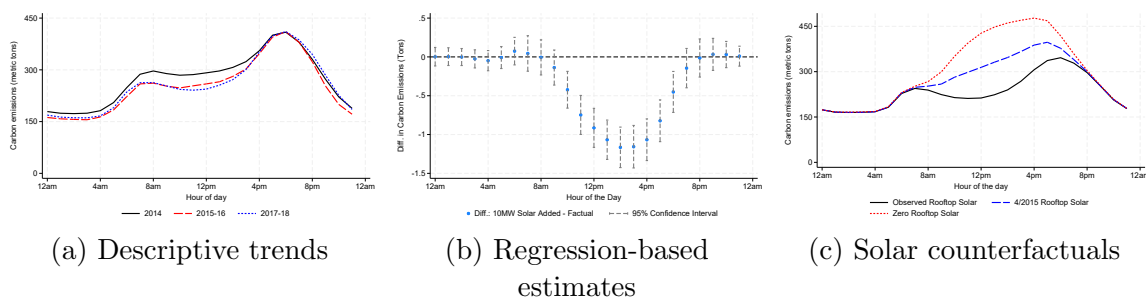
results provide further evidence that solar-induced increases in operating profits in the evening are attributable to increases in market power rents rather than increases in competitive rents.

Finally, Figures H.1c and H.1d plot hour-of-day specific averages of the hourly aggregate costs incurred due to starts and hourly total number of starts respectively. These figures document that increases in solar capacity lead to less units starting up in the middle of the day, but more units starting up in the late afternoon and early evening. This is intuitive: increases in solar output reduce the need for fossil-fuel units to start up during the day, but these displaced units must start up in the evening when solar stops producing.

# I The impacts of rooftop solar penetration on carbon emissions

This section presents results on the impacts of rooftop solar penetration on the aggregate carbon dioxide emissions emitted by the gas-fired fleet in Western Australia (WA). We present results from three methodologies: (1) descriptive trends coincident with the rooftop solar boom from 2014-2018, (2) differences in predictions of market outcomes from a 10MW increase in rooftop solar penetration from the regression framework in Section A, and (3) counterfactual analyses that solve the dynamic competitive benchmark for different levels of rooftop solar penetration (see Appendix Section H).

Figure I.1: Impacts of rooftop solar penetration on aggregate carbon emissions



**Notes:** The left panel of this figure plots hour-of-day specific averages of the hourly aggregate carbon emissions emitted by the gas-fired fleet for the year 2014, the years 2015-2016, and the years 2017-2018. The middle panel of this figure plots the hourly average differences in predictions of aggregate carbon emissions across the gas-fired fleet with versus without adding 10MW of rooftop solar capacity to the system. The bars presented are 95% confidence intervals based on standard errors clustered by month-of-sample. The right panel plots the hour-of-day specific averages of the hourly aggregate carbon emissions emitted by the gas-fired fleet across the year 2018 simulated using our dynamic competitive benchmark for the levels of wholesale demand implied by: (1) the observed level of rooftop solar capacity, (2) the level of rooftop solar capacity on 4/1/2015, and (3) zero rooftop solar capacity.

Across methodologies, we find that increases in rooftop solar penetration lead to

substantial declines in carbon emissions. Output from rooftop solar panels displaces gas-fired electricity production during the day, leading to reductions in fuel burned and thus carbon emissions. There are little to no changes in carbon emissions in the evening when rooftop solar stops producing, reflecting the relatively small share of fuel use attributable to solar-induced increases in starts.