Plugging in the Wind and Sun: Market Design for Shared Infrastructure Under Uncertainty *

Gautam Gowrisankaran, Ashley Langer and Nicholas Ryan.†

March 26, 2026

[Do not cite: Preliminary and incomplete]

Abstract

The past decade has seen a surge in new renewable power projects seeking to plug in to the electric grid. Plant entry on the grid generates cost externalities, whereby each new potential plant can change the costs that other projects nearby must pay to connect. We study the interconnection queue that admits new projects to the California electricity market. We find evidence that externalities matter for investment: (i) plant interconnection costs increase in nearby entry; (ii) projects with lower interconnection costs are more likely to connect; (iii) projects awarded the right to use spare grid capacity are much more likely to connect. We embed these descriptive findings in a dynamic equilibrium model of the interconnection queue that admits externalities between projects. In ongoing work, we analyze how counterfactual queue policies to correct externalities would alter queue behavior and social surplus.

*Thank you to Akshat Adlakha, Patricio Hernandez, Raymond Jin, Tanay Kondiparthay, Nicole Theinova, and Michael Ofodile for excellent Research Assistance. Thank you to Severin Borenstein, Jim Bushnell and Roy Shanker for helpful conversations. Thank you to Ron Borzekowski and Tom Langford at Yale for helping this project launch. Thank you above all to the CAISO team, including Robert Emmert, Jason Foster, Debi Le Vine, Danielle Mills, Bill Weaver and Linda Wright, for making this study possible and for your expert advice. Nicholas Ryan acknowledges financial support from NSF Career Grant #2143098.

†Ryan (corresponding author): Department of Economics, Yale University, Box 208269, New Haven, CT 06520-8269 (e-mail: nicholas.ryan@yale.edu.)

1 Introduction

Public resources are commonly allocated without markets. One reason why governments do not use markets is the complexity of market design when the use of a public good causes externalities. For example, using a public road has external costs for other drivers. Using electromagnetic spectrum can cause interference for other users. In these cases, the difficulty of the market design problem with externalities leads to allocation through rules, like open access or lotteries, which are simple but inefficient.[1]

The green energy revolution has run up against a similarly hard allocation problem. Green energy relies on the power grid because the best generation sources are set by geography and may be far from demand. The last 15 years have seen miraculous declines in the capital costs of solar power (75%), wind power (44%), and battery storage (94%). The backlog of renewable energy projects seeking to connect to the power grid in the United States has surged (Figure 1, panel A). Yet, only a small fraction of these potential entrants have actually plugged in. The 1227 GW of wind and solar capacity now *waiting* to connect in the United States represents an absurd 4.5 times the total capacity installed in the country (Rand et al., 2025; U.S. Energy Information Administration, 2025*b*).

What accounts for the gap between projects seeking to enter and those that actually do? The roots of the problem of connecting new power projects to the grid are externalities between different projects. Projects are rivals for the spare capacity of the power grid. The cost of connecting one plant depends on what other projects seek to connect nearby. Yet, projects are not charged for the costs they impose on other projects. There are therefore missing markets for the externalities that projects cause each other through the power grid. Projects do not know what it will cost them to connect to the grid, and the only way to

[1]Prior to US auctions of electromagnetic spectrum, which accounted for both interference and complementarities in spectrum use, blocks were allocated through committee hearings or by lottery (Milgrom, 1998).

resolve this uncertainty is to enter the process to connect. We conjecture that these missing markets interact with plant uncertainty about costs to drive a wedge between potential and realized entry.

This paper assesses the importance of these missing markets for renewable power in California. California leads the United States in renewable penetration, among large markets, with a renewable generation share of 57% in 2024 and a net-zero generation target for the year 2045 (U.S. Energy Information Administration, 2025*a*; California State Legislature, 2018). California is also important to study because the California Independent System Operator (CAISO), the grid operator in the state, has long used an interconnection process that explicitly accounts for externalities in costs between new power projects. This "cluster-based" process, explained below, is the main reform the Federal Energy Regulatory Commission has recently ordered *all* grid operators nationwide to emulate (Federal Energy Regulatory Commission, 2023*a*). Studying California is therefore a look into the future with respect to both renewable penetration and interconnection policy. This study is possible due to private-use microdata, shared by CAISO under a data use agreement, which covers the interconnection process over the last 14 years for every potential power project in the state, regardless of whether it was built or only proposed.

The paper's analysis proceeds in three main steps. First, we describe the interconnection process and the nature of the markets it is missing. Second, we provide empirical evidence on the importance of these missing markets. Third, we specify a model of the interconnection queue and use the microdata to estimate this model and run counterfactual policy evaluations. The present paper draft includes the model specification but omits model estimates and counterfactual results that are presently work-in-progress.

We begin by describing the interconnection process (Section 2). All projects that wish to connect to the grid must apply to CAISO's interconnection queue. The purpose of the queue is to discover what grid upgrades are needed to connect new projects to the grid

and to allocate the costs for those investments to projects. From a plant's point of view, the interconnection queue resolves uncertainty about the interconnection costs they would have to pay to connect to the grid. Projects that enter the queue receive, after one year, an initial ("Phase 1") estimate of their interconnection costs. These costs pay for investments to accommodate the projects' added generation, including both *point-of-interconnection* costs (POI), at the point the plant plugs in, and *network costs*, elsewhere in the transmission grid. Projects then decide to continue in the queue or exit. After two years, projects receive a revised ("Phase 2") cost estimate. At this point, projects may also receive an allocation of *transmission plan deliverability* (TPD), the right to use the spare capacity of the grid. They then decide again whether to continue, exit, or contract to connect to the grid. Projects that continue in the queue face the same choice again in the future. If projects continue in the queue, their costs may be further revised in later years.

The queue process explicitly accounts for externalities yet does not host markets to correct them. The queue accounts for externalities because CAISO's cohort-based approach calculates costs for the entire set of projects in the queue, together, at a given time. The cost estimates that projects receive to connect, at each step, therefore depend on whether other nearby projects decide to continue in the queue themselves or to exit. Yet, the queue is missing two prominent markets. The first missing market is that projects are not charged for the cost externalities they impose on other projects. The presence of one project may necessitate that another project contribute to a new transmission upgrade; the first project pays for its own share of this upgrade, but not the cost it has imposed on the second. The second missing market is that projects are not charged for the value of the spare capacity of the grid. CAISO allocates TPD, a valuable right to scarce grid capacity, by a scoring rule. Projects that receive this right enjoy an economic benefit but do not pay for the opportunity cost of denying the spare capacity to another plant.

The second part of the paper (Section 3) provides empirical evidence on the importance

of these missing markets for renewable costs and investment. Our data cover every step in the interconnection process. We observe projects that enter the queue, each successive estimate of interconnection costs that they received, at intervals over years, and their ultimate decisions about whether to exit the queue or to contract to connect their project to the grid. We use these data to establish four facts about the interconnection process.

First, interconnection costs are wildly unpredictable. The standard deviation of the log of network costs, the part of costs due to transmission grid upgrades, is fully 3.08 log points. These costs are not very predictable given project characteristics ($R^2 = 0.218$), so projects cannot tell what their costs will be before entering the queue to find out. Even conditional on a first cost estimate, the second cost estimate received varies enormously. This volatility in cost estimates creates option value for projects to remain in the queue, even if their initial cost estimates are high, to receive a new cost draw.

Second, there is strong statistical evidence of externalities in interconnection costs, as would be expected given CAISO's cohort-based cost allocation model. To study externalities, we group projects geographically into study areas, parts of the grid within which projects are especially likely to affect the costs of other projects. We find that the elasticity of initial network costs for a given plant with respect to the aggregate generating capacity in the queue in that plant's study area is 0.8, indicating strong congestion externalities at the time of initial entry in the queue.

Third, interconnection cost estimates matter for whether projects advance toward grid connections. The queue process is dynamic. Projects gain new information and must decide, at several points, whether to post a bond to remain in the queue to proceed. Overall, 75% of projects in our data continue from the Phase I cost study, at the first year, to the Phase II cost study, the year after. We run regressions for the determinants of project progression in the queue conditional on a battery of project-level controls and fixed effects. We find that projects that receive network cost estimates in the second, third, and fourth

quartiles are 10, 25 and 41 percentage points, respectively, less likely to proceed from Phase I to Phase II than a project receiving a network cost estimate in the first quartile. The combination of the second fact (cost externalities) and this third fact imply that the presence of other projects is a major determinant of queue progress.

Fourth and finally, the free allocation of grid capacity has a large effect on whether a project contracts to connect to the grid. The CAISO allocates TPD, the spare capacity of the grid, to projects using an administered scoring rule based on proxies for a project's progress to operation. TPD is valuable because it entitles projects to reimbursement for part of their interconnection costs and because it allows projects to sell their capacity in the capacity market. Among projects without TPD, 10% contract to connect to the grid within four years of queue entry. We find that projects that receive TPD *and* have low network costs are 20 percentage points more likely to connect. The allocation of spare grid capacity therefore *triples* the likelihood that a project will sign a contract to connect to the grid.

Our descriptive analysis of the market therefore shows that the two missing markets in the queue are profoundly important for projects' investment decisions. Cost externalities are quantitatively large. Costs affect project progression. Projects that receive TPD, by rule, are much more likely to connect. While this analysis shows externalities are important, it cannot quantify the efficiency losses caused by these externalities or consider the gains from counterfactual queue policies. To allow us to approach these questions, we build a model of the interconnection process around these stylized facts.

We model projects faced with CAISO's interconnection process playing a dynamic game. Projects enter the queue and receive CAISO's cost quotes . Each year, they must decide whether to exit the queue, continue to receive a new cost estimate, or contract to connect to the grid.[2] Part of the payoffs in this game, including private interconnection

[2]All projects must receive at least two cost estimates before contracting. The first, Phase I study estimate is a rough estimate that is not considered detailed or firm enough for contracting.

costs and bonds that projects must post to continue, are directly observed in the data, which enables us credibly to estimate the monetary value of contracting. Critically, project interconnection cost draws, at each step, depend on the decisions and progression in the queue of other projects in the same area of the grid. Similarly, project allocations of TPD depend on what other projects seek TPD in the same area. In equilibrium, the decisions of projects to enter and continue in the queue affect costs, which in turn affect project decision-making. We thus model cost externalities explicitly and embed them in an appropriate dynamic equilibrium.

In ongoing work, we estimate the model by maximum likelihood and use the model estimates to consider counterfactual policies. The counterfactuals of interest are motivated by policy reforms to decongest the queue by creating market incentives that correct for externalities in the queue. There are potentially many ways to correct these externalities, including raising the costs to enter the queue, charging projects that linger in the queue, and removing policy caps on project cost liabilities.

We use our model to explore two specific policies targeted at the two missing markets in the current interconnection process. First, we calculate project entry, costs and total surplus under a graduated tax that charges projects for each period they remain in the queue. This tax, in principle, should reduce congestion and therefore the cost externalities that projects impose on one another. Second, we implement an auction for transmission plan deliverability, rather than giving out this valuable asset for free. Under this auction projects reaching the Phase II study bid for the right to use a quantity of the TPD available in their area of the grid. The auction bid is a transfer from projects to CAISO and therefore would not affect efficiency directly. However, the prospect of having to bid for TPD removes a large source of project rents from the current interconnection queue. The removal of these rents may discourage marginal projects, which impose externalities but are not very profitable themselves, from entering the queue and imposing externalities in the first place.

This paper contributes to the economics of renewable energy. Renewable energy has physical characteristics, like place-specific generation potential and intermittent, non dispatchable generation, which make it economically unlike conventional fossil-fired power generation (Borenstein, 2012). A range of research has investigated how electricity markets are accommodating renewable growth, in domains including: wholesale energy market behavior (Ito and Reguant, 2016; Bushnell and Novan, 2021), transmission congestion and the value of renewable energy (Gonzales, Ito and Reguant, 2023; Fell, Kaffine and Novan, 2021), and the dynamics of energy market behavior (Jha and Leslie, 2025; Elliott, 2025; Butters, Dorsey and Gowrisankaran, 2025). We study the interconnection queue, an important bottleneck between privately-driven investment in new projects and the publicly-run power grid. Our study elucidates a new source of rivalry in energy markets, where competition acts not mainly through output markets but rather through access to scarce public infrastructure.

The prior paper closest to our research is Johnston, Liu and Yang (2023), which studies the interconnection queue in the Pennsylvania-New Jersey-Maryland (PJM) market using a dynamic model. Our analysis is distinct from that of Johnston, Liu and Yang (2023) in two main respects. First, the PJM interconnection queue, unlike CAISO's, does not use a cluster-based process that explicitly accounts for cost externalities. This kind of process is important to study because it is the main recommendation in a sweeping, nationwide FERC reform of interconnections (Federal Energy Regulatory Commission, 2023*a*). Second, we provide evidence for missing markets, through cost externalities and grid capacity allocation, that Johnston, Liu and Yang (2023) do not study in PJM. The two papers are complementary in that the main emphasis of Johnston, Liu and Yang (2023) is on how projects respond to waiting times in the queue, whereas we study physical externalities and rivalry in grid allocation between projects. While the queue plays out over many years, the forces we study would cause inefficiency even if it were run overnight.

Finally, our paper joins a broader literature in industrial organization on market design with externalities. The literature seeks to measure the efficiency of market outcomes in markets with externalities, such as shipping (Brancaccio, Kalouptsidi and Papageorgiou, 2024), airport slots (Marra, 2024; Bauer, 2025) and electromagnetic spectrum (Doraszelski et al., 2025).[3] In these examples, externalities often act through congestion in access to scarce infrastructure or resources. Our setting is novel in that externalities change the physical cost of investment. Moreover, our analysis benefits from directly observing the private costs of investment, which discipline our estimates of externalities.

The paper continues as follows. Section 2 describes the rationale for the interconnection process and the rules of the CAISO interconnection queue. Section 3 analyzes queue data to estimate externalities and the determinants of plant progress in the queue. Section 4 lays out our dynamic model of the interconnection queue and Section 5 how it is estimated. Section 6, presently empty, will exposit the results from the model. Section 7 briefly discusses the implications of our results for queue policy.

[3]Empirical market design studies also consider congestion effects for individuals in, for example, public housing (Waldinger, 2021) and urban transportation (Almagro et al., 2024).

2 Context and data

2.1 Rationale for interconnection process

The power grid is a public network. The interconnection process seeks to connect new projects to the grid at low cost so that their power can reliably be delivered to customers. Projects cause externalities on each other by congesting the power grid.[4] If projects could freely enter at any point, they could congest the system, and in the extreme might provide no value, if they were unable to deliver their power to customers or prevented other projects from doing so. Projects therefore need the consent of the grid operator to connect.

The interconnection process manages these externalities. The process: (i) discovers what grid upgrades would be needed for prospective projects to deliver their power and (ii) allocates the *interconnection costs* for these upgrades to projects. The process therefore both plans grid investments, on the part of the grid operator, and gives projects precise, location-specific cost signals, so that they can make more efficient decisions about where to invest.

To calculate interconnection costs, the interconnection process turns anticipated externalities in power flow into costs to connect to the grid in the first place. The grid operator uses a model of supply and demand on the power grid to calculate what upgrades new projects would necessitate. For example, if a plant, by connecting, would cause congestion on a transmission line, then that plant will be required to pay for an upgrade to that line, in addition to its own costs for building the plant and plugging in to the grid.

The main economic problem in the interconnection process is that the power grid creates interdependence in connection costs. If many projects are connecting in the same area

[4]Congestion refers to the condition when, at a common system price, the exports from (imports to) any one node of the grid would exceed the transmission capacity connecting that node. In this case, the flow of power between two nodes is rationed (either through separate nodal prices or direct control) so as not to exceed the capacity of the transmission grid.

of the grid, it is more likely that they will cause congestion, and therefore more likely they will require large network upgrades to be built. Externalities in power flow therefore imply positive externalities in interconnection costs. Conversely, because many grid upgrades come in discrete increments of capacity, it is possible that a new plant could split the cost of an existing upgrade—one already required by projects further along in the process—and thereby lower the interconnection costs of those projects.

2.2 Interconnection process within California

The California Independent System Operator (CAISO) uses a cluster-based interconnection process to try to capture these cost externalities. In a cluster-based process, a large cluster (meaning a *cohort*, as a cluster is defined by an annual entry deadline) of projects enters the interconnection process together. For each cohort, CAISO calculates interconnection costs based on a modeled scenario in which they assume all of the projects in the queue will be built. By grouping projects together, the process explicitly allows for cost externalities, since the interconnection costs for each plant can depend on what other projects are seeking to connect at the same time.

To explain how the interconnection process works, we will now describe the milestones in CAISO's cluster-based process. Figure 2, the extensive form game tree for our model, closely mimics the actual timing and information arrival in the interconnection process. This description pertains to the process from cluster 1 (entry year: 2009) to cluster 14 (entry year: 2022). In cluster 14, for reasons we will elucidate, the queue became so congested that CAISO filed an emergency order to the Federal Electricity Regulatory Commission (FERC) to suspend admission to the interconnection queue while it reformed the interconnection process (California Independent System Operator Corporation, 2023). We therefore take cluster 14 as the natural endpoint for our analysis.

Interconnection request.—A project enters the interconnection queue by filing an interconnection request (IR) (Figure 2, action *Enter*). This request details the characteristics of the plant and its desired entry location. For example, a 40 MW wind farm connecting at the 230 kV bus of the Silvergate substation. At the time of filing, a project has to put down a study deposit to cover CAISO's costs for running the interconnection study. These deposits are small relative to the scale of plant investment.[5] CAISO then adds new interconnection requests to its grid model along with both existing projects and entrants from prior cohorts that are still in the queue (have not yet exited). CAISO conducts engineering studies of what upgrades would be required to connect each plant in the cohort to the grid, conditional on all other projects in the queue also being connected. These studies take roughly one year.

Phase I and Phase II cost studies.— A project receives the first *estimates* of its interconnection costs, called the Phase I study costs, after this year (Figure 2, node *Phase I*). The project receives a Phase I study report and can meet with CAISO to discuss study results. The study breaks down interconnection costs into two main components: point-of-interconnection (POI) costs and network costs. Point-of-interconnection (also called facilities) costs are the costs of upgrades at the node of the grid where the project will connect. For example, adding a bay to a 220 kV substation. Network costs are the costs of upgrades elsewhere in the electrical grid. For example, adding transmission capacity in the high-voltage network where the project is expected to cause congestion. Because network costs for a given project depend on the flows of power from other projects that are not located at the same POI, network cost estimates are more likely than POI cost estimates to depend on the presence of other nearby projects. CAISO attributes all network costs to

[5]CAISO study deposits are meant to cover CAISO's human costs of carrying out cost studies, rather than to act as an incentive or disincentive to apply to the queue. Circa 2022, the interconnection study deposit rate was $50,000 plus $1,000 per MW of generating capacity, up to a maximum of $250,000 (CAISO, 2022).

the specific grid investments needed to deliver a project's power. For investments that help to deliver the power of multiple projects, CAISO allocates the total cost of the network investments across all projects whose power flow contributes to making that investment needed.

After learning the Phase I study results, the project has 30 days to decide whether to post a bond, called an Interconnection Financial Security (IFS), to continue in the queue (Figure 2, action *Continue*). If the project does not post the bond it must exit and is refunded any unused portion of its study deposit. If the project does post the IFS, the amount is calculated as a fraction of the interconnection costs. This entitles the project to then receive a Phase II study (Figure 2, node *Phase II*). The Phase II study takes a further year, after which the project receives updated and more detailed engineering plans for what upgrades its connection would necessitate along with updated cost estimates for both the POI and network cost components.

Generator interconnection agreement.—At any point after the Phase II study, the project can decide to sign a contract, called a Generator Interconnection Agreement (GIA), with CAISO to connect to the grid (Figure 2, action *Contract*). In a GIA a project consents to pay its interconnection costs in exchange for CAISO connecting the project to the grid.

The estimated interconnection costs for a project in the queue are subject to periodic *reassessments* (Figure 2, node *Reassessment*). Projects enter and exit the queue each year. The presence or absence of projects can affect the costs of other projects. CAISO therefore periodically issues cost reassessments to projects already in the queue. Projects in the queue are under no time limit to sign a GIA. Projects that linger in the queue without exiting or signing a GIA may see their own costs change, in reassessments, and their presence may also affect the costs of new entrants, through cost externalities.

Transmission plan deliverability.—The interconnection process adds upgrades piecemeal in response to the entry of new projects. The transmission operators (TOs) who run the California grid, including the three major utilities, Pacific Gas & Electric, San Diego Gas & Electric, and Southern California Edison, also plan upgrades on their own, as part of a large-scale transmission plan, to accommodate anticipated supply and demand growth. As part of the interconnection process, projects compete to be allocated part of the spare capacity of the grid created by the transmission plan. The rights to this spare capacity are called *transmission plan deliverability* (TPD). CAISO allocates TPD with a scoring rule based on project characteristics, like whether the project has proven control of the site where it is to be built.

Transmission plan deliverability has a high return, because it both lowers the costs and increases the revenues of potential projects. It lowers costs because projects granted TPD are reimbursed for their interconnection costs up to a reimbursement cap.[6] If a project has a network component of interconnection costs less than the cap, it will be reimbursed for these network costs after connecting to the grid (if network costs are higher than the cap, a project can still be built, but will not be reimbursed for the excess of its costs beyond the cap). The idea is that, because the transmission operators who run the grid are already investing in grid capacity, and charging the costs of these investments to all customers, projects can gain access to the spare capacity these projects create *for free*.

Getting TPD also increases the revenues of potential projects by allowing them to sell capacity in CAISO's capacity market. The CAISO electricity market has separate segments for energy (delivery of electricity) and capacity (the ability reliably to deliver electricity in the future). Utilities must buy capacity to ensure that they can serve all their future demand. Projects are only allowed to sell their capacity if that capacity can be reliably delivered over the electric grid. If a project receives TPD, then its allocation of this spare grid capacity

[6]The cap was $60 per kW up to 2019. Thereafter it rose each year to reach $96 per kW in 2024.

ensures it will be able to deliver its energy over the grid. The project is therefore able to sell its capacity into the capacity market and create an additional revenue stream, above and beyond the expected future value of energy sales.

Discussion.—The cluster-based interconnection process explicitly accounts for externalities in cost among multiple potential entrants. Yet, there are features of the process that suggest it falls short of internalizing the externalities between projects.

- *Uncertainty*. Projects do not know their costs and must enter to learn their costs. CAISO does not know what projects will be built, so cost estimates are conditioned on the presence of projects that may never connect.

- *Signal-jamming*. The presence of other projects affects the own cost of any given project. The queue presence of other projects therefore may affect my information and hence investment.

- *Free allocation of TPD*. TPD is a valuable right to the scarce resource of spare grid capacity, but is allocated for free.

For these several reasons, projects must enter the queue to learn their costs, and might also enter for the prospect of gaining valuable TPD, but, by doing so, projects are likely to change the queue progress and hence investment decisions of other projects. The combination of large option value and externalities creates massive congestion: excessive equilibrium entry (Figure 1). We provide empirical evidence on the strength of these forces in Section 3, after introducing our data.

2.3 Data

From the California Independent System Operator (CAISO).—Our main data are CAISO administrative records for all queue project entrants covering their characteristics,

costs, and decisions over a period of 14 years. This data comes from the Resource Interconnection Management System (RIMS), which is the database and web portal through which CAISO records queue progress and communicates study results to projects.[7]

The RIMS data contain several different sub-components. First, there are project characteristics submitted at the time of application, such as the type of power plant to be built, its entry location and its capacity. Second, there are updates on the status of each project in the queue, such as whether the project has withdrawn from the queue or signed a Generator Interconnection Agreement (GIA), and the dates on which it did so. These records allow us to reconstruct the complete history of decisions for each project at each point in the interconnection process (Figure 2). Third, within the RIMS, projects can access study results, which report to projects for every study the upgrades that would be required for their interconnection and the attendant POI and network costs of those upgrades. We observe these study results for all projects. We therefore know the private information of projects, on which they are basing investment decisions, at every stage of the process. Fourth and finally, we obtain separately from CAISO records on which projects received TPD in each allocation year.

Externalities are more likely for nearby projects. To define what is nearby, we use CAISO's study areas to classify queue entrant locations. Figure A2 illustrates these areas.[8] We validate that these study areas are highly predictive of cost externalities. The purpose of study areas is to approximate electrical neighborhoods, in which projects are more likely to exert externalities on each other. We construct for all projects an adjacency matrix indicating which projects share costs for any network upgrade with which other projects. (We can observe cost sharing completely, for the universe of projects, from the

[7]This project was made possible by a collaborative Data Use Agreement between CAISO and Yale University under which we gained access to the RIMS for all queue projects.

[8]Strictly, we observe study areas only at the locations of projects that entered. We extend the area boundaries to correspond with county boundaries, to cover the entirety of California, and to also cover parts of neighboring states that are technically on the CAISO-controlled transmission grid.

aggregated cost study reports.) We then apply a Louvain clustering algorithm to this matrix, to partition projects into mutually exclusive clusters that are likely to share electrical links. Appendix A.3 describes this algorithm in detail. Figure A2 shows that these groups of projects sharing network upgrades—that is, exerting cost externalities on one another—align fairly closely with our county-based study areas. Appendix Figure A3 shows that our area definitions are powerful predictors of the likelihood that two projects share any network upgrade costs. Two projects selected at random within a study area have a roughly 40% chance of both contributing costs to at least one common upgrade. By contrast, two projects randomly paired across areas have less than a 10% chance.

Sundry data sources.—We augment the CAISO data with additional sources on renewable energy costs and resource potential. We obtain capital cost data for solar, wind, and battery storage from the National Renewable Energy Laboratory's Annual Technology Baseline and related NREL/LBNL inputs (National Renewable Energy Laboratory, 2024; Wiser et al., 2025; Seel et al., 2025), as well as IRENA (International Renewable Energy Agency, 2025). We measure solar potential using Global Horizontal Irradiance (GHI) and wind potential using power density (National Renewable Energy Laboratory, n.d.; World Bank and Technical University of Denmark (DTU), 2025). We collect vacant land prices from CoreLogic (2021). These observable characteristics enter the payoff to building generation capacity across locations in California.

3 Missing markets in the interconnection queue

This section describes the interconnection logjam in the California electricity market and investigates its causes. In the current interconnection process, a large number of projects wait in the queue, neither contracting to connect to the grid nor withdrawing. The tendency

of projects to linger in the queue builds up a logjam of capacity that might be, but will not necessarily be, added to the grid in the future.

We then focus on two missing markets in the interconnection process that contribute to this logjam and raise its economic cost. The first missing market is for congestion. Because the amount of grid capacity needed to connect any one project depends on all the projects seeking to connect nearby, projects impose *cost externalities* on each other through their presence in the queue. The second missing market is for access to the spare capacity of the grid. CAISO gives access to the spare capacity of the grid, a valuable asset, for free, which induces entry by projects seeking to be awarded a share of this capacity. These two forces interact to generate excessive entry, as projects seeking a share of the grid enter the queue and remain in the queue, raising costs for other projects.

3.1 The interconnection logjam

Figure 1 shows the massive increase in renewable energy capacity entering California's interconnection queue with each new cohort. The surge of entry culminated in cluster 14 (entry year 2022), with 34 GW of new renewable energy capacity entering the queue, roughly 1.5 times the existing *stock* of renewable capacity on the grid. CAISO was overwhelmed by this surge of entry and unable to sort through the large volume of projects. It therefore petitioned the Federal Electricity Regulatory Commission (FERC) to temporarily suspend the interconnections process while it reformed the queue process (California Independent System Operator Corporation, 2023). FERC granted this petition (Federal Energy Regulatory Commission, 2023*b*). We discuss the current state of policy towards interconnections in the concluding Section 7.

The increase in queue entry in fact understates the severity of the logjam, because the interconnection queue, at a given time, consists both of newly entering projects and of projects that entered in prior years but *have not left*. A project can leave the queue in two

ways. First, it can withdraw the project from the queue without connecting to the grid (we call this "exit"). Second, it can sign a contract with CAISO, called a Generator Interconnection Agreement (GIA), to commit to pay the costs of interconnection in exchange for CAISO building out the new capacity (substations, transmission lines) needed to connect the project to the grid (we call this "contract"). However, projects are under no obligation to leave after a certain time.

Figure 4 shows the progression of projects through the interconnection queue by months since initial entry. The figure shows the share of projects either remaining in the queue (panel A) or signing a Generator Interconnection Agreement (GIA) (panel B) over time. In each panel, each line indicates a separate cohort (which CAISO calls "clusters") of queue entrants, from cluster 1 (entry year 2009) to cluster 14 (entry year 2022).

There are three main findings from the figure. First, many projects remain in the queue for a long time (panel A). For example, in clusters 11, 12, and 13, about 35% of cohort entrants are still in the queue, having neither exited nor contracted to connect to the grid, after four years. From earlier clusters, a substantial fraction of projects remain in the queue even after *eight* years; for example, about one-quarter of projects from clusters 8 and 9 are still in the queue after eight years. The fraction of projects remaining in the queue at a given queue age has generally increased across successive cohorts, as can be seen comparing across the lines for each cluster.

Second, contracting can take a long time, also, even for those projects that eventually do sign a contract (panel B). There is an initial jump in contracting from around 24 to 30 months after entry. However, additional projects trickle into contracts slowly after this point. For example, for cohort 11, about 10% of projects had contracted at 30 months, but this rose to over 20% at 4 years and approached 30% by 6 years after entry. The rate of contracting, in recent cohorts, is not appreciably higher or lower than in earlier cohorts. As more projects remain in the queue, but projects are not contracting at lower rates, the

implication is that fewer projects are exiting the queue.

Third, there is sharp evidence that queue progression depends on the arrival of information. The interconnection process gives information to projects at discrete points, notably with the arrival of the Phase I cost study (about 12 months after entry) and the Phase II cost study (about 24 months after entry). In panel A, we see a massive drop in projects in the queue after the Phase I study is released and a smaller, but visible drop after the Phase II study is released. Below, we will show that these drops are driven by projects that receive information that they are high cost exiting the queue. Conversely, after the Phase II study, projects are allowed to contract, and projects that receive favorable information, that they have low interconnection costs, begin to contract shortly afterwards (panel B).

3.2 Missing market #1: Cost externalities between projects

The purpose of the queue is to estimate and allocate interconnection costs. Projects face considerable uncertainty in learning about their costs over the interconnection process. In this part we first quantify this cost uncertainty. We then provide evidence for cost externalities: the cost draws for individual projects depend on congestion in the queue. Finally, we show that interconnection costs have a large effect on progression through the queue.

Interconnection cost draws and transitions.—Projects learn about their interconnection costs in stages (Figure 2). Figure 5 shows the evolution of costs over these stages. The panels in the left column show Point of Interconnection (POI) costs, the costs paid by one project alone to plug-in to the grid. The panels in the right column show network costs, the costs paid by one project towards transmission network upgrades that may also be shared with other projects. Within each column, the top panel shows the initial distribution of log costs drawn in the Phase I study. The second panel in the column shows a scatter plot of project costs in the Phase II study (vertical axis) against their initial draw in the Phase I study (horizontal axis). The third panel in the column shows a scatter plot of

project costs in a cost reassessment (vertical axis) against their draw in the Phase II study.

The panels show that there is an enormous range of possible interconnection costs across projects, especially so for the network component of costs. The high variability of network cost draws make sense, because network costs depend on what upgrades a project's entry might trigger elsewhere in the grid. In panel A, for point of interconnection (POI) costs, the standard deviation of costs is 2.82 log points and the inter-quartile range 0.71 log points. In panel B, for network costs, the standard deviation of costs is 3.08 log points and the inter-quartile range 1.15 log points. Therefore, for network costs, it is typical that cost estimates differ, across projects, by 3.1 orders of magnitude. Both the level and the dispersion of network costs are markedly higher than for POI costs. These interconnection costs comprise a meaningful portion of the total costs for a new renewable energy project. For a solar project in 2020, for example, the capital cost of the solar panel is $1,290 per MW (thousands of dollars). The median, p75 and p90 network interconnection costs are $35 per MW, $90 per MW and $322 per MW (thousands of dollars), respectively. Therefore, high network cost draws might represent as much as 7% or 25% of total project capital costs.

The uncertainty in interconnection costs is not resolved after the initial study. CAISO uses the Phase I studies as a preliminary benchmark to give projects guidance on whether to proceed through the queue. Phase II studies may yield different costs, because they do more detailed engineering analyses of connections *and* because the underlying grid conditions in CAISO's model—including the presence of other projects—change over time. Panels C and D plot POI and network costs, respectively, from the Phase II study, against the cost estimate for the same project from the Phase I study. There is large variation in the Phase II study cost even after a project knows its Phase I study cost. In panel D, the $R^2 = 0.45$ for the auto-regression of log Phase II network costs on log Phase I network costs, leaving a root mean square error of 0.61, in logs. Thus projects face large residual uncertainty even after receiving an initial cost estimate. For POI costs, this uncertainty is resolved completely in

Phase II, and POI costs basically do not change in subsequent cost reassessments (panel E). For network costs, this uncertainty is not resolved, and costs may change significantly from the Phase II study to future reassessments (panel F).

Interconnection cost externalities.—What drives the uncertainty in interconnection costs? In this part, we provide evidence that cost externalities across projects contribute.

From engineering principles, the sign of any cost externalities is unknown. The addition of a project to the queue can have two effects on the cost estimate received by another project.[9] First, the addition of a new project may cause a new grid constraint to bind, which then requires a new investment (such as a substation or transmission line upgrade). In this case, there is a positive externality whereby the entrant increases the costs of incumbent projects. Second, a new project share the costs of grid upgrades that CAISO had planned to carry the energy from projects already in the queue. In this case, there is a negative externality whereby the entrant decreases the costs of incumbent projects.

Table 1 regresses interconnection costs on project characteristics and the aggregate state of the queue. The first two columns regress the log of initial POI and network costs on project entry (as in Figure 5, panels A and B) on these covariates.

$$\log c_{jat}^{net} = \beta_1 \log AMW_{a,t-1} + X_j'\beta_2 + \alpha_a + \alpha_t + \varepsilon_{jat} \tag{1}$$

The next two columns regress Phase II study costs on their own lags (as in Figure 5, panels C and D) and the same set of covariates. We partition projects into areas as defined in Section 2.3. We measure the aggregate state of the queue $AMW_{a,t-1}$ with the MW of capacity in the queue within a given project's area of the grid at the time of queue entry for

[9] When CAISO studies the interconnection of new projects, they simulate, in their model of the grid, what transmission investments are needed to reliably deliver the electricity from projects to where it is demanded. When a new project is added to the queue, it is also added to CAISO's model of grid dispatch, and conversely, when a project exits the queue, it is removed from the model. The underlying model drives changes in needed interconnection investments and therefore costs.

that project's cohort.[10]

Table 1, columns 1 and 2 produce evidence of cost externalities. In the first column, with log point of interconnection costs as the dependent variable, the coefficient on Log AGW_τ is close to zero. In the second column, with network costs as the dependent variable, the coefficient on Log AGW_τ is 0.80 (standard error 0.27), meaning that the elasticity of a project's own costs with respect to the capacity of projects seeking to connect in its study area is 0.80, not so far below one. These results make engineering sense. While spillovers in point-of-interconnection costs are possible, if projects share the exact same POI, spillovers in network costs are much more likely, since many projects share expenses for transmission upgrades. We estimate large spillovers in network costs on initial entry. Other coefficients in the column 2 specification are also consistent with power system operation. Projects seeking full capacity deliverability status (FCDS) have higher network costs than other projects, as additional network upgrades are needed to achieve FCDS. Projects, like batteries, that are co-located with other projects have lower network costs because they require fewer transmission upgrades.

Table 1, columns 3 and 4 test for cost externalities in the transition from Phase I study costs to Phase II study costs. Recall from Figure 5 that there is significant volatility in costs across study phases, as CAISO refines their cost estimates and new projects move in and out of the queue. There are two main findings from the table. First, the lagged value of a project's own POI costs and network costs, from the Phase I study, are highly significant predictors of the respective cost components in Phase II (coefficients 0.55 (standard error 0.05) and 0.57 (standard error 0.05), respectively). Second, there is evidence of negative cost externalities in cost transitions, whereby more queue capacity in your area lowers

[10]The timing here is important. Interconnection cost studies that are given out to projects in some year $t+1$ are based upon the cohort of projects that were present in the queue in the year t, because CAISO conducts the studies over the year between entry at t and notifying projects at $t+1$. Additional entry in year $t+1$ itself does not immediately affect the cost studies provided in year $t+1$, though it may affect the revision of those costs in future studies and reassessments.

costs in Phase II relative to Phase I. As discussed above, negative cost externalities are possible if existing projects in the queue share upgrades with new entrants, so that adding new entrants lowers existing projects' costs. We find this to be the case: the coefficient on Log AGW_τ in column 4, for example, is -0.22 (standard error 0.06), representing the elasticity of the Phase II network cost with respect to additional capacity in a project's study area. These estimates of negative cost externalities, specifically between the Phase I and Phase II studies, are consistent with CAISO's own views on cost externalities.[11]

We therefore find evidence for both congestion and decongestion effects, at different stages of the process, though congestion effects are larger, on balance. On initial entry, having more capacity in a project's cohort leads to higher costs for network upgrades. Once a project is already in the queue and has received a phase I study, however, more capacity in the following cohort reduces a project's costs, relative to its own cost on entry.

Interconnection costs affect project advancement.—Cost externalities will matter for investment to the extent that costs affect progress through the queue. In this part we provide evidence that projects with high interconnection costs are more likely to drop out of the queue.

The first step of progress is from the receipt of a Phase 1 study to then exit or continue on to receive a Phase 2 study (Figure 2), second node from left). This decision has stakes, since projects are required to post a partially non-refundable bond to receive a Phase II study. As a consequence, we observe many projects dropping out after receiving their Phase I study results (Figure 4, panel A).

Table 2 regresses an indicator variable for a project receiving a Phase 2 study on the information set of the project at Phase 1, including both of its interconnection cost estimates

[11] See Federal Energy Regulatory Commission (2023*b*), which states "CAISO asserts that because the most common change in the queue is the withdrawal of some interconnection customers, the Phase II interconnection studies and annual reassessments typically remove no longer needed upgrades from interconnection customers' studies, thereby reducing remaining interconnection customers' cost responsibilities."

as explanatory variables. Columns 1 to 3 use continuous measures of costs and column 4 uses quartiles of the cost distributions (omitting the lowest). We find strong evidence that high network costs from the Phase 1 study make it less likely for a project to advance. Both point of interconnection costs and network costs are estimated to have similar, negative effects on a project proceeding to Phase 2, though the coefficients on network costs are more precisely estimated (columns 1 to 3), because network costs are more variable (Figure 5, panels A and B). The effects are large. In column 4, moving from the first to the second quartile of network costs reduces the probability of advancement by 11 pp (standard error 4 pp). Moving from the first to the fourth quartile of network costs reduces the probability of advancement by fully 41 pp (standard error 5 pp). Hence receiving information about high network costs strongly discourages projects from proceeding in the queue. By implication, other projects entering in the area of a given project, by raising congestion and costs (Table 1), exert externalities on the investment decisions of their queue neighbors.

3.3 Missing market #2: Free allocation of access to transmission

The lack of pricing for cost externalities is an important missing market. There is also a *second* market missing from the interconnection queue process, which is a market for access to the spare capacity of the grid. We estimate both a rule for how this spare capacity is allocated and the effect receiving spare capacity on project advancement.

Appendix Table E1 estimates the probability a project will receive transmission plan deliverability (TPD, see Section 2) as a function of its characteristics x_j. There are two economically relevant results. First, battery projects are more likely to receive TPD. This result makes sense because batteries, by choosing when to discharge their energy, can better utilize the existing capacity of the grid. Second, projects in areas where there is more capacity in the interconnection queue are less likely to receive TPD. The spare capacity of the grid is rivalrous. If many projects are seeking TPD, then it reduces the probability that

any given project will receive it.

Figure 6 plots the fraction of projects that have contracted, at a given point after queue entry, based on whether they are allocated TPD ($TPD = 1$) or not ($TPD = 0$). Projects are only allowed to contract after they receive their second cost study, 2 years after queue entry. The TPD allocation typically happens shortly after the Phase II study, between the second and third year after project entry. The figure shows that the chance that a project contracts leaps up immediately, at this time, for projects that receive TPD, and continues to increase steeply until about 4 years after entry, before slowly leveling off. By contrast, only a trickle of projects that do not receive TPD choose to connect to the grid. The raw gap in connection probability between projects with and without TPD is around 30 percentage points at 4 years after entry.

To refine this raw estimate, Table 3 estimates the effect of receiving TPD on project contracting to connect to the grid within four years of entry (i.e., signing a Generator Interconnection Agreement). Because TPD is only allocated after the Phase 2 study, the sample is restricted to projects that receive a Phase 2 study (and for which we have 4 years of data on follow-up, to construct the dependent variable).

The main finding is that receiving TPD—the right to use the spare capacity of the grid—has massive effects on project contracting. Amongst projects that do not receipt TPD, the probability of signing a contract within four years of entry is 10 pp (see table footer). With appropriate controls, in column 2, the effect of TPD receipt on contracting is 25 pp (standard error 4.7 pp), more than tripling the baseline probability. The column 3 and 4 specifications test whether the receipt of TPD interacts with the level of a project's network costs. In column 3, we classify a project as low cost if its network costs are less than \$60 per kW. In column 4, we classify a project as low cost if its network costs are below the (time-varying) threshold for TPD to offer full reimbursement of network costs. We find that projects that *both* receive TPD *and* have low costs are 20.2 pp more likely to

contract (standard error 5.7 pp), while projects that receive TPD but have higher costs are 9.7 pp more likely to contract (standard error 6.5 pp), relative to projects that do not receive TPD. It makes sense that the effect of TPD on contracting is greatest for projects with low costs, because for such projects receipt of TPD acts as a guarantee of reimbursement of *all* of their network costs. For projects with higher costs, TPD reduces network costs of interconnection, but does not eliminate them altogether.

3.4 Discussion: Missing markets cause equilibrium congestion

The above suite of results suggests that the missing markets in the interconnection process may have large effects on investment. We have shown that: (a) project costs depend on congestion; (b) project advancement depends on costs; (c) TPD depends on congestion; (d) project contracting depends on TPD.

These force imply that there could be an inefficient, congested equilibrium. We illustrate the basic idea of this equilibrium in Figure 7. At the top node, projects enter the queue to learn their costs and to learn whether they will get TPD. Because projects enter for this option value of information, more projects enter the queue than will eventually be built, causing congestion (right node). This congestion raises cost estimates for other projects via cost externalities (bottom node). Cost externalities feed back to increase the option value that caused projects to enter in the first place (left node). Externalities increase costs, but also *uncertainty* about costs (Figure 5), because my own costs are a function of the decisions of other projects. If other projects drop out, my costs might get revised downwards. This enhanced option value induces projects to wait in the queue, even though their presence will affect costs for later cohorts of entrants.

The analysis to this point shows the importance of missing markets in the interconnection process. The main limitation is that, by taking a static, one-step view of project decisions, we cannot quantify the strength of externalities overall, throughout the inter-

connection queue, or consider the effect of counterfactual policies on project costs and investment. We now therefore specify a dynamic model suitable for these purposes.

4 Model

Our empirical model incorporates the salient features of CAISO's cluster interconnection process. We consider the problem of energy projects that want to connect to the grid. The model the interconnection process with two stages: entry and queue. In the entry stage potential entrants decide whether to enter the interconnection queue. In the queue stage, projects in the queue learn about their interconnection cost quotes and TPD status over time and periodically make decisions about whether to exit the queue, continue without connecting, or connect to the grid.

We model CAISO as including multiple geographic areas $a = 1, \ldots, A$ (Figure A2). We define areas to allow for cost externalities within study areas, but minimize externalities across areas. We find that our study areas approximate actual cost interdependencies reasonably well (Appendix Figure A.3). This approach allows us to treat each area a as an independent market. We omit area subscripts below to describe one area across periods.

We index each entry cohort, or *cluster*, by $y \in \{1, \ldots, 14\}$. We index periods by $\tau = 0, 1, \ldots, T_{max}$, where $T_{max} = 7$ is the maximum number of periods a project can remain in the queue. Each period of model time corresponds to one year. Subsequent clusters start one period after the previous one and so, as an example, actions for cluster 1 at $\tau = 3$ occur simultaneously with actions for cluster 2 at $\tau = 2$. We use a discount factor of $\beta = 0.9$ corresponding to a hurdle rate that a firm might use.

For each cluster, in period $\tau = 0$, a fixed number of potential entrants in each area simultaneously decide whether to join the queue. In each later period $\tau \geq 1$, projects in the queue may choose whether to *continue* waiting, to *exit* the queue or, for periods $\tau \geq 2$, to *contract* to connect to the grid (Figure 2). Because interconnection costs depend on which other projects are in the queue, these decisions generate the congestion problem described in Section 3.4. The role of the CAISO, in each model period, is to study the state of the

queue and provide projects with new estimates of cost or an allocation of TPD.

The remainder of this section describes the queue stage ($\tau \geq 1$), entry stage ($\tau = 0$), and equilibrium of our model in turn.

4.1 Queue stage

Each project $j = 1, \ldots, J$ has a set of time-invariant characteristics, x_j, that are observable in our data. These characteristics include its geographic area, entry cluster, capacity in MWs, and technology type (solar, wind, or battery). In addition, each project has a privately observable profit shock, $\omega_j \sim N(0, \sigma_\omega)$, which captures persistent unobservable project-level heterogeneity in the value of connecting to the grid. This heterogeneity could be due, for example, to projects having access to sites with higher or lower generation potential, or having different costs of capital.

Each period has two phases. First, the ISO updates the information set of projects by assigning projects' costs. In practice, we can think of this ISO phase as happening either at the start of each period or, equivalently, in the time between projects' decision nodes. For example, CAISO takes time to study the costs of projects and update their information before projects get to move again. Second, with updated information, projects act. We discuss these two phases starting with the regulator phase.

ISO phase.— The ISO runs studies and assigns project costs and TPD status to projects. We will not model the ISO as an optimizing player but rather model it as setting rules for determining interconnection fees and TPD status and conducting cost studies. The ISO performs different cost studies in each period. In period $\tau = 1$, the ISO assigns projects initial estimates of interconnection costs, $c_{j\tau}^{POI}$ and $c_{j\tau}^{Net}$ (Figure 5, panels A and B). In period $\tau = 2$, the ISO updates these costs with refined estimates (Figure 5, panels C and D). At this point, the ISO also assigns TPD to projects. In later periods $\tau \geq 3$, the ISO may update the network costs of projects again, depending on the cluster and the period, in

what is called a cost reassessment. The POI costs and TPD status are fixed after $\tau = 2$.

The cost estimates are based on engineering studies from the ISO's analysis of a cluster (Section 2.1). The model represents this study process as a stochastic mapping from the project's characteristics, including its own prior cost estimate, and queue congestion to an assignment of costs. We view this mapping as a structural relationship from project and queue characteristics to cost estimates, in the sense that it is invariant to the counterfactual policy changes we consider. The ISO assigns transmission plan deliverability to projects based on their characteristics. We similarly view the regulator's TPD determination process as a stochastic mapping from a project's characteristics into a binary variable, TPD_j, which equals one for projects assigned TPD.

CAISO rules limit the exposure of projects to increases in interconnection costs between periods. Accordingly, a project's payoffs depend on its current cost quotes but also the lowest cost quotes that it has received in the past. Let $c_{j\tau}^{Net,min}$ and $c_{j\tau}^{POI,min}$ denote the minimum earlier network and POI cost estimates. In the first period, these are the same as the current cost estimates.

Project phase.— Each period, following the ISO phase, there are a set of active projects that simultaneously choose actions. The potential actions for project j for $\tau \geq 2$ are $a_{j\tau} \in \{\text{exit}, \text{continue}, \text{contract}\}$. At $\tau = 1$ projects are not allowed to contract since the first cost estimates are seen as too uncertain.

At the time projects act we assume they know x_k and cost quotes for all other projects k in the queue, including those from previous clusters. Project j also has private information on their own persistent per MW profit shock ω_j and idiosyncratic per MW cost shocks, $\varepsilon_{j\tau}^{Exit}$, $\varepsilon_{j\tau}^{Continue}$, and $\varepsilon_{j\tau}^{Contract}$, on each action. We let these residuals be distributed type 1 extreme value. We assume that the only potential actions at the terminal period $\tau = T_{max}$ are *Exit* or *Contract*.

We let the idiosyncratic shocks, $\varepsilon_{j\tau}$ and ω_j, be scaled values per MW of capacity. This assumption allows us to compare projects that differ in size by orders of magnitude within the same model. Since costs are measured directly in dollars, we can identify the scale of the ε shocks.[12]

Payoffs.—If a project *continues* in the queue it needs to post an interconnection financial security (IFS) deposit. The decision to proceed in the queue has real stakes because projects only get part of this deposit back if they later *Exit*. In general terms, we can write the per MW flow payoff from continuing as

$$-Deposit(MW_j, c_{j\tau}^{POI}, c_{j\tau}^{POI,min}, c_{j\tau}^{Net}, c_{j\tau}^{Net,min}, \tau) + \varepsilon_{j\tau}^{Continue} \tag{2}$$

where *Deposit* is the additional per MW deposit required to be posted to continue as of period τ. The function that governs deposit amounts is set by CAISO policy. Roughly speaking, projects have to post 15% of their estimated interconnection costs to proceed from $\tau = 1$ to $\tau = 2$, another 15% to proceed from $\tau = 2$ to $\tau = 3$, and the balance on choosing to contract. Appendix D.1 gives exhaustive details.

If a project *exits*, it receives back a fraction of the previous deposits that it put down. The flow payoff per MW from exiting is then

$$Refund(c_{j\tau}^{POI}, c_{j\tau}^{POI,min}, c_{j\tau}^{Net}, c_{j\tau}^{Net,min}, \tau) + \varepsilon_{j\tau}^{Exit}. \tag{3}$$

The previously posted deposit can be inferred from a project's current and lagged cost estimates. Roughly speaking, the refunded deposit includes the entire POI cost deposit but only 50% of the network cost deposit. Because only part of the deposit is refunded, and the deposit is linear in network costs, proceeding in the queue is more costly for projects with high network costs. Again, Appendix D.1 describes the refund function.

[12] In practice, we normalize the scale of ε, and then use the inverse of this scale as a parameter on costs that are measured in dollars per MW.

Finally, if a project *contracts*, it earns the profits from connecting the plant to the grid and generating electricity, but pays the additional interconnection costs which have not already been covered by its deposits. We write this payoff function as

$$\begin{aligned} & \pi(x_j, TPD_j, c_{j\tau}^{POI}, c_{j\tau}^{POI,min}, c_{j\tau}^{Net}, c_{j\tau}^{Net,min}, \tau) + \varepsilon_{j\tau}^{Continue} \\ = \; & \overline{\pi}(x_j, TPD_j, \tau) - \\ & \frac{1}{\sigma_\varepsilon}\left(ICF(MW_j, c_{j\tau}^{POI,min}, c_{j\tau}^{Net}, c_{j\tau}^{Net,min}, \tau) + IFSDeposit(\cdot)\right) + \\ & + \frac{\sigma_\omega}{\sigma_\varepsilon}\omega_j + \varepsilon_{j\tau}^{Continue}. \end{aligned} \tag{4}$$

Here $\overline{\pi}$ is the mean profit for a project based on its observables. The second term is the additional interconnection cost paid by the project. This cost is measured in dollars and scaled by the standard deviation of $\varepsilon_{j\tau}$. The third term is the contribution of project and project-period specific unobservable shocks to profit. The profit shock, ω_j, enters here and is multiplied by a scale parameter.

4.2 Entry Stage

We now turn to how potential projects decide whether to enter the queue or not. This margin is important to model because counterfactual policies may change the volume of projects entering the queue, which ultimately impacts the quantity of renewable projects connecting to the grid.

In the entry stage for each cluster, at $\tau = 0$, there are fixed numbers of potential entrants of technology types $k \in \{\text{solar}, \text{wind}, \text{battery}\}$. Each potential entrant receives a one-time, simultaneous opportunity to enter with its energy project. We allow the cost of entry to vary by technology type, denoting it η_k. We assume that there is an idiosyncratic shock to the cost of entry and its distribution varies by type, $\varepsilon^\eta \sim N(0, \sigma_{\eta k})$.

Before entry, potential entrants observe projects from earlier clusters that are already in the interconnection queue. They do not yet know their own project-specific characteristics

x_j, other than their type. On entry, projects learn their project-specific characteristics x_j, such as their capacity and ω_j. They draw these characteristics from a known distribution.

4.3 Equilibrium

We assume that potential and actual projects in the queue play a Bayesian Markov Perfect Equilibrium (BMPE). Project decisions are interdependent because interconnection costs depend on the entry and continuation of all other projects. The assumption of Markov perfection implies that actions are a function only of payoff-relevant state variables. We use *Bayesian* Markov perfection because projects do not know other projects' persistent profit shocks, ω.

The state variables include the characteristics of all projects in the queue as well as public expectations regarding other projects' ω values. For a given project, the payoff relevant state variables include its fixed characteristics, interconnection costs, time in queue, period τ, as well as persistent and transitory payoff shocks.

$$\mathbf{s}_{j\tau} = \left[\tau \ \mathbf{x}_j \ \omega_j \ c_{j\tau}^{Net} \ c_{j\tau}^{Net,min} \ c_{j\tau}^{POI} \ c_{j\tau}^{POI,min} \ TPD_{j\tau} \ \varepsilon_{j\tau}\right]. \tag{5}$$

Focusing first on fixed project characteristics, some enter profits directly, notably technology, size, and the unobserved persistent shock ω_j. Others, such as voltage, affect payoffs indirectly because they enter interconnection cost estimates. Interconnection costs enter both the payoffs from contracting but also the amount of deposits projects have to post as deposits. Both the current and lagged minimum cost estimates are state variables because the lagged minimum fixes these prior deposits. The time-varying and serially correlated state space for a project therefore has six dimensions (excluding the fixed x_j and the i.i.d. $\varepsilon_{j\tau}$ shocks).

In principle, the BMPE depends on the complete set of these time varying state variables for all projects in the queue across clusters as well as public expectations regarding their

ω values. There are hundreds of projects in the CAISO queue. Given the large number of renewable energy projects, modeling every potential project would imply a huge number of states.

To address this problem and circumvent the curse of dimensionality, we use an equilibrium concept that follows Ifrach and Weintraub (2017)'s Moment-based Markov equilibrium (MME) and Gowrisankaran, Langer and Zhang (2025)'s Approximate Belief Oligopoly Equilibrium (ABOE). As in an MME, we specify an aggregate state that distills the impact of other projects' information into a tractable number of dimensions. In our case, we use the aggregate project capacity (or area-gigawatts AGW_{ay}) either in the interconnection queue or connected to the grid as the aggregate state for each area. This includes capacity from the current cluster and capacity from prior clusters that is either still in the queue or connected. The assumption is that the projects do not have additional information beyond the aggregate state that could affect their action-specific payoffs or state transitions.

The idea of using AGW_{ay} to capture the aggregate state is that it captures cost externalities as well as expectations regarding other projects' ω value. Our AGW_{ay} measure includes the capacity of the projects that enter CAISO's cost study models and are therefore likely to affect own project costs. Higher AGW_{ay} results in both congestion that necessitates additional infrastructure and the potential to share infrastructure costs across projects, thereby affecting projects' interconnection costs (Table 1).

We extend MME with the ABOE approach. This approach allows all projects to act as large players who understand that their decisions influence the aggregate state. Under ABOE, if a project chooses to continue in the queue, it appreciates that AGW will be higher than it would be in expectation, where there is some chance that that project would exit. Similarly, if a potential project enters the queue, it understands that AGW will be higher in the future than expected, because it conditions on its *own* entry.

With this state space reduction, the time-varying serially correlated state space for any

project becomes the six states listed above plus the single aggregate state, AGW_{ay}. We parameterize projects' expectations of the AGW evolution with an AR(1) regression. When an existing project chooses to enter or continue, it adjusts its expectations of the evolution of the aggregate state with $Pr_j(exit) \times MW_j$ to correct for the fact that the AR(1) regression included some probability of the project exiting. We estimate $Pr_j(exit)$ with a logit regression. Similarly, when a potential entrant project chooses to enter, we estimate the expected additional MWs it contributes to AGW with a log-linear regression.

5 Estimation

We estimate the model in three parts. In the first part we estimate the structural engineering cost estimates, and endogenous regressions outside of the dynamic game, mainly with panel regressions. These include TPD—which is a policy function—the AGW transition, the probability of exit in the queue stage, and the expected additional MWs a project contributes at entry. In the second part, we form the quasi-likelihood of choices in the queue stage of the game conditional on these initial functions and estimate queue parameters via quasi-maximum likelihood. In the third part, we estimate the entry stage parameters via quasi-maximum likelihood taking the expected payoffs conditional on entry from the queue stage and the initial regressions as given. We describe these steps at a high level here and in detail in Appendix D.

5.1 Parameters estimated outside the dynamic game

We first estimate regressions of costs and the other values listed above on observed states. These estimates use rich data on projects' advancement through the queue, such as the data on cost transitions described in Figure 5.

Table 4, panel A summarizes the model parameters. At the project level, we estimate

a suite of state arrival and transition processes for: (i) initial cost estimates for network and POI costs; (ii) cost transitions for network and POI costs; (iii) the arrival of TPD. We assume that these regression estimates represent structural parameters that will not change with projects choices, since they are either determined either by engineering studies or by regulatory choices. We also estimate projects' propensity to exit the queue using a static logit regression and the expected MWs of entry for each project in the entry stage as a log-linear function of the entry state. At the aggregate level, we estimate the transition of the aggregate GWs of capacity in each area over time as an AR(1) function. The exit, expected MWs of entry, and aggregate GWs functions are all endogenous, meaning that their parameters may change in counterfactuals.

5.2 Parameters estimated within the dynamic game

We estimate the dynamic model by quasi-maximum likelihood, conditional on the parameters estimation outside of the dynamic game. To form the quasi-likelihood, we first define a finite terminal period T_{max}, after which all projects are assumed to be forced to make a choice between exiting or connecting. We then calculate the value function by backward induction over a gridded representation of the state space. With the value function, we can then use interpolation to calculate choice probabilities for each action at any continuous point in the state space, including the states observed in the data. For a given ω_j and a project that remains in the queue up to $\overline{\tau}_j$, the likelihood of the sequence of actions chosen by a project in the queue is

$$L_j(\omega_j, \theta) = \prod_{\tau=1}^{\overline{\tau}_j} P(a_{j\tau} = a_{j\tau}^{obs} | \mathbf{s}_{j\tau} \backslash \omega_j, \omega_j, \theta), \tag{6}$$

where θ represents the consolidated vector of dynamic parameters in Table 4, panel B.

The model contains persistent unobserved heterogeneity via ω_j. We integrate out this unobserved heterogeneity using quadrature. Let ω^l for $l = 1, \ldots, L$ be the points of support

on the equally-weighted quadrature grid. The likelihood can then be written as

$$\log L(\theta,\Omega) \quad = \quad \sum_{j=1}^{J} \log \left(\sum_{l=1}^{L} (1/L) \prod_{\tau=0}^{\overline{\tau}_j} P(a_{j\tau} = a_{j\tau}^{obs} | \mathbf{s}_{j\tau} \backslash \omega_j, \omega_j = \sigma_\omega \omega^l, \theta) \right). \quad (7)$$

We search over this likelihood with a combination of a Nelder-Mead algorithm and random starting values.

6 Model Results [Preliminary and subject to change]

6.1 Estimation Results

This section presents quasi-maximum likelihood estimates of the structural parameters governing project behavior in both stages of the California interconnection process: the entry stage and the dynamic queue stage. These estimates and simulations are preliminary. We have fully completed coding for estimation and nearly done so for counterfactuals. However, we are actively working to modify the model and code and expect the results here to change.

Entry Stage.—Table 5 reports estimates of the entry cost distribution for each technology type. Following Section 4.2, the entry cost for a project of technology type $k \in \{\text{solar, wind, battery}\}$ is drawn from a normal distribution with mean μ^k and scale σ_η^k. All six parameters are highly statistically significant. Our estimates are all scaled by the entry cost scale term, which means that we recover mean entry costs by dividing mean entry cost parameter by the inverse entry cost scale for any given technology. These parameters describe the full distribution of entry costs among *potential* entrants. Because only projects with sufficiently low costs actually enter the queue, it is particularly informative to combine the estimates with the observed entry rate for each technology type to understand the mean entry cost conditional on the project entering. We observe that potential solar projects enter the queue 20.5% of the time, potential wind projects enter 12.2% of the time, and potential battery projects enter 18.9% of the time. This implies that the mean entry cost conditional on entry for solar is \$46.5M, for wind is \$203.4M, and for battery is \$60.7M.

These estimates are much larger than fees CAISO collects for entering, which are generally around \$50k per project. This is likely because we do not observe the number of potential entrants of each technology type in each region and therefore need to make an

assumption about how many potential entrants are present. We currently assume that the number of potential entrants in each area of each technology type is always equal to 125% of the maximum number of observed entrants in each area of each technology type. We are in the process of investigating how these estimates change with alternative assumptions about the number of potential entrants.

Interconnection Stage.—Table 6 reports quasi-maximum likelihood estimates of the interconnection stage payoff parameters, which determine whether projects will continue in the queue, exit, or connect to the grid. The parameters are organized into those that scale project's profits, common components that affect all technology types, and technology-specific revenue and capital cost components.

These results are still highly preliminary as we explore alternative model specifications. That said, there are a few key results that have been consistent across model specifications. First, we find that the scale of the logit payoff shocks that vary over choices and periods for each project is large relative to the scale of the project-specific persistent profit shock. [13] This variation may come from changes in the opportunity cost of capital or the availability of capital or vacant land for sale. Second, we find that receiving TPD status substantially increases a project's value of connecting to the grid, which projects value at $398 thousand dollars per MW. This is consistent with the reduced-form evidence in Section 3, which showed that projects that receive TPD status are substantially more likely to connect to the grid, and is likely the result of these projects being able to bid into electricity capacity markets. Finally, we find that there are substantial additional costs that we are not accounting for in our model that vary by technology type.

[13] We normalize the coefficient on interconnection fees paid to CAISO to be in millions of dollars, so the scale of the logit shocks, σ_ε has a standard deviation of $1/3.054 = 0.327$

6.2 Model Fit

In order to assess how well our model fits the data, we simulate data for the continuation stage of the model starting at the observed projects in the first year they appear in the data and simulating forward using the model. Table 7 presents summary statistics for the observed data and the data simulated from the model. In general, the simulated data replicates the observed data well, with projects staying in the queue only slightly longer (2.41 years vs 2.05 years), being slightly more likely to sign a General Interconnection Agreement to connect to the grid within four years (13.9 percent vs 13.3 percent), and being slightly more likely to receive TPD (46.9 percent vs 44.3 percent). However, we do see that the minimum POI and network costs over all of the years that the projects are in the data are lower in the simulated data than in the observed data. We are in the process of revising our model to better reflect the fact that cost revisions only arrive probabilistically in later periods, which would reduce the number of new cost quotes that CAISO provides projects, and thereby reduce the probability that projects receive a particularly low cost quote.

Figure 8 plots the progression of projects in model simulations evaluated at the fitted parameters. In broad strokes, the evolution of decisions in the model matches the evolution in the data (Figure 4). Many projects exit early on, and then exit slows to a trickle over later years. Conversely, the share of initial entrants contracting to connect to the grid rises from nothing, in the first period, to reach nearly 30% at the end of the queue stage time horizon. One early finding from the model is the large value of transmission plan deliverability to project progression and contracting. Figure 9 mimics Figure 6 in plotting the fraction of projects that have contracted to connect to the grid by model year, separately for projects that have and have not received TPD. We find that a yawning gap opens between the low rate of contracting for projects without TPD and the higher rate with TPD. At the present model estimates, the size of this gap is roughly in accord with the data, but the level of

contracting for both types of projects is higher than in the data (Figure 6).

These results demonstrate that the model matches the broad patterns of project progression through the queue and contracting behavior. They also show the viability of running dynamic counterfactuals despite the underlying complexity of the dynamic economic environment with externalities.

In ongoing work we are making a few substantial changes to the model that we anticipate will better capture projects' payoffs from connecting to the grid. First, we are adjusting the propensity for projects to receive cost revisions after period 2 to better account for the probabilistic nature of these revisions. Second, we are incorporating new data on whether a project has signed a power purchase agreement with an electric utility, whether a project has already secured control of a site via land purchase or leasing, and the level of the state-wide renewable portfolio standard, which increases over our sample period.

6.3 Counterfactuals

The queue has been the object of a number of major recent policy reforms. The efficacy and social value of these reforms is essentially unknown and will take years to play out. The model can therefore play a critical role in informing the direction of queue policy.

We are in the process of using our estimated model to evaluate a set of counterfactuals that examine how equilibrium outcomes would change under alternative interconnection queue mechanisms. We plan to simulate two different counterfactuals that act on the two *missing markets* that our analysis has identified (Section 3).

Cost externalities: Queue persistence tax.—The first missing market we identify is due to unpriced cost externalities between projects. Projects are not charged for the costs they impose on other projects. This suggests that an appropriate intervention is to raise the cost of waiting in the queue without taking any decision.

We therefore propose to test a modified interconnection deposit policy where projects

are required to post nonrefundable, escalating deposits for each period that they remain in the queue without contracting. For example, the queue congestion tax would start at 10% of their combined network and POI costs in period $\tau = 3$ and rise by this amount for each period a project remained thereafter wihtout connecting. These payments would reduce the option value of continuing in the queue and will likely lead to projects choosing either to exit or contract more quickly after period two—put up, or get out. The case for social efficiency is that they will also reduce the cost estimates of new generations of projects entering at the same time and later.

Free allocation of access to transmission: Grid capacity auction.—The second missing market we identify is due to CAISO's free allocation of the scarce spare capacity of the transmission grid. TPD is extremely valuable. Projects that enter the queue are competing for the rents they could earn from a free allocation of this scarce TPD. This suggests that an appropriate intervention is to reduce these rents by auctioning off grid capacity for its fair market value.

We therefore propose a second, focal counterfactual that simulates a Vickrey auction for TPD status in period 2 for each cohort. In these multi-unit Vickrey auctions, projects would bid a per-MW value of receiving TPD. Projects that bid the highest values would receive TPD in an area, up to the same amount (MW) of TPD that CAISO actually awarded in practice, and pay the bid price from the highest-value bid of a project that did not receive TPD in the auction. The beauty of this proposal is that, since the Vickrey auction is incentive compatible, we can use the value function in our dynamic model to easily compute the appropriate bid for each project, which is the *difference* in the continuation value for that project between having $TPD_j = 0$ and $TPD_j = 1$ in period $\tau = 2$.

This proposal would be expected to act on *both* missing markets. First, it may allocate TPD to higher-value projects than does CAISO's present administrative allocation rule. The

rule is based on project characteristics and may therefore be correlated with payoffs, but it is unlikely to capture the range of variation in payoffs across projects. Second, running an auction for TPD removes a major source of rents for projects. If there is no prospect of gaining TPD for free, projects will be less likely to enter and persist in the queue at earlier periods. This anticipated reduction in congestion will reduce cost externalities, though the auction does not target this externality directly.

7 Conclusion

The advent of cheap renewable energy is one of the most important technological changes of recent times. Because renewable energy is site-dependent, building renewables requires building out the grid. Grid interconnections have become a major bottleneck in new renewable energy construction.

This paper studies the role of missing markets in the interconnection queue as the fundamental causes of this bottleneck. We describe two prominent missing markets: for cost externalities and for the spare capacity of the power grid. These missing markets interact with project cost uncertainty, in equilibrium, to congest the queue. Projects must enter the queue to find out their costs. By doing so, they raise the costs for other potential projects. Because costs are uncertain, projects that receive higher cost estimates linger in the queue, owing to the option value of getting better cost draws for a project. As these projects wait, they raise the cost estimates of other new entrants. The culmination of this congestion equilibrium led CAISO to suspend the queue admission process and appeal to FERC to reform (California Independent System Operator Corporation, 2023).

We use novel, detailed data from CAISO to analyze the importance of these missing markets. We find that missing markets matter a great deal for queue progression and investment: (i) interconnection cost uncertainty is enormous; (ii) projects cause cost externalities

on each other; (iii) interconnection cost estimates affect project advancement and, by implication, real investment; (iv) the free allocation of spare grid capacity (i.e., *transmission plan deliverability*) is the single most important determinant of whether a project connects to the grid. This suite of results shows that the theoretical forces underlying the queue's missing markets are quantitatively important for investment.

In ongoing analysis we use a dynamic equilibrium model built around these empirical results to consider the effects of counterfactual policies that correct for externalities in the queue. Two prominent policies of interest are study charges that incentivize projects to reduce externalities and using an auction, rather than free allocation, to allocate transmission plan deliverability. We aim for this analysis to contribute to an ongoing national discussion of how best to reform the interconnection queue to get needed power onto the grid at a low cost (Federal Energy Regulatory Commission, 2023*a*).

References

Almagro, Milena, Felipe Barbieri, Juan Camilo Castillo, Nathaniel G. Hickok, and Tobias Salz. 2024. "Optimal Urban Transportation Policy: Evidence from Chicago."

Bauer, Sebastian. 2025. "Competition and Welfare in Airport Slot Allocation."

Borenstein, Severin. 2012. "The Private and Public Economics of Renewable Electricity Generation." *Journal of Economic Perspectives*, 26(1): 67–92.

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou. 2024. "Investment in Infrastructure and Trade: The Case of Ports."

Bushnell, James, and Kevin Novan. 2021. "Setting with the Sun: The Impacts of Renewable Energy on Conventional Generation." *Journal of the Association of Environmental and Resource Economists*, 8(4): 759–796. Publisher: The University of Chicago Press.

Butters, R. Andrew, Jackson Dorsey, and Gautam Gowrisankaran. 2025. "Soaking up the Sun: Battery Investment, Renewable Energy, and Market Equilibrium." *Econometrica*, 93(3): 891–927. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA20411.

CAISO. 2022. "Fifth Replacement FERC Electric Tariff (Open Access Transmission Tariff)."

California Independent System Operator Corporation. 2023. "Tariff amendment – interconnection process enhancements 2023 track 1." United States of America, Federal Energy Regulatory Commission and California Independent System Operator Corporation Tariff Filing Docket No. ER23-2058-000.

California State Legislature. 2018. "Senate Bill No. 100 (SB-100): California Renewables Portfolio Standard Program."

CoreLogic. 2021. "Residential Property Parcel Data." Proprietary dataset.

Doraszelski, Ulrich, Katja Seim, Michael Sinkinson, and Peichun Wang. 2025. "Ownership Concentration and Strategic Supply Reduction." *American Economic Review*, 115(3): 903–944.

Elliott, Jonathan. 2025. "Investment, Emissions and Reliability in Electricity Markets." *Mimeo, Johns Hopkins University*.

Federal Energy Regulatory Commission. 2023*a*. "Improvements to generator interconnection procedures and agreements." United States of America, Federal Energy Regulatory Commission Final Rule Order No. 2023. tex.howpublished: 184 FERC ¶ 61,054.

Federal Energy Regulatory Commission. 2023*b*. "Order accepting tariff revisions to suspend cluster admissions after cluster 14." United States of America, Federal Energy Regulatory Commission Order 184 FERC ¶ 61,069.

Fell, Harrison, Daniel T. Kaffine, and Kevin Novan. 2021. "Emissions, Transmission,

and the Environmental Value of Renewable Energy." *American Economic Journal: Economic Policy*, 13(2): 241–272.

Gonzales, Luis E., Koichiro Ito, and Mar Reguant. 2023. "The Investment Effects of Market Integration: Evidence From Renewable Energy Expansion in Chile." *Econometrica*, 91(5): 1659–1693. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA20769.

Gowrisankaran, Gautam, Ashley Langer, and Wendan Zhang. 2025. "Policy Uncertainty in the Market for Coal Electricity: The Case of Air Toxics Standards." *Journal of Political Economy*, 133(6): 1757–1795. Publisher: The University of Chicago Press.

Ifrach, Bar, and Gabriel Y. Weintraub. 2017. "A Framework for Dynamic Oligopoly in Concentrated Industries." *The Review of Economic Studies*, 84(3): 1106–1150.

International Renewable Energy Agency. 2025. "Renewable Power Generation Costs in 2024." International Renewable Energy Agency, Abu Dhabi.

Ito, Koichiro, and Mar Reguant. 2016. "Sequential Markets, Market Power, and Arbitrage." *American Economic Review*, 106(7): 1921–1957.

Jha, Akshaya, and Gordon Leslie. 2025. "Start-up Costs and Market Power: Lessons from the Renewable Energy Transition." *American Economic Review*, 115(2): 690–724.

Johnston, Sarah, Yifei Liu, and Chenyu Yang. 2023. "An Empirical Analysis of the Interconnection Queue."

Marra, Marleen. 2024. "A Market for Airport Slots."

Milgrom, Paul. 1998. "Game theory and the spectrum auctions." *European Economic Review*, 42(3): 771–778.

National Renewable Energy Laboratory. 2024. "2024 Annual Technology Baseline." National Renewable Energy Laboratory, Golden, CO.

National Renewable Energy Laboratory. n.d.. "Physical Solar Model (PSM)."

Rand, Joseph, Nick Manderlink, Steven Zhang, Chris Talley, Will Gorman, Ryan Wiser, Joachim Seel, Julie Mulvaney Kemp, Seongeun Jeong, and Fredrich Kahrl. 2025. "Queued Up: Characteristics of Power Plants Seeking Transmission Interconnection." Lawrence Berkeley National Laboratory.

Seel, Joachim, Kemp Julie Mulvaney, Anna Cheyette, Will Gorman, Naim Darghouth, Dana Robson, Joe Rand, and Seongeun Jeong. 2025. "U.S. Utility-Scale Solar." Lawrence Berkeley National Laboratory.

U.S. Energy Information Administration. 2025*a*. "California state energy profile."

U.S. Energy Information Administration. 2025*b*. "Short-Term Energy Outlook."

Waldinger, Daniel. 2021. "Targeting In-Kind Transfers through Market Design: A Revealed Preference Analysis of Public Housing Allocation." *American Economic Review*,

111(8): 2660–2696.

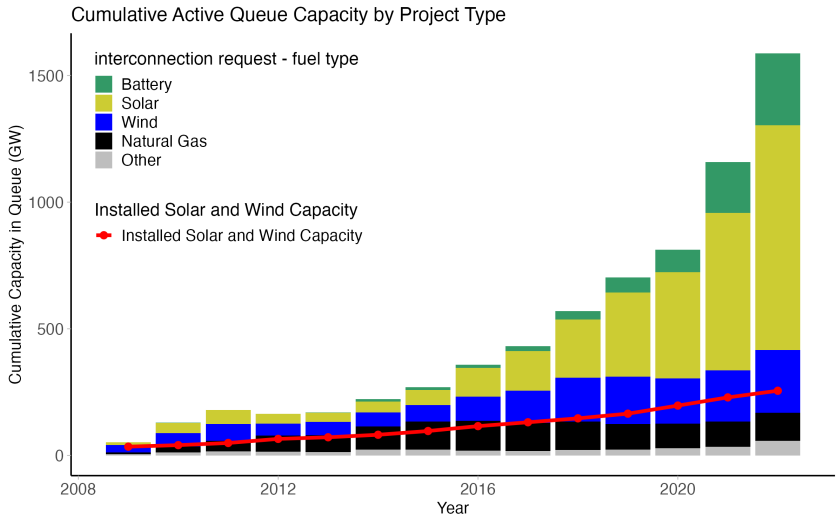
Wiser, Ryan, Dev Millstein, Ben Hoen, Will Gorman, Joe Rand, Joachim Seel, Galen Barbose, Anna Cheyette, Naïm Darghouth, Seongeun Jeong, Julie Kemp, Eric O'Shaughnessy, and Ben Paulos. 2025. "Land-Based Wind Energy Technology Data Update." Lawrence Berkeley National Laboratory.

World Bank, and Technical University of Denmark (DTU). 2025. "Global Wind Atlas."

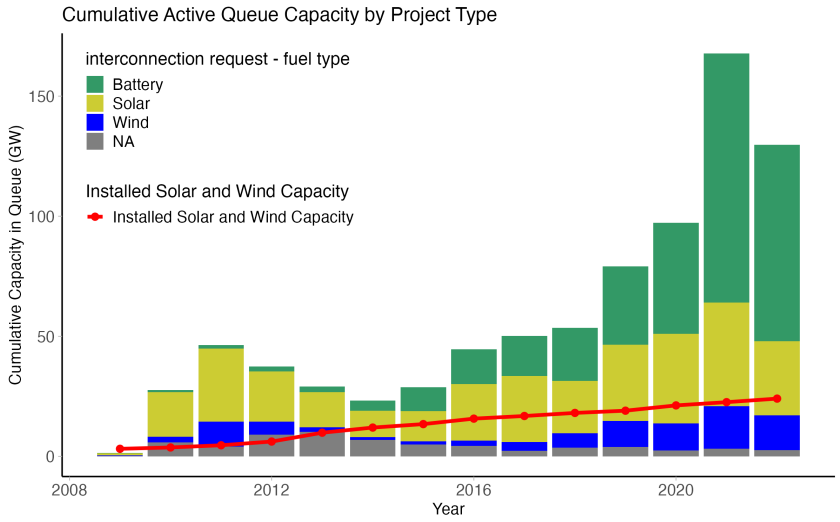
8 Figures

Figure 1: Interconnection Queue Entry by Year

A. All projects in the United States

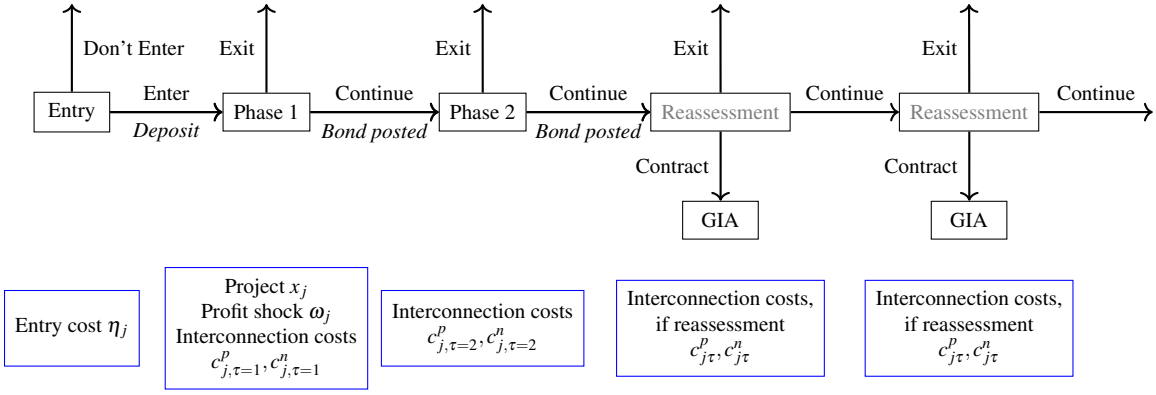


B. Projects in CAISO



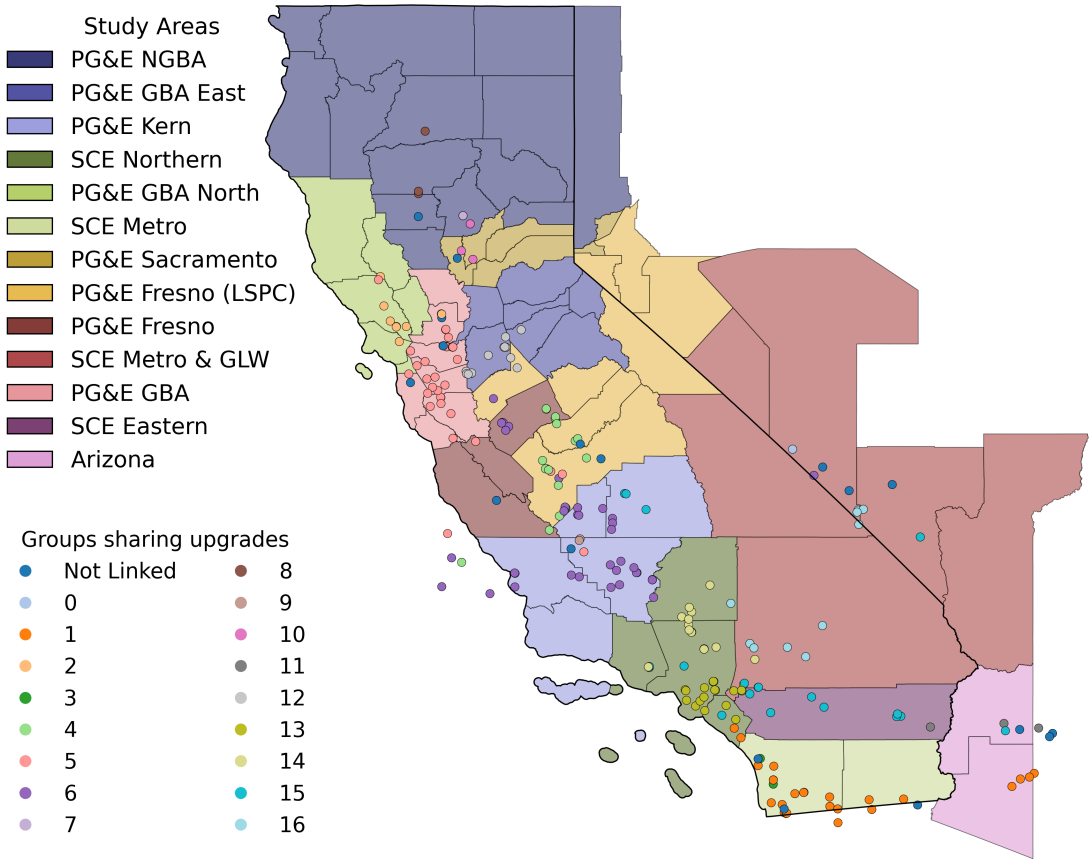
Notes: The top figure shows the total cumulative capacity (GW) of projects in interconnection queues across the U.S. based on data from the LBNL as published in 2023. The red line plots the cumulative capacity of wind and solar projects ($\geq$ 1 MW) across the U.S. for each year. The data come from the EIA Electric Power Annual reports (2019, 2024). The bottom figure shows the total nameplate capacity (GW) of all projects in the California ISO interconnection queue. The red line plots the cumulative installed capacity of solar and wind projects ($\geq$ 1 MW) operational in California, at each year of the queue process. The sample includes all queue requests submitted to CAISO from 2009 through 2022. Installed capacity data are drawn from the California Energy Commission's Electric Generation Capacity & Energy dataset

Figure 2: Interconnection Game Tree for One Project



Notes: The figure shows the extensive form of the interconnection game for a single project. Potential entrants decide whether to enter the queue (i.e., file an interconnection request) or not. In Phase 1 they receive an initial cost estimates and can continue in the queue or exit. In Phase 2 they receive a revised cost estimate and can continue, exit, or contract to connect their generator to the grid. They face the same choices in periods thereafter. If projects continue in the queue, their costs are subject to reassessments as other projects enter and exit, affecting their network cost component.

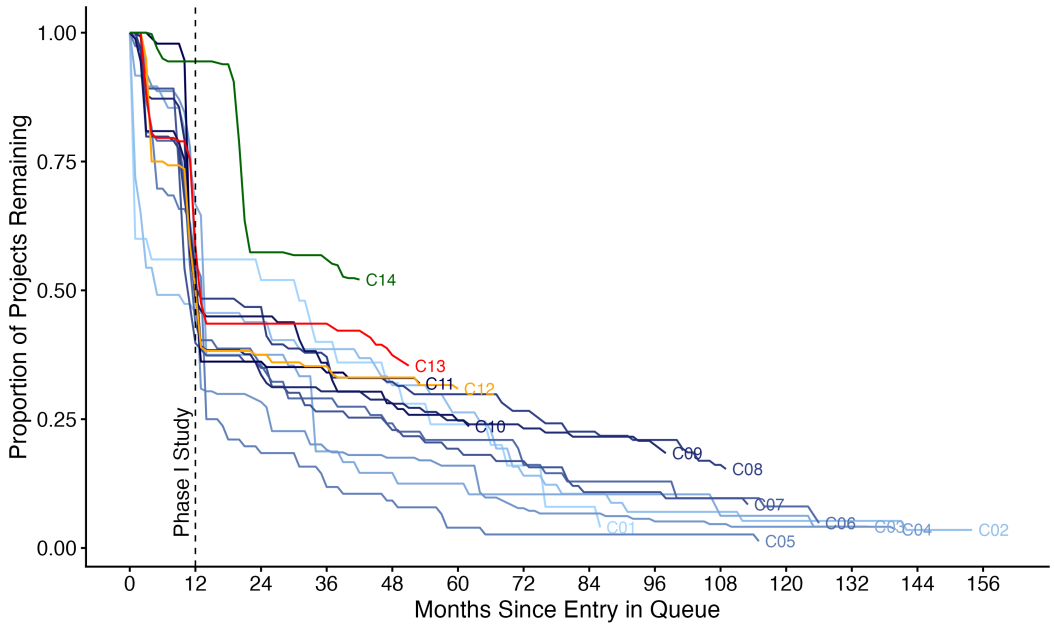
Figure 3: Map of Study Areas and Groups of Projects that Share Network Upgrades (Cluster 14)



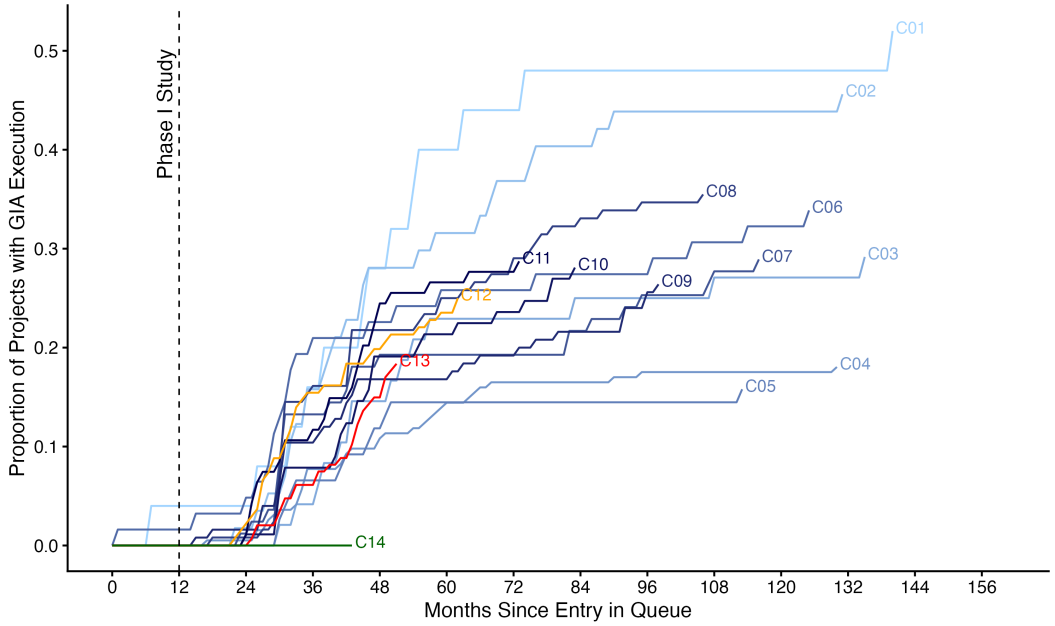
Notes: The figure shows each cluster of projects as a result of using Louvain clustering using points. Points which have no common upgrades with others are considered in the Not Linked cluster. The black outline highlights California's state boundary. Each shaded portion constitutes an area in the model.

Figure 4: Progression through the Interconnection Queue, by Cohort (cluster)

A. Projects remaining in the queue

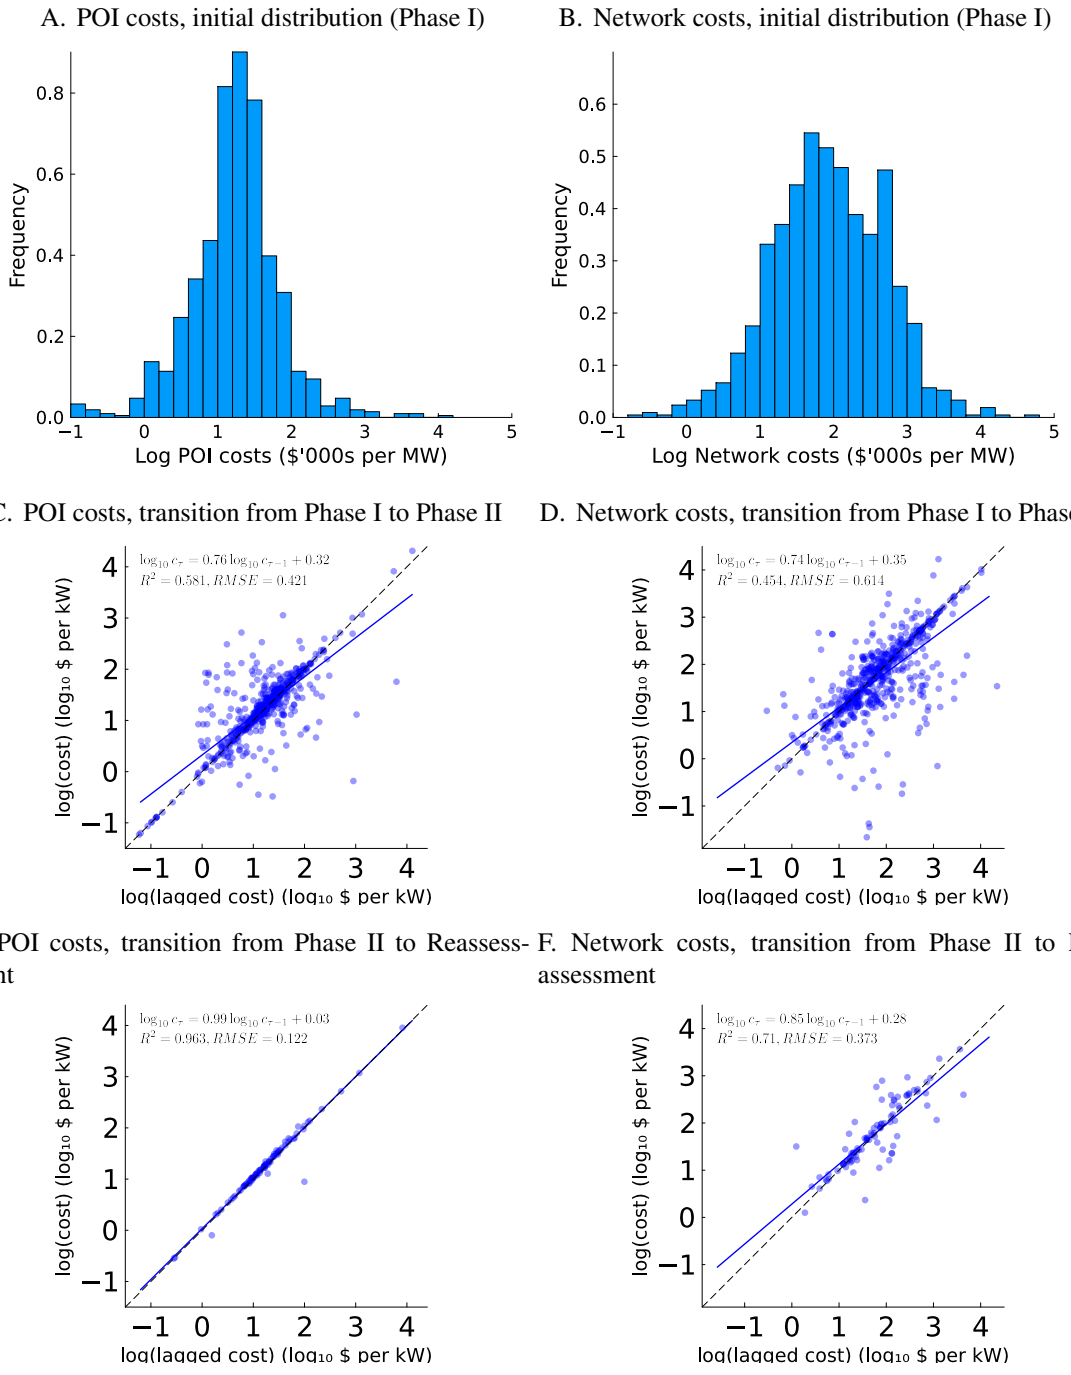


B. Projects signing a Generator Interconnection Agreement (GIA)



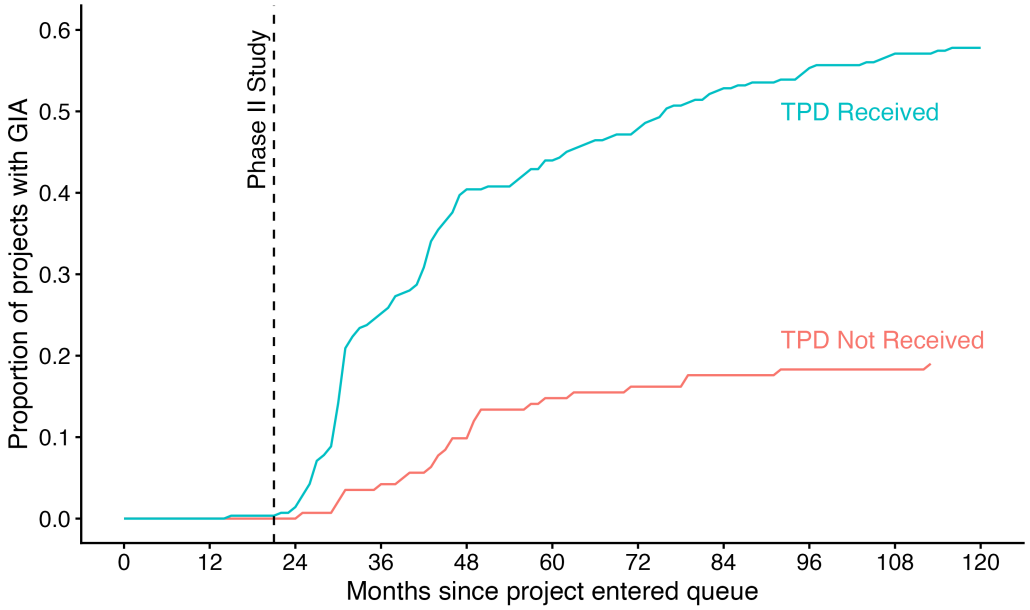
The figure shows the share of projects remaining in the queue (panel A) or signing a Generator Interconnection Agreement (GIA) (panel B) over time. In each panel, each line indicates a separate cohort (which CAISO calls "clusters") of queue entrants, from cluster 1 (entry year 2009) to cluster 14 (entry year 2022). In panel A, the lines show the fraction of projects in that cohort remaining in the queue over time. Projects remain in the queue if they have neither withdrawn (exited) nor signed a Generator Interconnection Agreement (GIA) to connect to the grid. In panel B, the lines show the fraction of projects that have signed a GIA. A GIA is a contract between the project and the grid operator CAISO in which the project agrees to pay interconnection costs in exchange for CAISO connecting the project to the grid.

Figure 5: Evolution of Private Information on Interconnection Costs



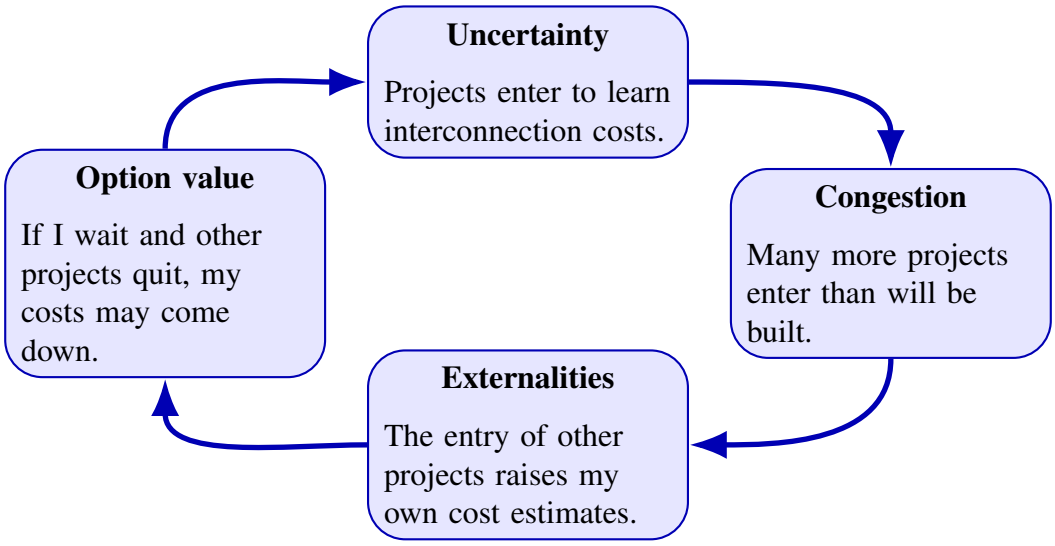
The figure shows the initial cost distribution and its evolution over time for the two different components of interconnection costs. The panels in the left column show Point of Interconnection (POI) costs, the costs paid by one project alone to plug-in to the grid. The panels in the right column show network costs, the costs paid by one project towards transmission network upgrades that may also be shared with other projects. Within each column, the top panel shows the initial distribution of log costs drawn in the Phase I study. The second panel in the column shows a scatter plot of project costs in the Phase II study (vertical axis) against their initial draw in the Phase I study (horizontal axis). The third panel in the column shows a scatter plot of project costs in the first cost reassessment (vertical axis) against their draw in the Phase II study. On the scatter plots we include a dashed 45-degree line and a solid blue linear fit, the equation for which is printed on the panel at top left.

Figure 6: Probability of a Project Contracting to Connect to the Grid by Transmission Plan Deliverability Status



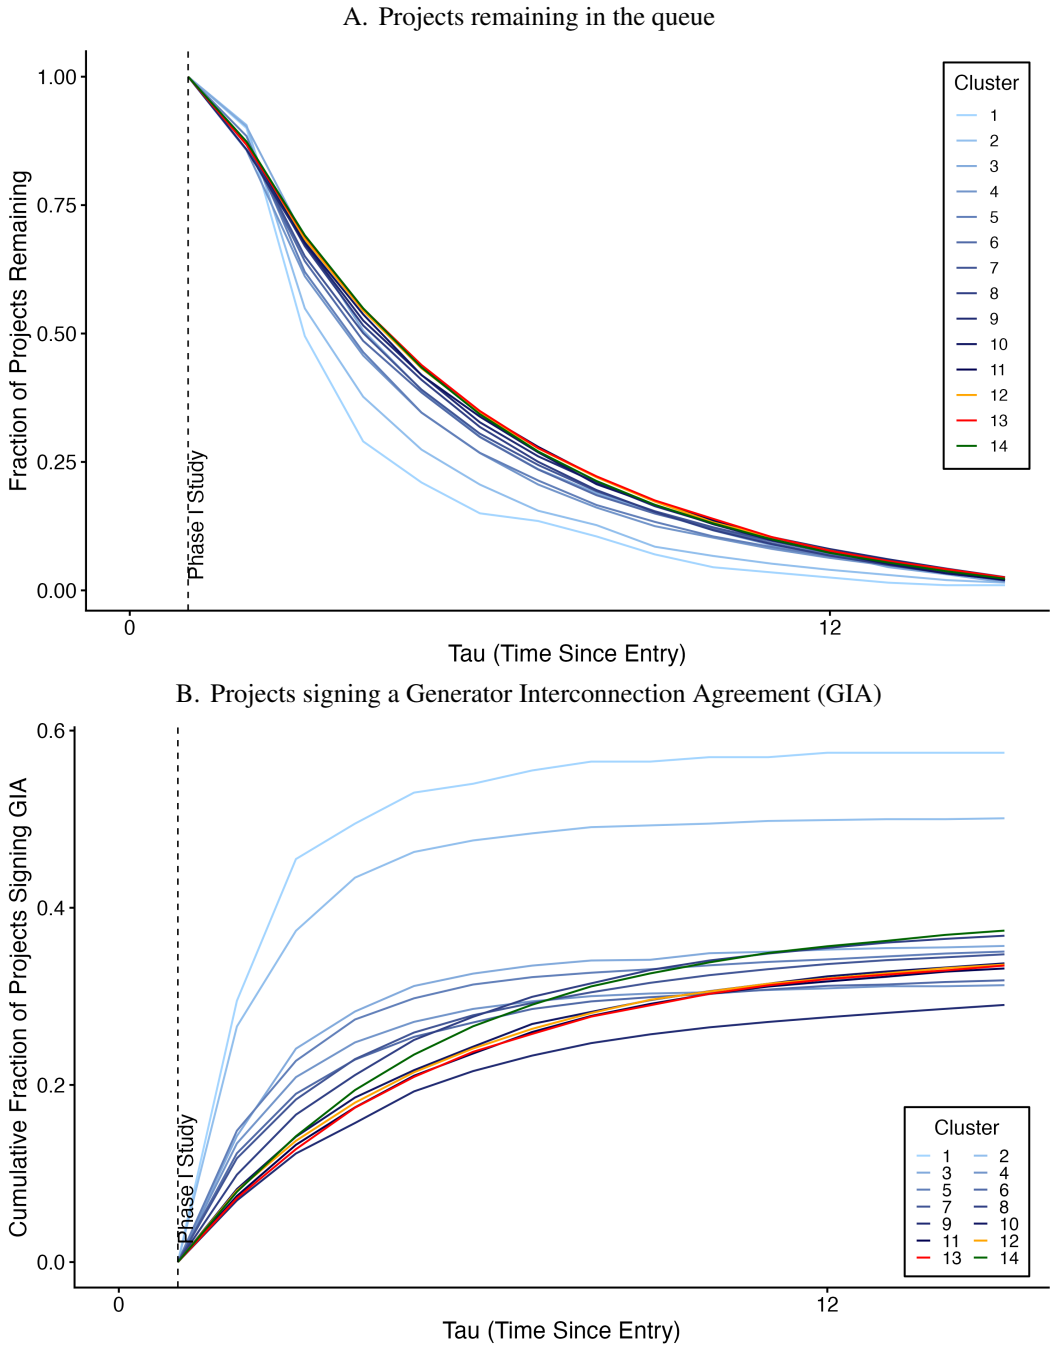
The figure plots the cumulative proportion of projects that have executed a Generator Interconnection Agreement (GIA) against the number of months since the project entered the interconnection queue. A GIA is a contract to connect to the power grid. Separate curves are shown for projects that received Transmission Plan Deliverability (TPD = 1) and projects that did not (TPD = 0). Transmission Plan Deliverability is the right to use some of the spare capacity of the power grid.

Figure 7: Schematic of Equilibrium Congestion



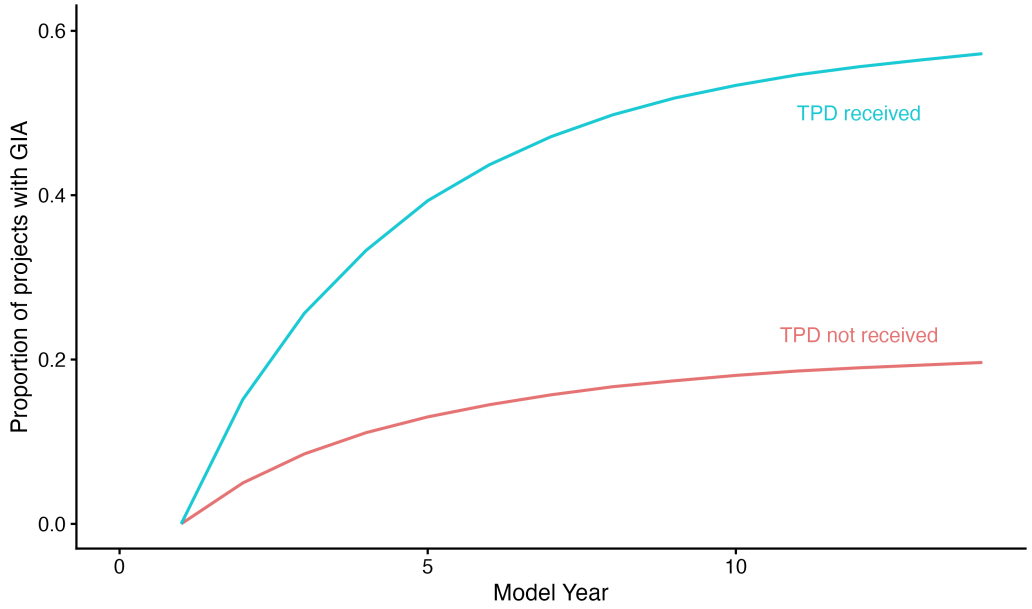
The figure illustrates the forces behind equilibrium queue congestion. Projects enter the queue to learn costs. Entering the queue increases the volume of projects studied. The higher volume of projects increases costs for other projects. These increases in costs further raise the option value of entering and continuing in the queue, exacerbating congestion.

Figure 8: Progression through the Interconnection Queue Simulated with Fitted Model



This figure emulates figure 4 with simulated data, based on preliminary estimates. The sample of unique projects is 99,700 comprised of a single queue of 997 projects, simulated 100 times. It shows the share of projects remaining in the queue (panel A) or signing a Generator Interconnection Agreement (GIA) (panel B) over time in the simulated data. In each panel, each line indicates a separate cluster of queue entrants, from cluster 1 to cluster 14. In panel A, the lines show the fraction of projects in that cohort remaining in the queue over time. Projects remain in the queue if they have neither withdrawn (exited) nor signed a Generator Interconnection Agreement (GIA) to connect to the grid. In panel B, the lines show the fraction of projects that have signed a GIA. A GIA is a contract between the project and the grid operator CAISO in which the project agrees to pay interconnection costs in exchange for CAISO connecting the project to the grid.

Figure 9: Probability of a Project Contracting to Connect to the Grid by TPD Status, Simulated Data



This figure emulates figure 6 with simulated data, based on preliminary estimates. The sample of unique projects is 99,700 comprised of a single queue of 997 projects, simulated 100 times. The model It plots the cumulative proportion of projects that have executed a Generator Interconnection Agreement (GIA) against the number of model years since the project entered the interconnection queue. A GIA is a contract to connect to the power grid. Separate curves are shown for projects that received Transmission Plan Deliverability (TPD = 1) and projects that did not (TPD = 0). Transmission Plan Deliverability is the right to use some of the spare capacity of the power grid.

9 Tables

Table 1: Final initial cost and PI→PII transition regressions (Total congestion)

	Log POI Cost($/kW)	Log Network Cost($/kW)	Log POI Cost($/kW)	Log Network Cost($/kW)
	(1)	(2)	(3)	(4)
Lagged Log POI Cost ($/kW$)			0.542^{***}	
			(0.055)	
Lagged Log Network Cost ($/kW$)				0.602^{***}
				(0.045)
Log AGW_τ	0.250^{*}	0.815^{***}	-0.159^{***}	-0.287^{***}
	(0.137)	(0.273)	(0.043)	(0.058)
MW (Own Capacity)	-0.002^{***}	-0.0003	-0.0004^{**}	-0.00002
	(0.0003)	(0.0003)	(0.0002)	(0.0003)
FCDS (=1)	0.024	0.028	0.538^{***}	0.667^{***}
	(0.201)	(0.199)	(0.144)	(0.187)
Wind (=1)	0.164	0.180	-0.176	-0.243
	(0.150)	(0.260)	(0.173)	(0.261)
Battery (=1)	0.340^{***}	-0.038	-0.220^{**}	-0.226
	(0.088)	(0.125)	(0.095)	(0.141)
Co-located (=1)	-0.171	-0.375	0.194	0.013
	(0.168)	(0.232)	(0.130)	(0.188)
Voltage 100–230 kV (=1)	0.316	-0.089	-0.659^{***}	-0.394^{**}
	(0.193)	(0.220)	(0.153)	(0.178)
Voltage 500+ kV (=1)	0.448^{***}	0.252^{**}	-0.081	-0.077
	(0.090)	(0.113)	(0.089)	(0.108)
Cluster fixed effects	Yes	Yes		
Area fixed effects	Yes	Yes		
R^2	0.278	0.184	0.397	0.423
N Clusters	15	15	15	15
N Areas	14	14	14	14
N Projects	1061	1061	646	645

The dependent variables are log(cost + 1) for POI and network costs, expressed in dollars per kW. Dependent variables include project characteristics: project capacity in megawatts (MW), an indicator for Full Capacity Deliverability Status (FCDS), indicators for Wind and Battery (with Solar omitted), an indicator for co-located projects, and voltage-group dummies; fixed effects are included as indicated in the table. Transition columns model Phase II costs and include the lagged log cost from Phase I as an AR(1) term. Queue congestion regressors (AGW) are defined at the area–year level in gigawatts: AGW Entering is the annual flow of new capacity entering the queue; AGW Backlog is the stock of active capacity remaining in the queue; Standard errors are HC1-robust. ***Significant at the 1 percent level. **Significant at the 5 percent level. *Significant at the 10 percent level.

Table 2: Linear Probability Model for Reaching Phase 2

	Dep. Var.: Reaching Phase 2			
	(1)	(2)	(3)	(4)
Phase 1 network cost ($/W)	-0.018^{**}	-0.026^{***}	-0.026^{***}	
	(0.011)	(0.008)	(0.008)	
Phase 1 POI cost ($/W)	-0.039	-0.027	-0.026	
	(0.076)	(0.064)	(0.065)	
Phase 1 network cost quartile 2				-0.101^{***}
				(0.039)
Phase 1 network cost quartile 3				-0.251^{***}
				(0.039)
Phase 1 network cost quartile 4				-0.408^{***}
				(0.043)
Phase 1 POI cost quartile 2				-0.084^{**}
				(0.040)
Phase 1 POI cost quartile 3				-0.080^{*}
				(0.042)
Phase 1 POI cost quartile 4				-0.069
				(0.045)
Project x_j Controls	Yes	Yes	Yes	Yes
Cluster fixed effects		Yes	Yes	Yes
Area fixed effects			Yes	Yes
R^2	0.059	0.176	0.191	0.267
Projects	1100	1095	1091	1091
Clusters		13	13	13
Areas			14	14

This table reports coefficients from regressions of reaching Phase 2 on Phase 1 network and POI costs. Full capacity deliverability status (FCDS) means that the generator is requesting that its entire output be delivered to the grid even under peak load conditions. Project x_j controls includes fuel type, capacity, indicators for various voltage levels, Site Exclusivity and FCDS. The data come from CAISO's event checklist, resource summary page, scraped cost PDFs, and TPD spreadsheets. We limit our sample to projects with Phase 1 costs and projects in clusters 3-14. Statistical significance at certain thresholds is indicated by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Linear Probability Model for Receiving GIA Within 4 Years

	Dep. Var.: GIA within 4 years of entry		
	(1)	(2)	(3)
TPD received (=1)	0.281^{***}	0.312^{***}	
	(0.048)	(0.059)	
Network low cost ($60/kW$)			-0.037
			(0.065)
TPD received (=1) $\times$ Network cost Quartile 1			0.288^{***}
			(0.088)
TPD received (=1) $\times$ Network cost Quartile 2			0.418^{***}
			(0.084)
TPD received (=1) $\times$ Network cost Quartile 3			0.283^{***}
			(0.089)
TPD received (=1) $\times$ Network cost Quartile 4			0.230^{**}
			(0.092)
Site Exclusivity (=1)	0.098	0.079^{*}	0.071
	(0.043)	(0.044)	(0.044)
Project x_j Controls	Yes	Yes	Yes
Cluster fixed effects		Yes	Yes
Area fixed effects		Yes	Yes
Average GIA Receipt (No TPD)	0.207	0.207	0.207
R^2	0.157	0.226	0.23
N Projects	439	433	433

This table reports coefficients from regressions of GIA receipt within 4 years on TPD receipt, an indicator for Phase 2 network costs greater than \$80/kW, an indicator for Phase 2 network costs less than \$80/kW, and their interaction with TPD receipt.A Generator Interconnection Agreement (GIA) is a pro forma contract among a party requesting interconnection, the CAISO, and the Participating TO that owns the transmission facility with which the requesting party wishes to interconnect. Transmission plan deliverability (TPD) means that the project was allocated a portion of the available capacity on the transmission network. Full capacity deliverability status (FCDS) means that the generator is requesting that its entire output be delivered to the grid even under peak load conditions. Project x_j controls includes fuel type, capacity, indicators for various voltage levels, and FCDS.The data come from CAISO's event checklist, resource summary page, scraped cost PDFs, and TPD spreadsheets. We limit our sample to projects with Phase 2 costs and projects in clusters 3-14. Statistical significance at certain thresholds is indicated by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Model parameters

Step	Variable	Description of parameters
Panel A. Parameters estimated outside the dynamic game		
Private states	$\delta^{net,init}, \sigma_v^{net,init}$	Initial network cost estimate
	$\delta^{poi,init}, \sigma_v^{poi,init}$	Initial POI cost quote
	$\delta^{net}, \sigma_v^{net}$	Network cost transitions
	$\delta^{poi}, \sigma_v^{poi}$	POI cost transitions
	δ^{tpd}	Transmission plan deliverability allocation
Aggregate state	δ^{agw}, σ_v^{agw}	Queue state transition
Ancillary	$\phi_{k0}, \phi_{k1}, \sigma_{kF}^2$	Capital cost process for wind, solar and batteries
Panel D. Parameters estimated within the dynamic game		
Dynamic	β_π	Coefficients on observables in profit equation
	σ_ω	Scale of unobserved profitability shock
	σ_ε	Scale of action-specific payoff shock
Panel C. Entry stage parameters		
Entry	σ_η	Scale of action-specific entry cost shock

Table 5: Entry Stage Parameter Estimates

	Coefficient	Std. Error
Solar mean entry cost	1.720^{***}	0.056
Inverse solar entry cost scale	7.188^{***}	0.640
Wind mean entry cost	2.381^{***}	0.190
Inverse wind entry cost scale	3.559^{***}	0.815
Battery mean entry cost	2.182^{***}	0.056
Inverse battery entry cost scale	12.352^{***}	0.660

Notes: Maximum likelihood estimates of entry cost distribution parameters. The entry cost for a project of technology type k is drawn from $N(\mu_k/\sigma_k^\eta,\ 1/(\sigma_k^\eta)^2)$, measured in billions of dollars. Implied mean entry cost equals μ_k/σ_k^η and implied standard deviation equals $1/\sigma_k^\eta$. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 6: Interconnection Stage Parameter Estimates

	Coefficient	Std. Error
Scale parameters		
Inverse scale of profit shocks	3.054^{***}	0.636
Scale of persistent profit shock	-0.001	828.658
Common payoffs		
Transmission Plan Deliverability (0/1)	1.217^{***}	0.292
Land cost (1/acre)	2.800	3.025
Capital cost	-0.608^{***}	0.129
Technology-specific profits		
Solar	-2.345^{***}	0.499
Wind	-1.753^{***}	0.591
Battery	-0.634^{**}	0.321

Notes: Quasi-maximum likelihood estimates of interconnection stage payoff parameters. All coefficients are measured in millions of dollars per σ^{ϵ} per MW. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 7: Model Fit

Statistic	Data	Simulation
N projects	997	997
Avg entry year	2017.6	2017.6
Avg years in queue	2.05	2.41
Share GIA $\leq$4 yrs (%)	13.3	13.9
Avg initial POI cost ($mn)	53.8	53.8
Avg min POI cost ($mn)	48.9	34.4
Avg initial network cost ($mn)	276.4	276.4
Avg min network cost ($mn)	234.2	160.6
Share receiving TPD (%)	44.3	46.9

Online Appendix

Plugging in the Wind and Sun: Market Design for Shared Infrastructure Under Uncertainty

Gautam Gowrisankaran, Ashley Langer and Nicholas Ryan

A Appendix: Data

A.1 Data from CAISO

A.2 Other Data Sources

Capital costs of renewable installations.— We use as an input to our model the capital installation costs for the three main technologies entering the queue: utility-scale solar photovoltaic (PV), onshore wind, and battery storage. Because our model is forward-looking we need both historical cost data and expectations of future cost trajectories.

We use three sources of cost data. From the Lawrence Berkeley National Laboratory, we use data on total fixed installation costs (including capital, balance-of-systems and development costs) for utility-scale solar PV and for onshore wind. From the International Renewable Energy Agency (IRENA), we use analogous capital costs for battery storage. Finally, from the National Renewable Energy Laboratory (NREL), we take projections of these cost series out to 2035 (*Annual Technology Baseline (ATB) 2024*).

We process these data to splice together a continuous and monotonic cost series for each technology. First, we smooth the series by interpolating some spikes (such as 2022 storage costs). Second, we project costs for future years in which projections are not reported in source data using the function

$$\log(\text{CapEx}_t) = a \cdot e^{-b(t-2035)} + c. \tag{8}$$

This functional form captures rapid early-stage cost declines that gradually flatten as technologies mature. Parameters are estimated using nonlinear least squares on a data set including the historical data series and the 2035 projection from NREL. We then use the fitted models to generate annual projections for 2025–2039. Figure A1 shows these projections alongside the original cost series. The fit of this simple functional form is remarkably good.

A.3 Definition of study areas

Study areas were defined by combining network-based clustering with basic geographic grouping. We started with the PTO regions published by CAISO (for example `PG&E GBA`, `PG&E Fresno`, `PG&E Kern`, `SCE Metro`, and others), which gave us an initial sense of how the state is usually divided for planning. These PTO regions do not have precise county boundaries, so we used them mainly as a starting point and then refined them based on how projects were actually connected.

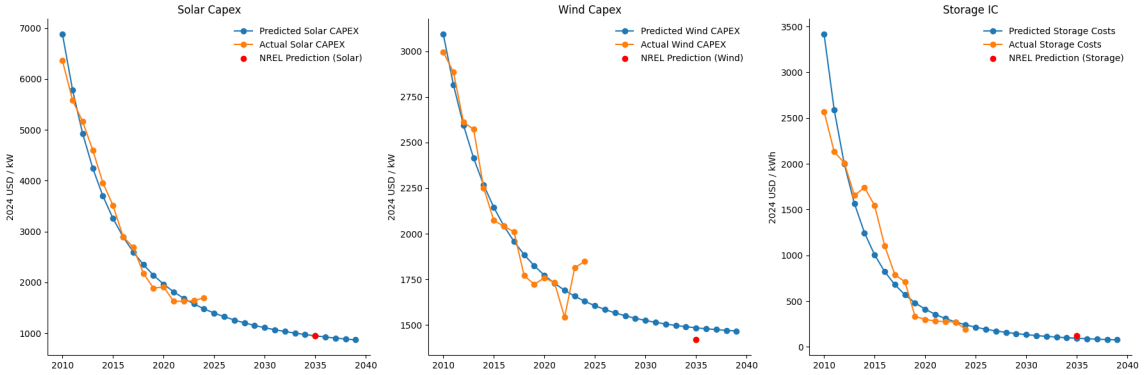


Figure A1: Capital Cost Time Series

To understand which projects were linked, we used the itemized upgrade lists from the Phase 1 cost studies. Because upgrade names are not consistent across projects, we used fuzzy string matching to identify when two upgrade descriptions referred to the same substation, transformer bank, or line. We only treated the upgrade as "shared" if the matched component represented more than 20% of both projects' total upgrade cost, so minor overlaps did not create connections. If this threshold was met, we recorded a link between the two projects in our adjacency matrix. Running Louvain clustering on this matrix grouped projects into clusters that shared major upgrade dependencies.

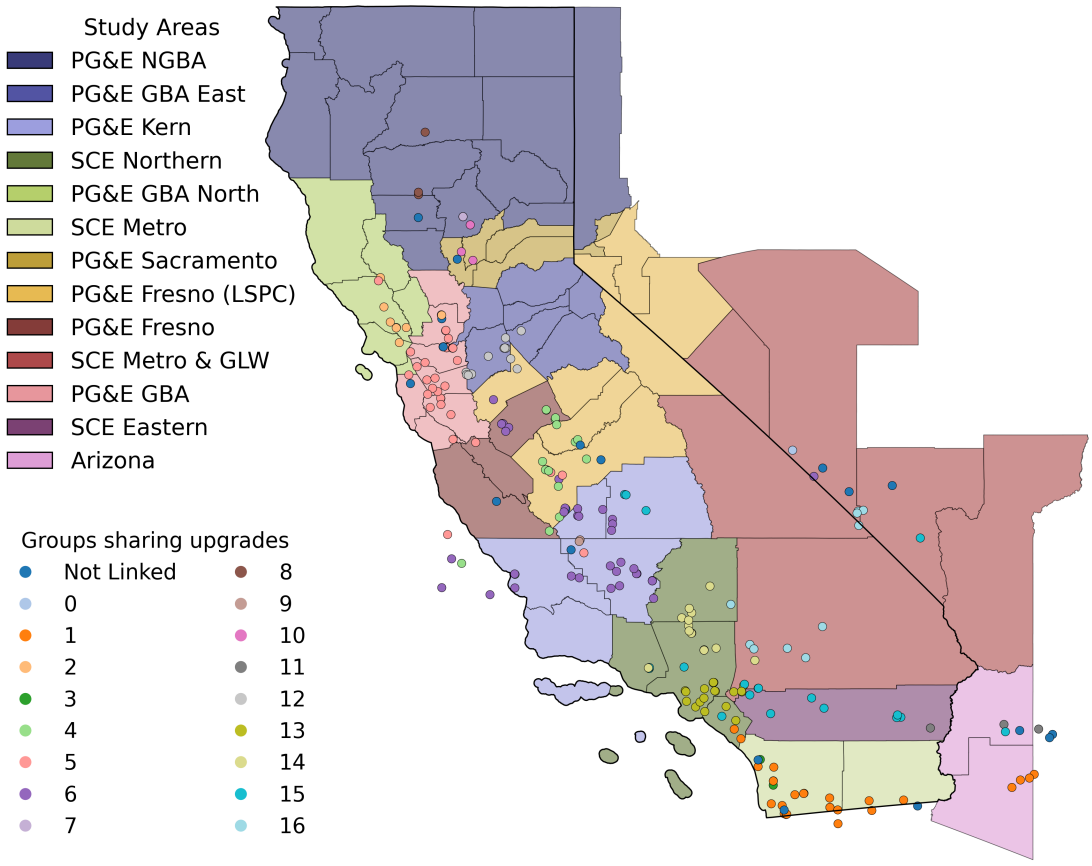
Cluster and Zone Methodology.—Once clusters were identified, we matched each project to its coordinates and county and used counties as the base unit for defining zones. Counties that contained multiple clusters, or counties that were adjacent and shared clusters, were grouped together to form larger areas. Counties without active projects or isolated single-project clusters were given simple placeholder zone IDs.

We then compared these cluster-based groupings to the original CAISO PTO regions and adjusted the boundaries where needed. In most places the clusters lined up well with the PTO structure, but in a few cases we made manual edits to keep the zones consistent with both the geography and the network connections. A clear example is Kern County: the clustering produced two distinct groups of projects on opposite sides of the county that did not share upgrades with each other. To reflect this, we split Kern into two study areas. In other locations where clusters were very fragmented, we merged adjacent counties to create a more coherent zone.

Overall, the final zones were created by starting from CAISO's PTO regions, refining them using the cluster results from the upgrade network, and making small manual adjust-

ments where the clustering and geography did not fully align.

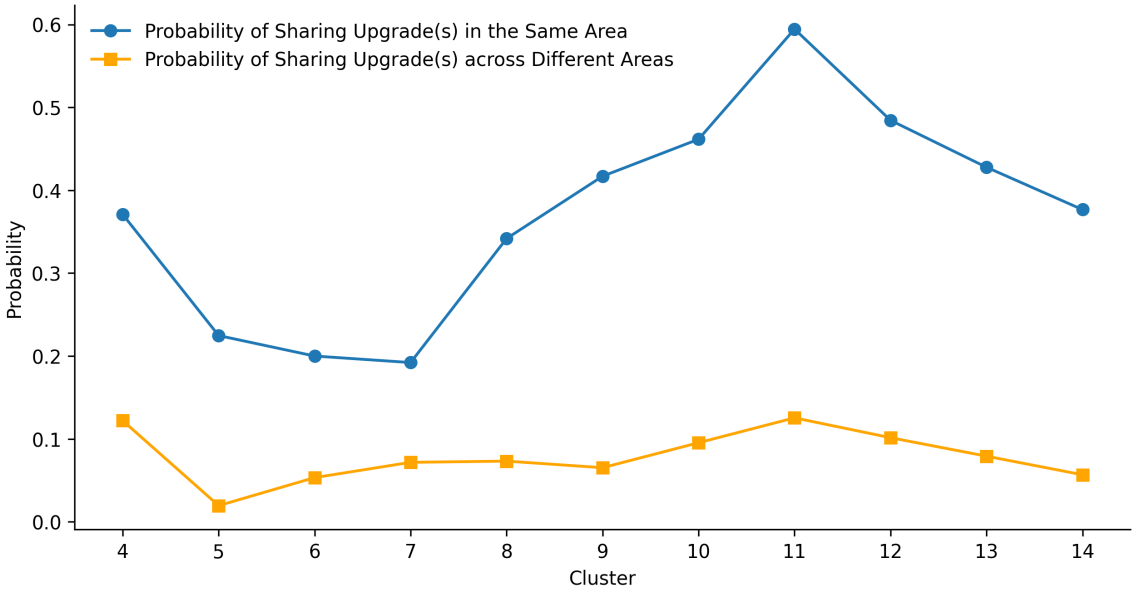
Figure A2: Correspondence of Study Areas to Shared Network Upgrades (Cluster 14)



Notes: The figure shows each cluster of projects as a result of using Louvain clustering using points. Points which have no common upgrades with others are considered in the Not Linked cluster. The black outline highlights California's state boundary. Each shaded portion constitutes an area in the model.

Validation of area definition.—With our detailed data on interconnection costs we are able to validate the definition of study areas above. The raw cost data that we use is observed at the level of the upgrade, an individual piece of transmission or substation infrastructure that is required to connect a project to the grid. Some pieces of infrastructure are used by multiple projects and therefore shared. We can observe sharing in this data when two different projects are both assigned part of the costs for one network upgrade. For example, if two projects both contribute to congesting a transmission line, they both will be assigned part of the cost of an upgrade to that line.

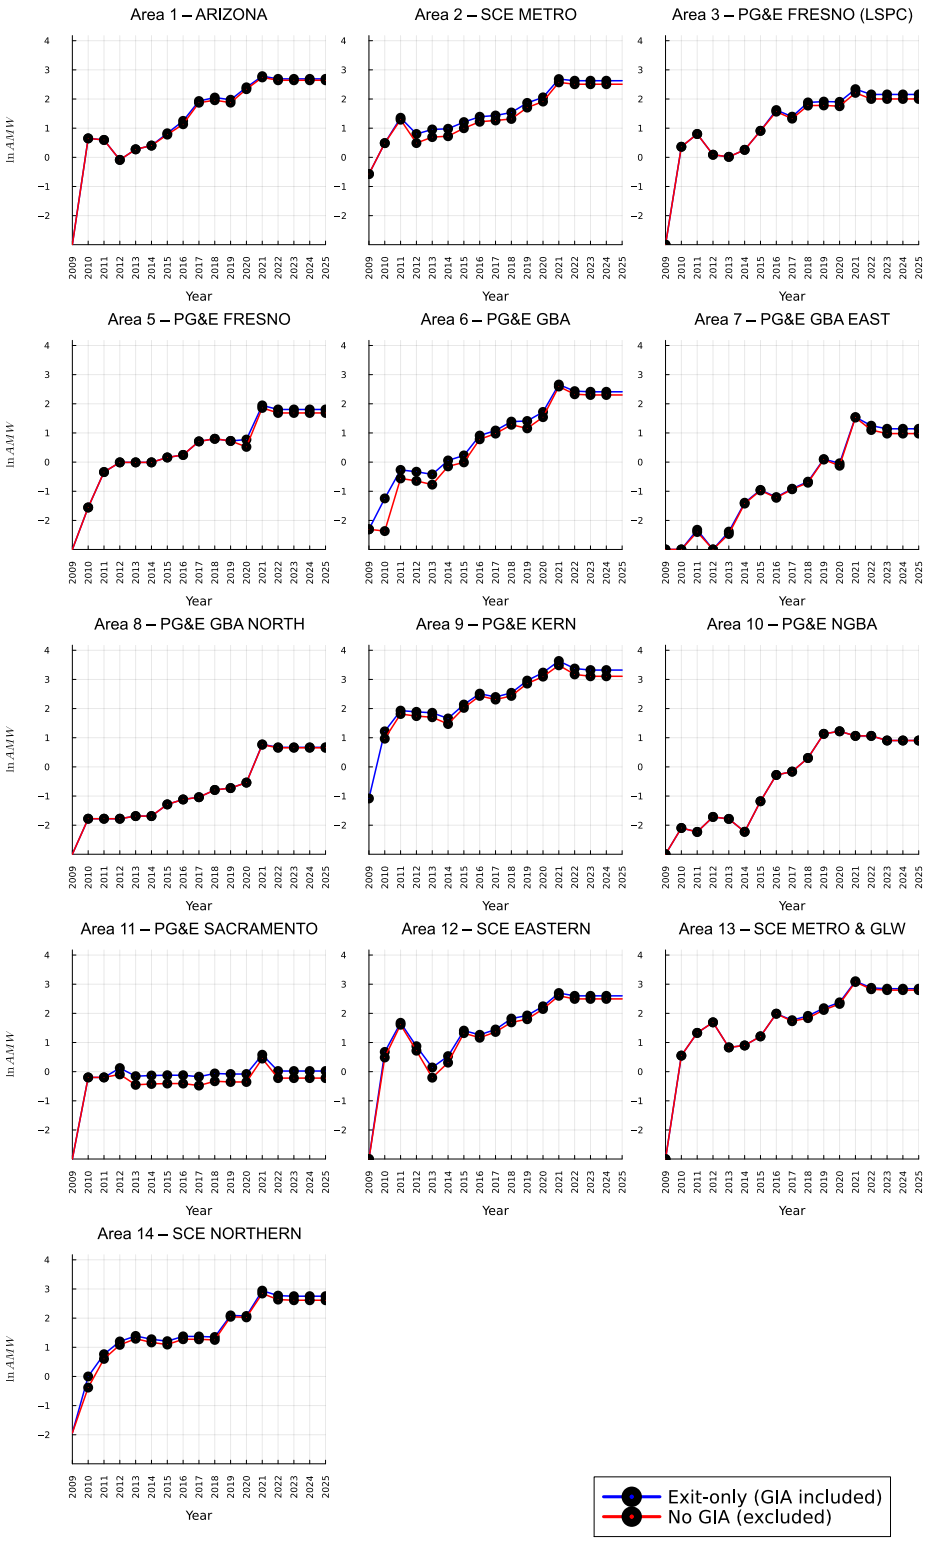
Figure A3: Probability of Projects Sharing Upgrade Conditional on Areas



Notes: This figure plots the probability that two energy projects share a transmission upgrade, conditional on geographic area. The blue line reports the within-area linkage probability, while the orange line reports the cross-area linkage probability. Two projects are defined as sharing a link if either project contributes at least 5% of the total cost to the same upgrade. Clusters 1–3 contain too few projects to reliably estimate these probabilities, resulting in unstable estimates; these clusters are therefore excluded from the figure.

Figure A3 shows the probability that a project shares an upgrade with a randomly-selected project either in the same study area (blue line) or a different study area (orange line). The probability of sharing an upgrade with a project in the same study area fluctuates over time, but is typically around 0.40. By contrast, the probability of sharing an upgrade with a project in a different area is typically around 0.10 or below. (Even distant projects can share upgrades when some projects, like long-distance transmission lines, are critical for many different projects in the grid.) This finding validates our assumption that congestion externalities are much more likely to operate within a study area. It also shows that our definition of study area is a good proxy for the geographic scope of these externalities.

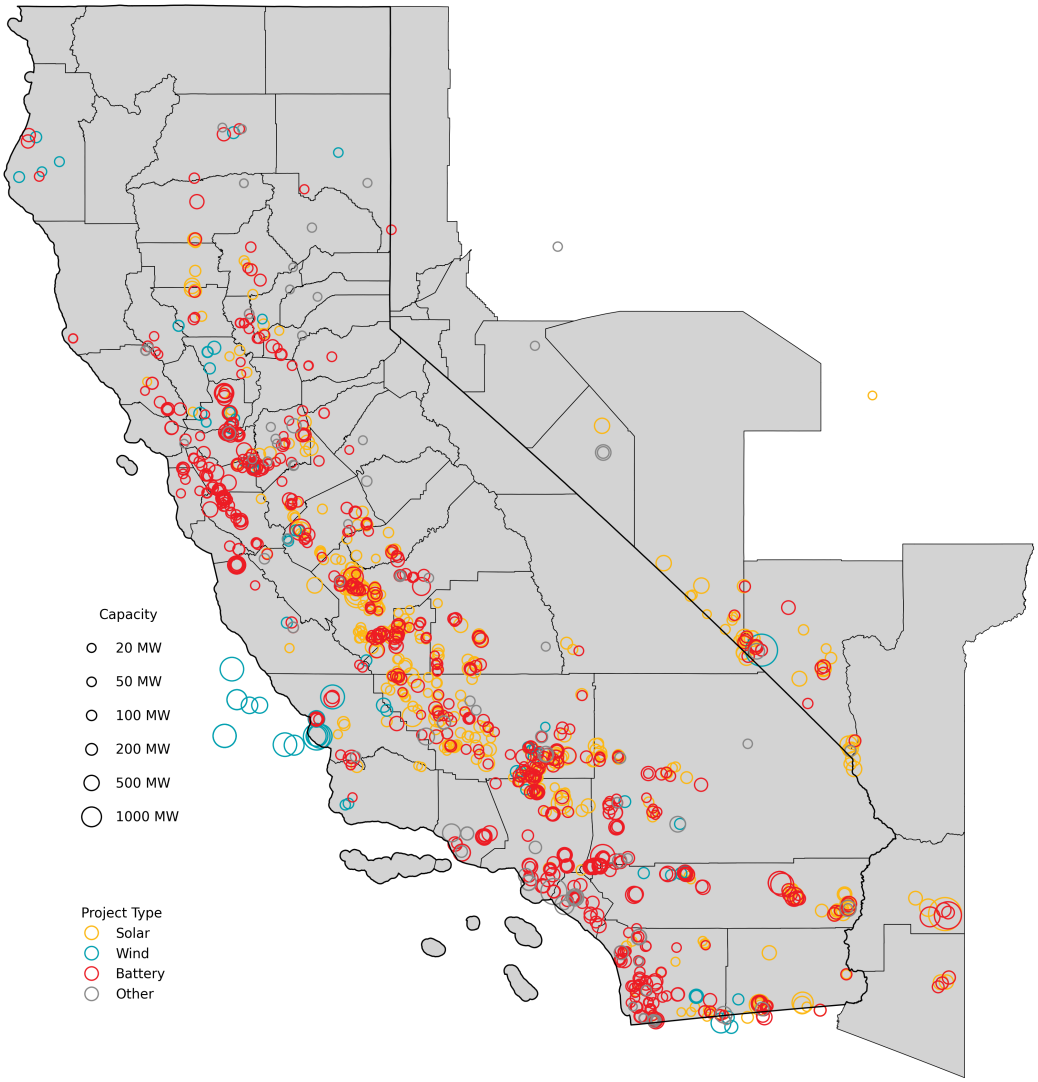
Figure A4: AMW by Area



Notes: *AMW* is the active stock of MW in an area in the queue up to calendar year t. The figure plots the trends in queue congestion ($\ln AMW$) evolution as the queue progresses over time.

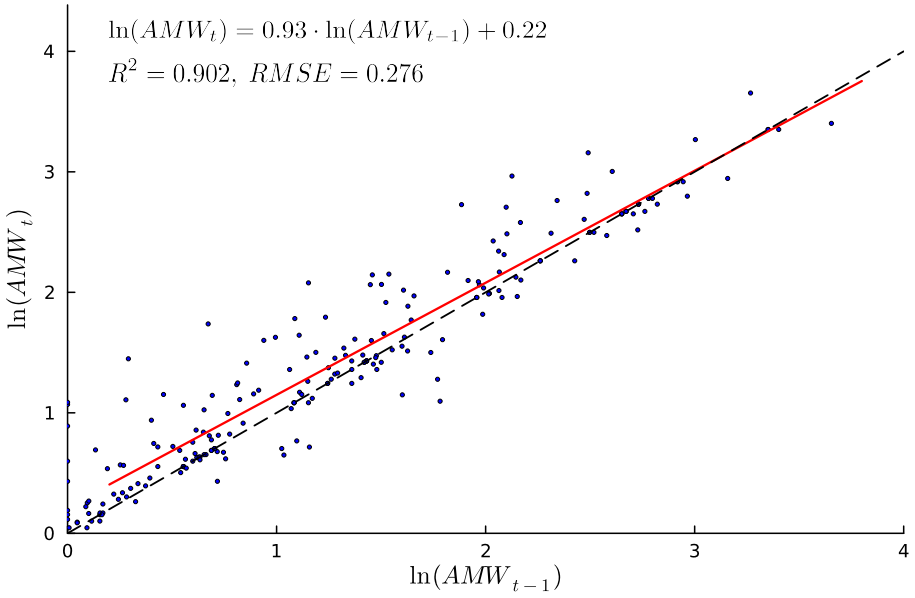
B Appendix: Supplementary Descriptive Analysis

Figure B5: Locations of Projects in the Interconnection Queue, Clusters 1 to 14



Notes: The figure shows capacity, type, and location for projects proposed to be linked to the CAISO grid. The black outline highlights California's state boundary. Each point is scaled by their capacity at the point of interconnection. The sample includes all valid coordinates of the project from CAISO's resource summary, and from scraped files of electricity projects.

Figure B6: AR(1)



Notes: AMW_t is the active stock of MW in the queue up to calender year t. The figure shows the persistence in AMW_t

C Appendix: Conceptual model

We consider the problem of energy projects that want to connect to the grid, focusing on one area of California.[14] We denote projects, $j = 1, \ldots, J$. In this simple version of the model, each individual project, j, faces three stages to this connection process.

First, project j learns its type (solar, wind, or battery) and receives a one-time opportunity to enter the interconnection process. If it enters, it must pay an entry cost, η_j, that is drawn from a distribution $N(\mu, \sigma_\eta)$. Upon entry, the project receives its characteristics, x_j, that indicate its size in MWs and other relevant information. Second, the ISO assesses the project's per MW interconnection costs, c_j, and decides what benefits to offer the project, principally TPD_j in our case. It reports these variables to the project, and they ultimately determine the interconnection fees that the project will pay to the ISO if it chooses to connect, $MW_j \times ICF(c_j)$, where MW_j is one of the characteristics in x_j. Finally, the project obtains a private information profit shock ε_j and then decides whether or not to build its generation facility and connect to the grid. The project knows that, if it connects to the grid, it will earn net profits per MW from this point on of $\pi_j = \bar{\pi}(x_j, TPD_j) - ICF(c_j) + \varepsilon_j$, where $\overline{\pi}$ is the profits from operating in electricity generation and capacity markets.

The central complication of the interconnection process is that projects' interconnection costs are interdependent and that, in an era when most new projects are renewables or batteries, there are many potential entrants. The problem facing the ISO is therefore how to map from assessed interconnection cost quotes c_j to fees, ICF_j, understanding that this mapping will affect which projects connect. This is complicated because only a subset of projects will actually connect, and projects may impose heterogeneous marginal costs.

Consider first the simple case where the stages happen simultaneously across projects, so that the ISO can observe the full set of projects that have chosen to enter before providing projects interconnection fee quotes. In this case, c_j is the size-weighted average interconnection cost over the set of projects $\mathscr{J} \subseteq \{1, \ldots, J\}$ that have chosen to enter the queue. The ISO then needs to evaluate how to map the combined total costs of all projects that enter, $\mathscr{C}(\mathscr{J})$, to interconnection fees. A simple mapping might just be to charge each project its average costs, $ICF(c_j) = c_j$. However, it could instead use marginal costs, in which case $ICF_j(\mathscr{J}) = (\mathscr{C}(\mathscr{J}) - \mathscr{C}(\mathscr{J} \setminus j))/MW_j$.

[14] Our estimation considers multiple areas $a = 1, \ldots, A$ but we suppress the a notation here in the interest of brevity.

These simple mechanisms have substantial problems because of the interdependencies between projects, and because projects choose whether to connect after receiving their fee quotes. For example, consider two solar farms in the Mojave Desert that would share interconnection infrastructure to bring their power to Los Angeles. On the one hand, it is possible that a transmission line is needed if either project connects. In this case, each project's marginal cost is small. If CAISO assigns each project a fee equal to its marginal cost, then the sum of the fees will not cover the cost of the line. On the other hand, the line may only be necessary if both projects connect. In this case, each project's marginal interconnection cost is the entire line cost. Here, fees based on marginal costs would underincentivize connecting since total fees would exceed total costs. Using average costs to set fees has different issues that are equally problematic. For instance, if one project were substantially more profitable than the other, the less profitable project may not have sufficient profits to always connect. Even if the more profitable project always connects, this may mean that collected fees are insufficient to cover costs, and the ISO may need to raise consumer rates to cover the difference.

The fact that projects can choose whether to connect after receiving their fee quote is what makes this process difficult because this decision can change both the marginal and average costs of remaining projects. However, if the ISO did not allow projects to drop out, then many of them may not want to enter in the first place, since the interconnection fees that they would have to pay are uncertain and may imply negative profits. Overall, this results in the interconnection process being a very challenging mechanism design problem. Yet, the process is hugely important in ensuring that we achieve the next energy transition as efficiently as possible. We next consider different mechanisms that have been proposed.

C.1 Original Approach: Sequential Interconnection Queues

The original approach to the interconnection process was queueing, which was used by different ISOs, such as PJM. Unlike our above example, the queueing process was designed during an era with fossil fuel generation projects, when there were many fewer projects and hence externalities across projects were less salient. Once a project enters the queue, the ISO assesses interconnection costs based on the additional infrastructure that is needed taking the existing grid capacity as given, but not accounting for potential entrants behind it in the queue. Thus, to consider the queueing approach, we specify a different timing of our above model where the three stages of each project are sequential.[15]

[15] In reality, these stages for projects can be overlapping, but this does not change our main insights.

To understand the issues with the queueing approach, consider again the case of two potential projects that may receive opportunities to enter the interconnection queue and connect to the grid. For simplicity, assume that the projects are ex ante symmetric. Consider the case where the transmission line needs to be built if either or both projects enter. In this case, $\mathscr{C}(\{1\}) = \mathscr{C}(\{2\}) = \mathscr{C}(\{1,2\})$.

Consider the case where project 1's idiosyncratic opportunity to enter arises before project 2's opportunity. In this case, the ISO would quote project 1 a fee equal to the entire cost of the transmission line, i.e., $ICF_1 \times MW_1 = \mathscr{C}(\{1,2\})$. This would then lead project 1 to have too little of an incentive to enter the queue and build its generation capacity conditional on entering. Moreover, if project 2 then receives a subsequent opportunity to enter the queue, it would also have too little an incentive to enter in the case where project 1 declined to enter.

Overall, the sequential nature of the queueing process led to substantial hold-up in building transmission infrastructure and thereby slowed the green energy transition. While this simple analysis considers the case of purely sequential projects, if the three stages for the projects overlap, this complicates any mathematical analysis of the problem but fundamentally does not solve the hold-up issues.

C.2 Updated Approach: Cluster Analyses

Understanding the basic point that the interconnection queue results in hold-up, the Federal Energy Regulatory Commission (FERC) recommended that all ISOs switch to analyzing projects in clusters. CAISO implemented its cluster process in 2008 and was at the forefront of the move from sequential interconnection queues to cluster analyses.

The idea of a cluster analysis is that, by assessing the projects jointly, the ISO can address the externalities across projects, which can reduce the hold-up problem. The timing of the model then changes. First, as in the sequential queue, potential entrants obtain idiosyncratic entry opportunities during the period. If they decide to enter the queue, they do not receive any additional information until the end of the period. At this point, the ISO provides cost assessments and fee quotes to projects based on the entire set of queue entrants. Projects then use this information and their profit draws to decide whether to connect.

Because entry into the queue is still sequential and costs are assessed jointly, projects need to form expectations over the odds of other projects' entry decisions in deciding whether to enter. However, conditional on the set of projects that choose to enter the queue,

the decision process for the ISO looks similar to the simple simultaneous model that we started with. In particular, ISOs have the choice of whether to quote fees based on average costs, marginal costs, or other metrics.

Despite the fact that the ISO can base its fee quotes on more information, the cluster process does not fully solve the hold-up problem. This is fundamentally because projects can still choose to exit the queue once costs are assessed, which creates endogeneity.

D Appendix: Model specification

D.1 Interconnection financial security

The function for IFS posting is known and set by policy. The amount of IFS posting depends on project size MW_j, interconnection costs $c_{j\tau}$ and the time period τ. We let $IFS(x_j, \tau)$ be the cumulative amount of IFS posting by project j required to continue in the queue past time τ. We define $\Delta IFS_{j\tau} = IFS_{j\tau} - IFS_{j,\tau-1}$ as the incremental IFS posting required to continue in the queue at time τ, beyond what was already posted.

The IFS posting function is approximately

$$IFS(\mathbf{x}_j, \tau) = \begin{cases} 0.15 \cdot MW_j \cdot c_{j\tau} & \text{for } \tau = 1, \text{ to } Continue \\ 0.30 \cdot MW_j \cdot c_{j\tau} & \text{for } \tau = 2, \text{ to } Continue \\ 1.00 \cdot MW_j \cdot c_{j\tau} & \text{for } \tau \geq 2, \text{ to } Contract. \end{cases} \tag{9}$$

Projects have to post a progressively larger portion of their interconnection costs to move through the queue: first 15%, then 30%, then 100% of costs when they contract for the project.

We let $IFS_{j,\tau-1}$ be the cumulative IFS posted by project j before time τ. If the project *Exits*, then it receives all of the IFS for point of interconnection upgrades and approximately half of the IFS for network upgrades back. The precise refund for the network component is given by

$$IFS_{Refund}(\mathbf{x}_j, \tau) = \begin{cases} IFS^{net}_{j\tau} - \min\{10{,}000 \cdot MW_j, 0.5 \cdot IFS^{net}_{j\tau}\} & \text{for } \tau = 1, 2 \\ IFS^{net}_{j\tau} - \min\{20{,}000 \cdot MW_j, 0.5 \cdot IFS^{net}_{j\tau}\} & \text{for } \tau \geq 3. \end{cases}$$

Projects get about half of their IFS network posting back on exit. After the period $\tau = 3$, there is no further change in the refund function. Projects earn only a small amount of interest on their postings. The main cost of lingering in the queue is the opportunity cost of capital on the posting, rather than further increases in the share of the IFS retained by the grid operator.

The function (9) is an approximation to the true IFS bond for three reasons.

1. The IFS posting has components for both the network cost $c^{net}_{j\tau}$ and the point of interconnection cost $c^{poi}_{j\tau}$. The total IFS posting amount is the sum of these two components.

2. The true function imposes minimum and maximum costs for each of the POI and network postings.

3. Because interconnection cost estimates themselves change over time, the amount of the second and subsequent IFS postings will depend on both the current cost estimates and the *prior* cost estimates, which determined the prior bond.[16]

We code the complete and accurate IFS posting function in our model including these details. Because $IFS(\mathbf{x}_j, \tau)$ is a function only of variables that are already in the project state, we do not need to keep track of the posted IFS itself as a state. We instead just calculate it from the project state in each period.

[16]For example, if cost estimates declined between postings, a project may have to put up less than the specified 0.30 share of current cost, because its initial posting was larger than required, due to a high initial estimate. This reduced posting is equivalent to posting the specified amount in each period but also getting a refund for the estimation error in the prior period.

E Appendix: Model supplementary results

Table E1: Logit Model for Receiving TPD

	TPD Received			
	(1)	(2)	(3)	(4)
MW	-0.000	0.000	0.001	0.001
	(0.000)	(0.000)	(0.001)	(0.001)
FCDS	0.719	0.832*	0.900*	0.921*
	(0.370)	(0.381)	(0.454)	(0.450)
Wind (=1)	0.161	0.024	0.118	0.129
	(0.468)	(0.482)	(0.554)	(0.554)
Battery (=1)	0.406	0.511*	0.785**	0.798**
	(0.239)	(0.254)	(0.301)	(0.299)
Co-located (=1)	2.362***	2.263***	1.888**	1.867**
	(0.609)	(0.615)	(0.642)	(0.642)
Log(AGW)		-0.217	-0.850*	-0.823*
		(0.174)	(0.370)	(0.363)
TPD Capacity (GW)		-0.029	0.072	
		(0.118)	(0.172)	
N	406	406	406	406
Pseudo R^2	0.068	0.077	0.181	0.181
Area FE	No	No	Yes	Yes
Backlog block p-value		0.081		0.003
FE block p-value			0.000	0.000

Notes: Standard errors in parentheses. * p<0.1, ** p<0.05, *** p<0.01. This table reports logistic regression results for the receipt of TPD, on observable project characteristics, backlog measures, and fixed effects. The dependent variable is a binary indicator for whether a project received TPD. (1) includes MW capacity, FCDS status, project type indicators (wind, battery), and co-location status. (2) adds backlog measures: log-transformed area MW and total TPD capacity (GW). (3) incorporates area fixed effects and uses continuous allocation years. (4) is similar to (3) but excludes total TPD capacity as a predictor. Standard errors are clustered at the area-code level where applicable.

F Appendix: Congestion Effects on Costs

Table 1 highlights a negative effect of higher congestion in the queue. To investigate this movement, we created a dataset of matched upgrade-level data and cost change multipliers to investigate the mechanism behind lower costs in highly congested areas.

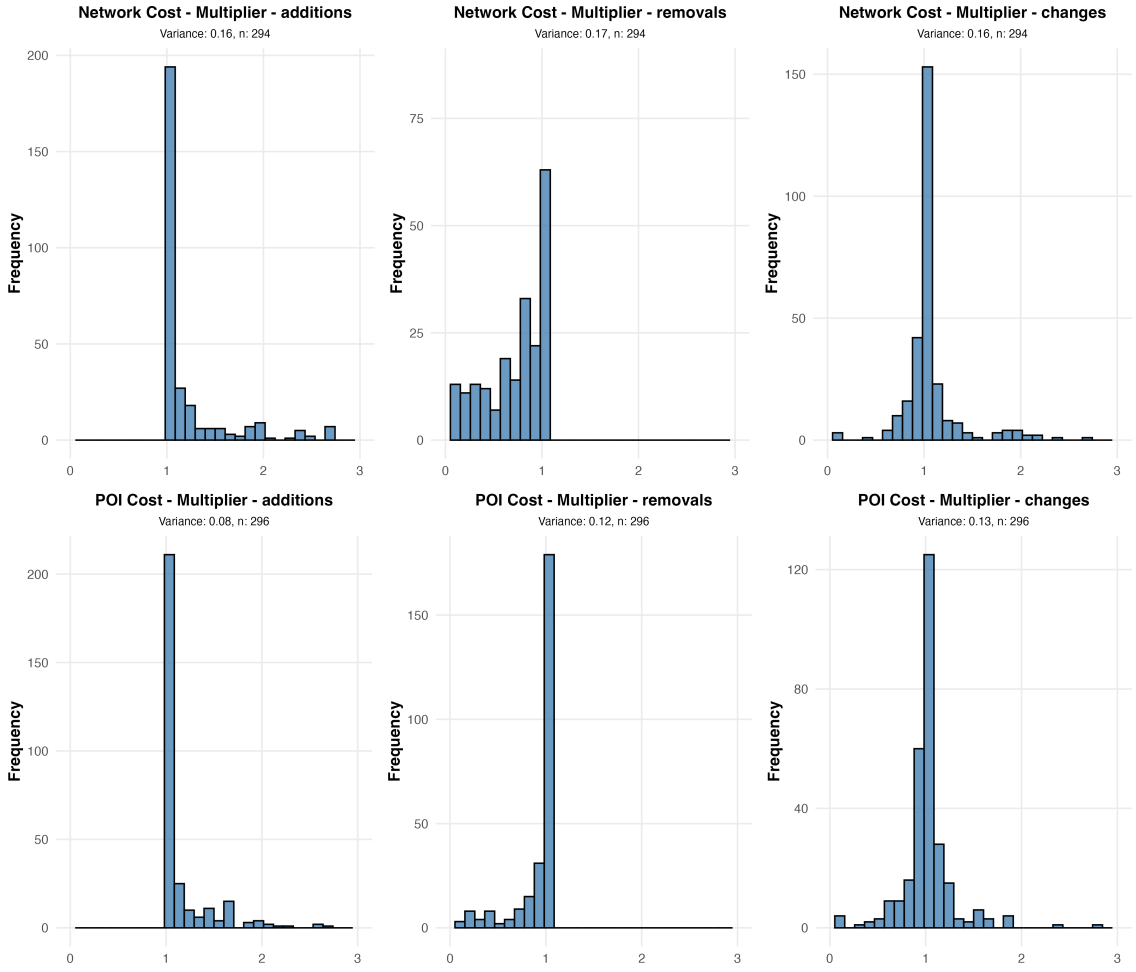
First, we match the upgrades based on the project id, type of upgrade and the name of the upgrade. We use Jaro-Winkler string distance matching within project id and upgrade category. Each upgrade item is classified as either persistent (matched across phases), removed (present in Phase 1 but absent in Phase 2), or added (absent in Phase 1 but present in Phase 2).

For each project–upgrade-category pair, the Phase 1 total cost is defined as the sum of all Phase 1 upgrade costs within type (Network or POI costs). Cost changes are decomposed into three components: (i) added cost, defined as the sum of costs associated with Phase 2-only upgrades; (ii) removed cost, defined as the sum of costs associated with Phase 1-only upgrades; and (iii) revised cost, defined as the difference between Phase 2 and Phase 1 costs for upgrades that persist across phases. Each component is normalized by the Phase 1 total cost to form a corresponding multiplier. Specifically, the addition multiplier equals added cost divided by Phase 1 total cost, while the removal multiplier equals removed cost divided by Phase 1 total cost. Intuitively, holding all else constant, each multiplier represents the proportional change in Phase 1 costs.

Figure F7 shows the distribution of these mutlipliers and Table F2 and F3 show the effect of higher AMW on these multipliers.

Table F4 reruns the transition regression, but for a subsample of projects that have been matched at the upgrade level. In accordance with F2, the effect of congestion on these projects has minimal effect on the network costs in phase 2.

Figure F7: Distribution of Upgrade Cost Multipliers



Notes: This figure shows the distribution of the cost multipliers of the upgrade matched sample

Table F2: Regression Results for Network Cost Multipliers

	Additions Multiplier	Removals Multiplier	Changes Multiplier
	(1)	(2)	(3)
Lagged log(AMW)	-0.044^{**}	0.006	0.022
	(0.021)	(0.023)	(0.026)
MW	-0.0001	-0.0001	0.00001
	(0.0001)	(0.0001)	(0.0001)
FCDS	0.003	0.072	0.069
	(0.069)	(0.066)	(0.066)
Wind Indicator	-0.166^{***}	0.196^{**}	0.048
	(0.057)	(0.087)	(0.127)
Battery Indicator	-0.079^{*}	0.144^{***}	-0.055
	(0.045)	(0.046)	(0.055)
Co-located	-0.038	-0.086	-0.062
	(0.053)	(0.060)	(0.049)
Voltage <140kV	0.038	0.132^{*}	-0.027
	(0.087)	(0.079)	(0.092)
Voltage >=500kV	0.033	-0.115^{**}	-0.036
	(0.048)	(0.051)	(0.054)
Cluster fixed effects	No	No	No
Area fixed effects	No	No	No
N	371	371	371

Table F3: Regression Results for POI Cost Multipliers (Total Data)

	Additions Multiplier	Removals Multiplier	Changes Multiplier
	(1)	(2)	(3)
Lagged log(AMW)	−0.018	−0.010	0.016
	(0.014)	(0.017)	(0.025)
MW	0.00001	−0.0001	−0.0001
	(0.00004)	(0.0001)	(0.0001)
FCDS	−0.111*	0.162**	0.025
	(0.065)	(0.081)	(0.076)
Wind Indicator	−0.126**	0.137*	−0.024
	(0.055)	(0.070)	(0.109)
Battery Indicator	−0.112***	0.126***	−0.064
	(0.032)	(0.042)	(0.043)
Co-located	−0.033	−0.023	−0.014
	(0.028)	(0.056)	(0.053)
Voltage <140kV	−0.043	−0.064	−0.023
	(0.038)	(0.091)	(0.105)
Voltage >=500kV	0.106***	−0.071	−0.012
	(0.038)	(0.048)	(0.049)
Cluster fixed effects	No	No	No
Area fixed effects	No	No	No
N	369	369	369

Robust standard errors are reported in parentheses. Columns represent POI Cost multipliers for total data.

Table F4: Initial Cost and Phase I→II Transition Regressions (Total Congestion) [Matched Sample]

	Log POI Cost($/kW)	Log Network Cost($/kW)	Log POI Cost($/kW)	Log Network Cost($/kW)
	(1)	(2)	(3)	(4)
Lagged Log POI Cost ($/kW$)			0.703***	
			(0.057)	
Lagged Log Network Cost ($/kW$)				0.680***
				(0.053)
Log AMW Backlog $_{\tau=-1}$ (GW)	0.707	1.090**	−0.060	−0.046
	(0.436)	(0.533)	(0.049)	(0.080)
MW (Own Capacity)	−0.001***	−0.0002	−0.0004	−0.0004*
	(0.0004)	(0.001)	(0.0003)	(0.0003)
FCDS (=1)	0.407	−0.012	0.528***	0.727***
	(0.288)	(0.245)	(0.185)	(0.244)
Wind (=1)	0.475	0.377	−0.149	0.311
	(0.303)	(0.400)	(0.189)	(0.322)
Battery (=1)	0.456***	0.017	−0.109	0.118
	(0.167)	(0.173)	(0.103)	(0.166)
Co-located (=1)	−0.141	−0.286	0.040	−0.003
	(0.210)	(0.235)	(0.119)	(0.194)
Voltage 200-500 kV (=1)	0.347	−0.306	−0.498**	0.165
	(0.279)	(0.276)	(0.219)	(0.226)
Voltage 500+ kV (=1)	0.400**	0.145	−0.164	−0.005
	(0.158)	(0.146)	(0.118)	(0.145)
Cluster fixed effects	Yes	Yes	No	No
Area fixed effects	Yes	Yes	No	No
Number of Clusters	12	12	12	12
Number of Areas	14	14	14	14
Number of Projects	346	346	344	346

This table represents the transition regressions for the part of the sample where the upgrades are matched. The dependent variables are log(cost + 1) for POI and network costs, expressed in dollars per kW. Models include project characteristics: project capacity in megawatts (MW), an indicator for Full Capacity Deliverability Status (FCDS), indicators for Wind and Battery (with Solar omitted), an indicator for co-located projects, and voltage-group dummies. Columns (3)-(4) model Phase II costs and include the lagged log cost from Phase I as an AR(1) term. Queue congestion regressor (Log AMW Backlog) is defined at the area–year level in gigawatts at $\tau = -1$ (previous year). Standard errors are HC1-robust. ***Significant at the 1 percent level. **Significant at the 5 percent level. *Significant at the 10 percent level.