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

Gautam Gowrisankaran, Ashley Langer, and Nicholas Ryan.†

March 26, 2026

[Please do not cite: Preliminary and incomplete]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, Lucas Davis, and Roy Shanker for helpful conversations. Thank you to Ron Borzekowski and Tom Langford at Yale for helping this project launch. This project was made possible by a collaborative Data Use Agreement between CAISO and Yale University. 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

Scarce public resources are commonly allocated without markets. For example, roads and ports are often open access, while healthcare is frequently rationed via queues. In the U.S., the electromagnetic spectrum was originally assigned via expert panels, before regulators turned to market-based allocation in the form of auctions. Governments may use non-market mechanisms to distribute these resources for many reasons including the complexity of market design when the use of a public resource causes externalities.

We are in the middle of a remarkable energy transition away from fossil fuels and towards new and renewable energy technologies that face a difficult public resource allocation problem. Over the last 15 years, the capital costs of solar power, wind power, and battery storage have declined 75%, 44%, and 94% respectively. In the U.S., recent legislation including the 2021 Inflation Reduction Act and the 2025 One Big Beautiful Bill Act have subsidized investment in these technologies. However, this energy transition requires significant shared public infrastructure in the form of upgrades to the electric grid that transports power. Designing policies to build and pay for these upgrades is difficult, because the externalities created by shared infrastructure interact with the uncertain future actions of projects seeking to potentially connect.

This paper investigates the role of cost externalities and uncertainty in the allocation of the shared grid resources required for renewable investment in California. To illustrate the complexity of this market design problem, consider the case of two potential solar generation projects in California's Mojave Desert. This is a good place for solar generation: it is sunny and land costs are relatively low. Yet, the electricity that is generated there needs to be transported to serve load in areas like Los Angeles. On the one hand, suppose a new transmission line is necessary to serve either project, but that both projects could share the new infrastructure. In this case, there could be a hold-up problem where interconnection costs would appear too high unless the regulator recognizes that the line could serve both projects. On the other hand, suppose that existing lines are adequate to serve one project but not both. Then, if the first project connected to the grid, it would raise costs for the second project by increasing congestion, potentially causing hold up to occur at that point.

Designing processes to account for both shared costs and congestion is challenging. In order to fund public infrastructure upgrades necessary for projects to connect, grid regulators—which are generally called independent system operators (ISOs)—use interconnection queues. The purpose of queues is to evaluate the infrastructure upgrades that are needed to connect new projects and to allocate the costs of these upgrades. Potential entrants join the queue to signal their interest in building new generation. ISOs then periodically evaluate necessary upgrades and quote projects required connection costs and fees. Projects incorporate updated cost information in their decision of whether to build and connect.

Yet the interconnection queue process appears to be flawed. Figure XX shows that there is currently 4.5 times as much renewable capacity waiting to connect in the U.S. as there is *total* capacity installed. Further, the amount of installed renewable capacity has increased only slowly, which indicates that projects are waiting in the queue, without connecting, much longer than they used to. Increases in the numbers of projects entering the queue are also contributing to a large backlog of projects waiting in the queue. While declining capital costs and increased government subsides might be driving some of the extra entry, the interconnection process may also be causing projects to enter for option value.

As partially illustrated by the Mojave Desert example, the combination of endogeneity and incomplete information make designing mechanisms to address over-entry and lengthy queue persistence challenging. Projects' profits are highly uncertain at the time of entry, and, in most ISOs, this uncertainty does not resolve until projects have continued through multiple cost studies spanning years. Because entry and remaining in the queue are costly, projects incur sunk costs to learn about their profitability. If projects' costs were independent, creating markets to generate optimal incentives for project entry and completion would be relatively tractable. However, project interconnection costs are interdependent, in that projects may congest or share transmission infrastructure. The fact that other projects make contemporaneous entry and completion decisions complicates this missing markets problem. This is because ISOs cannot fully predict whether projects will connect or exit the queue in response to cost quotes, but these actions affect other projects' costs.

Our investigation of the interconnection queue process in California is of widespread interest for at least three of reasons. First, California is a worldwide leader in renewable penetration, with a generation share of 57% in 2024 and a net-zero carbon generation target in 2045 (U.S. Energy Information Administration, 2025; California State Legislature, 2018). Second, the grid regulator, California Independent System Operator (CAISO), has long used an interconnection process that accounts for externalities in costs between new generation projects. Their cluster process evaluates projects' initial interconnection costs in annual clusters rather than individually, and insures projects against cost increases stemming from other projects' entry or exit decisions. This process was seen as an improvement over the earlier queueing processes. Indeed, in 2023, the Federal Energy Regulatory Commission (FERC) ordered grid operators nationwide to emulate CAISO's cluster process (Federal Energy Regulatory Commission, 2023). Yet, in that same year, California suspended its interconnection process after massive entry and the resulting project cost interdependencies made providing projects timely and informative cost quotes impossible. Finally, our investigation of CAISO's queue process provides an opportunity to study the importance of missing markets. California uses a scoring rule—rather than a market mechanism such as auctions—to allocate spare grid capacity for free via the transmission plan deliverability (TPD) process. TPD is a valuable public resource as it allows projects to bid into capacity markets, but the lack of a market allocation mechanism may mean that it is inefficiently allocated.

We base our analyses on private-use microdata—shared by CAISO under a data use agreement—which covers all projects that entered the interconnection queue over 14 annual clusters. CAISO provides projects initial quotes and updates for two types of interconnection costs: network costs—which reflect transmission infrastructure and are often interdependent—and point of interconnection costs, which are typically specific to a project. The data track projects through their annual receipt of updated cost quotes and TPD and their decisions to exit the queue, connect to the grid, or continue in the queue. We also merge CAISO data with information on local profit shifters and technology capital costs.

Our analysis proceeds in three main steps. First, we use the detailed data to provide empirical evidence on the importance of uncertainty and cost externalities in the California interconnection queue process. Second, we specify and estimate a model of a dynamic game between projects that choose whether to enter the queue and then whether to exit, continue waiting, or build the project and connect to the grid. Finally, we use this structural model to examine counterfactuals including changing the fees to enter or remain in the queue and solving the TPD missing markets problem by auctioning off spare grid capacity that is currently given away for free.

Beyond documenting the substantial queue backlog, we establish four facts about the role of uncertainty and externalities in the interconnection process.

First, initial interconnection costs are highly unpredictable ex ante—the coefficient of variation of network costs per MW is 2.60—and updates to cost quotes are also uncertain, with an R^2 of 0.45 in a regression of network costs per MW on its lag and observable characteristics. Second, externalities in network costs are important: increases in the aggregate generating capacity in the queue in a project's local area lead to substantial increases in initial network cost quotes but decreases in network cost updates. Third, interconnection costs significantly affect projects' decisions to remain in the queue and contract. Fourth, free allocation of excess grid capacity through TPD also has a large effect on projects' decisions to remain in the queue and eventually connect.

These facts imply that uncertainty and cost externalities in the interconnection queue process may be leading to the renewable project backlog, but they do not help us understand how alternative approaches may affect the backlog. We therefore specify and estimate a dynamic model of projects' entry into and progression through the queue, which allows us to better understand equilibrium queue outcomes. In our model, potential projects of three technology types (solar, wind, and battery) in each local area and cluster choose whether to enter the queue and then observe their characteristics. One year later, CAISO provides initial cost quotes and the projects choose whether to exit or remain in the queue. Two years after entry, CAISO provides updated cost quotes and allocates TPD. At this point, projects have the opportunity to exit, continue in the queue, or sign a contract and connect to the grid. In following years, CAISO may continue to update costs and projects make the same annual decisions, until a terminal period when they must commit. Projects receive *i.i.d.* profit shocks at each stage and also a persistent private unobservable shock to the profits from connecting.

Cost externalities make this a dynamic game, and in principle, projects play a Perfect Bayesian Equilibrium where strategies depend on all projects' characteristics as well as perceptions regarding rivals' persistent shock values as derived from their actions. Our estimation and counterfactuals impose that projects make decisions based on their own characteristics and a reduced aggregate state of MWs in the local queue (Ifrach and Weintraub, 2017; Gowrisankaran, Langer and Zhang, 2025). This allows for both the possibility of sharing or congesting transmission resources.

Our estimation approach starts by recovering the distribution of the CAISO engineering cost estimates as a function of the state in a pre-stage with standard panel data regressions. We then estimate the structural model using a nested fixed point quasi-maximum likelihood approach. The structural parameters include the value of TPD, the scale of the unobservables, and the profits from contracting across technology types. For each candidate parameter value, our estimator solves the dynamic decision process for every project. These decisions account for expectations of projects' behaviors using the aggregate state evolution with a correction for the project's own choices, which we implement using estimated policy functions. Identification of most of the determinants of profits follow standard discrete choice arguments. However, the scale of the persistent profit shock relative to the idiosyncratic shocks is identified by the extent to which projects persist in the queue, all else equal.

We find that TPD is valuable, worth $468,000 per MW in increased profits. However, there is little evidence of serially correlated shocks to profitability. We also find that entry costs are large and highly variable, and that these costs are lower for solar and battery than they are for wind.

Finally, we perform equilibrium counterfactuals that examine how well the current system performs, and how alterations to the system may influence projects' connection decisions and affect time in the queue. We consider charging fees to projects that remain in the queue without contracting, auctioning off TPD, and removing CAISO's insurance against interconnection cost increases. To empirically measure the scale of the externalities imposed by a project, we also exogenously increase capacity in the local queue by the size of the average project and examine how project decisions change.

Literature: Our paper principally contributes to three literatures. First, we extend a broad literature on market design with externalities, which has considered how policies affect the equilibrium allocation of scarce resources in industries such as shipping (Brancaccio, Kalouptsidi and Papageorgiou, 2024), airline slots (Marra, 2024; Bauer, 2025), and public housing (Waldinger, 2021). This literature has a long history of working directly with policymakers to improve the allocation of public resources such as the electromagnetic spectrum (**?**) and donor organs (**?**). Our setting is unusual in that physical externalities interact with endogenous uncertainty from agents' decisions, making it difficult to achieve first-best outcomes with simple market mechanisms.

Second, we add to the growing literature on the estimation of structural models of electricity markets (e.g. **??**). We model potential renewable projects' dynamic interconnection decisions, extending the literature on the dynamics of electricity markets, which includes Elliott (2025); Butters, Dorsey and Gowrisankaran (2025), and **?**. Closest to our research is Johnston, Liu and Yang (2023), which studies the interconnection queue in the Pennsylvania-New Jersey-Maryland (PJM) market with a dynamic model. Our approach differs from this work by directly incorporating the interdependence between projects' costs and specifying a dynamic game between projects.

Third, our paper contributes to the literature on the economics of renewable energy. Renewable energy is distinct from earlier fossil fuel sources in that generation potential varies over space and is often intermittent, with output increasing when the sun shines or the wind blows (Borenstein, 2012). Many papers have investigated how electricity markets accommodate renewable growth, including wholesale energy market behavior (Ito and Reguant, 2016; Bushnell and Novan, 2021) and the impact of transmission congestion on the value of renewable energy (Fell, Kaffine and Novan, 2021; Gonzales, Ito and Reguant, 2023). We study the interconnection queue, which can serve as a bottleneck between private investment in new projects and public investment in the power grid.

The paper continues as follows. Section 2 provides a description of and empirical evidence on CAISO's current interconnection process. Section 3 specifies our dynamic model of CAISO's interconnection queue. Section 4 provides estimation details. Section 5, presents the model's parameter estimates and counterfactual simulations. Section 6 concludes.

2 CAISO's Queue Interconnection Process

2.1 Institutional Background

The power grid is a network that ensures that electricity flows reliably from generators to users. Balancing power flows is a complex engineering endeavor that requires sophisticated algorithms, expensive equipment, and for new generation sources to be integrated with the existing grid. If new sources could freely connect, congestion could limit their ability to deliver power to customers and even cause grid instability or blackouts.[1] In California, projects therefore need CAISO approval and potentially infrastructure upgrades to connect.

The goal of CAISO's interconnection process is to manage these externalities while keeping interconnection costs low. CAISO conducts cost studies that evaluate necessary grid upgrades using advanced engineering models that take as inputs the current grid infrastructure and locational supply and demand. CAISO divides costs into point of interconnection (POI) costs—which specify the costs to connect the project to the nearest grid location—and network costs, which specify necessary upgrades to the rest of the grid.

Externalities imply that assessed costs will vary with the set of projects being considered. CAISO conducts cluster-based cost studies that account for interdependencies across projects. In each annual cluster, projects enter by a deadline by posting a deposit and by filing an interconnection request that documents project characteristics (e.g. capacity, technology) and desired entry location.[2] For each cluster, CAISO calculates interconnection costs under the assumption that all projects in the queue will be built.

[1] Congestion is when exports or imports from some node exceeds the transmission capacity connecting that node. In this case, the flow of power is rationed (either through separate nodal prices or direct control) to maintain grid stability.

[2] In 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).

Projects receive initial cost assessments approximately one year after the cluster entry deadline. They then decide whether to exit or continue in the queue, which requires paying a deposit that is increasing in their assessed costs. A year later, projects receive a second cost assessment that incorporates earlier queue exit and provides more detailed engineering information. At this point, projects can build their project, pay the remainder of the required fees, and connect to the grid. Or, as in the previous period, they can exit or continue in the queue by paying an additional deposit.[3] After the second year in the queue, projects face the possibility of further cost revisions but do not need to pay additional deposits to remain in the queue. CAISO insures projects against cost increases by specifying that interconnection fees are based on the lowest POI and network cost assessments the project receives. Figure 2 lays out the timing of the CAISO cluster process.[4]

California's three major utilities—Pacific Gas & Electric, San Diego Gas & Electric, and Southern California Edison—and other transmission operators make large-scale transmission plans to accommodate long-run anticipated supply and demand growth. Electricity charges fund the investments required by these plans. As part of the interconnection process, projects compete to be allocated transmission plan deliverability (TPD), which guarantees them access to part of the grid capacity created by the transmission plan. Immediately after the second interconnection cost study, CAISO allocates TPD with a scoring rule based on project characteristics, such as whether the project has proven control of its site.

2.2 Data

Our main data are from CAISO, principally from their Resource Interconnection Management System (RIMS) database. They cover characteristics, costs, and decisions for all

[3] Projects that exit the queue after the first year are refunded a portion of their deposit.

[4] Appendix A provides additional detail on the timing and terminology of CAISO's interconnection process.

queue project entrants over 14 years, from 2009–22. We summarize our data construction here with details in Appendix B.

Our data contain several different sub-components. First, we observe project characteristics submitted at the time of application, notably the type of power plant to be built, its entry location, and its capacity. Second, we observe updates on the status of each project in the queue, specifically whether the project has exited or decided to connect to the grid, 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. Once a project decides to connect, it must still sign a Generator Interconnection Agreement (GIA) and build its facility, but we refer to these actions in combination as the "connect" or "build" decision. Third, we observe information from CAISO's cost studies, which report the network and POI costs required for projects to connect to the grid as well as the extent to which these costs are shared with other projects. Finally, we obtain information from CAISO on which projects received TPD.

Our analysis partitions California into mutually exclusive areas. The idea is to approximate electrical neighborhoods, within which projects are more likely to exert externalities on each other. Using our data on cost sharing, we construct an adjacency matrix which contains indicators for the sets of projects that share substantial costs for at least one network upgrade for the 2022 cluster. We then apply a Louvain clustering algorithm to this matrix to obtain 13 mutually exclusive areas. Appendix B.3 details this algorithm, and Appendix Figure B3 shows that our area definitions are strong predictors of the probability that two projects share any network upgrade costs.

We augment the CAISO data with three additional data sources. We obtain capital cost data for solar, wind, and battery storage from the National Renewable Energy Laboratory's (NREL) Annual Technology Baseline and related Lawrence Berkeley National Laboratory inputs (National Renewable Energy Laboratory, 2024; Wiser et al., 2025; Seel et al., 2025), as well as the International Renewable Energy Agency (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).

2.3 Reduced-Form Evidence on CAISO Queue Performance

Before turning to our structural model, we use our data to substantiate five points regarding CAISO's interconnection queue. The points illustrate the queue backlog, the extent of interconnection cost uncertainty and externalities, and the importance of interconnection costs and TPD to projects' connection decisions.

California Interconnection Backlog.—Focusing first on the backlog, Figure 1 panel (b) shows the total capacity in CAISO's queue from 2009-2022. We see a pattern of increasing renewable capacity persisting in California's interconnection queue over time. This increasing backlog results from the combination of increasing queue entry and projects remaining in the queue longer over time. In 2021 (cluster 14), 38.4 GW of new renewable capacity *entered* the queue, which was nearly as large the total installed base of renewable capacity in California (40.6 GW).

Further, Figure 4 shows the progression of projects through the interconnection queue by months since initial entry, separately by cluster. The figure shows the share of projects either remaining in the queue (panel A) or connecting to the grid (panel B) over time. We see from panel A that projects remain in the queue for a long time, and that this pattern is increasing over time. Recent clusters have around 35% of projects still in the queue after four years. This queue persistence exacerbates the backlog as more projects enter the queue. Panel B shows that this increased persistence in the queue is *not* leading to higher contracting rates, with only around 20% of projects contracting within four years of queue entry. In combination, this means that the 379 GW of capacity in the California queue in 2021 was 9.3 times higher than the installed base of renewable energy capacity and 4.6 times higher than *total* installed capacity in California.

Interconnection Cost Uncertainty.—Turning next to uncertainty, projects that enter and continue in the queue receive information about the costs that they would be required to pay to connect to the grid. Entry and queue persistence provide more option value when these costs are more uncertain. Figure 5 shows the evolution of POI costs (left column) and network costs (right column) over years in the queue. The top panel shows that both POI and network costs from projects' first cost study are highly uncertain: the standard deviation of POI costs is 2.82 log points and the standard deviation of network costs is 3.08 log points, suggesting that it is quite common for some projects to receive cost assessments that are three orders of magnitude larger than others'. Network costs are generally larger than POI costs, with the 90th percentile of network costs at $322 thousand per MW, which is 25% of the 2020 capital cost of a MW of solar, meaning that this uncertainty is important for profits.

The middle panel of Figure 5 plots costs from the first cost study against costs from the second cost study. It shows that this uncertainty is not resolved after the initial cost study: the initial log costs and a constant only explain 58.1% of the variation in log POI costs and 45.4% of the variation in log network costs. The lower panel of Figure 5 shows that this uncertainty is largely resolved after the second study for POI costs but not for network costs. For this reason, our model will assume that projects only receive two POI cost assessments, but may receive additional network cost assessments.

Cost Externalities.— These cost uncertainties are particularly important because they interact with the presence of externalities across projects. Table 1 regresses interconnection costs on project characteristics and the aggregate capacity—in the queue or connected during the cluster process—in the project's local area. The first two columns investigate initial POI and network costs respectively, and the third and fourth columns investigate cost reassessments between the first and second cost studies. All four columns show evidence of cost externalities, although externalities are larger and more statistically significant for network costs than for POI costs. For initial cost quotes, more capacity in the queue increases projects' POI and network costs, with the effect for network costs being statistically indistinguishable from 1. This suggests that there is substantial congestion between projects that increases network costs as more projects enter the queue. After the initial cost assessment, lagged cost assessments are an important determinant of cost revisions, but the capacity of projects remaining in the queue is also statistically and economically important. Conditional on lagged assessments, local queue capacity actually *decreases* interconnection costs, suggesting that the ability to share costs with other projects is valuable.[5] We therefore find evidence for both congestion and cost sharing externalities at different stages of the process.

Importance of Interconnection Costs.—We next investigate the overall importance of these costs in influencing projects' queue decisions. Table 2 presents linear probability models of the probability of deciding to remain in the queue after the first cost study, conditional on cost assessments and various controls. Since continuing to the second cost study requires paying a partially non-refundable deposit, this decision reflects projects beliefs about their potential profitability. We see in columns one through three that, with increasing controls, higher initial network cost assessments always lead to lower probabilities of continuing in the queue. Column four presents similar results, but allowing for non-linear effects by cost quartile. The effects are large. In column 4, moving from the first to the

[5]Other coefficients in the column 2 specification are also consistent with power system operation. The lowest voltage projects have the lowest network and POI costs. Projects like batteries, that are co-located with other projects have lower network costs because they require fewer transmission upgrades.

second quartile of network costs reduces the probability of advancement by 11 percentage points. Moving from the first to the fourth quartile of network costs reduces the probability of advancement by fully 41 percentage points. Since Table 1 showed that these costs are impacted by nearby projects, this suggests that the externalities imposed by nearby projects can be critical for project decision-making.

Transmission Plan Deliverability.—We have shown that uncertainty and externalities are important for projects decisions, but CAISO also allocates spare grid capacity via TPD. Since this program allocates TPD for free with only imperfect information about projects' profitability, this program creates a "missing markets" problem if projects find TPD valuable. Inefficient allocation of a valuable resource could lead to over-entry into the queue and over-persistence in the queue, which may exacerbate the queue backlog.

Figure 6 plots the fraction of projects that have connected to the grid over time, conditional on whether they have received TPD. Projects can only connect after two years in the queue, at which point they also know whether they have TPD. The figure shows that projects receiving TPD immediately start connecting at much higher rates than those that do not, and this pattern continues for years. By ten years after queue entry, projects receiving TPD are over 30 percentage points more likely to have connected than those without TPD. Table 3 shows that these patterns persist even conditioning on project characteristics. Column 2 of the table shows that TPD receipt increases the probability of connecting within four years of queue entry from 20.7% to 51.9% with a full set of controls. This effect is consistent across network cost assessments, but is particularly strong for projects in the second network cost quartile, suggesting that TPD is most valuable for projects that might be on the margin for connecting.

3 Model

Our reduced form evidence shows that uncertainty, externalities, and delay are important in understanding CAISO's interconnection queue performance. However, this evidence is not sufficient to understand the extent to which these forces affect queue outcomes or how alternative market designs might improve the queue backlog.

To illustrate the issues with market design, we begin by presenting a simple conceptual model that explains the key forces at play. We then turn to the empirical model we estimate, that includes the complexities of CAISO's actual cluster interconnection process, explaining the different stages of this process and our concept of equilibrium.

3.1 Simple Conceptual Model

We consider the problem of energy projects that want to connect to the grid in a geographic area. Each project faces three stages to this process. First, the project learns its type (solar, wind, or battery) and receives a one-time opportunity to enter the interconnection process. It receives a draw from an entry cost distribution, which it must pay if it enters. Upon entry, the project learns its size in MWs and other information that may affect its payoffs. Second, the ISO assesses the project's interconnection costs. It reports these variables to the project, and they ultimately determine the interconnection fees that the project will pay if it chooses to connect. Finally, the project obtains an idiosyncratic shock to the profit of connecting and decides whether or not to build its generation facility and connect to the grid.

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 to fees, 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 externalities.

Original Approach: Sequential Interconnection Queues.—The original approach to the interconnection process was queueing, which was used by many different ISOs. 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. In a sequential process, once a project enters the queue, the ISO assesses its 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.

To understand the issues with the queueing approach, consider the case of two potential projects that receive sequential opportunities to enter the interconnection queue and connect to the grid, with project 1's opportunity before project 2's.[6] First, suppose that a transmission line needs to be built if either or both projects enter. In this case, the ISO would quote project 1 a fee equal to the entire cost of the transmission line. This would then lead project 1 to have too little of an incentive to enter the queue and build its generation capacity. Moreover, if project 1 declined to enter, project 2 would then also have an underincentive to enter the queue and build.

Now suppose that the transmission line needs to be built only if both projects enter. In this case, a market failure still exists, but its impact is more subtle. Project 1 would be overincentivized to enter and build because it would not bear the congestion cost to project 2. But, project 2 would have to pay the full cost of the transmission line and thereby have too little an incentive to build. Thus, depending on the values of the entry and profit unobservables, we may end up with one project when the social optimum would have both projects being built.

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

Updated Approach: Cluster Analyses.—Overall, the sequential nature of the queueing process is leading to substantial hold-up in building transmission infrastructure and thereby slowing the green energy transition. Understanding this concern, CAISO implemented its cluster process in 2008 and was at the forefront of the move from sequential interconnection queues to cluster analyses. In a cluster process, projects enter by a deadline, the ISO evaluates all projects in a cluster simultaneously, and projects then make their connection decisions simultaneously.

Cluster analyses allow the ISO to aggregate information across projects and thereby better account for cost interdependencies. However, a cluster analysis does not fully solve the hold-up problem because projects need to form expectations over other projects' entry and connection decisions.

Considering the same example, and following CAISO's actual practices, projects are quoted costs at the second stage equal to their capacity share of the total costs. If a project exits at the third stage, CAISO insures the remaining project against any cost increase. When infrastructure is necessary for either project, each project does not care if the other project drops out since this it is protected against fee increases. Hence, CAISO's policy would result in projects getting built, but at the cost of ratepayers having to subsidize infrastructure investments in the case where one project exits the queue. When infrastructure is only necessary if both projects connect, then allocating costs to both projects may result in inefficient expected exit at the third stage. If entry at the first stage is costly, projects also will not have the correct entry incentives.

Our reduced-form evidence showed that projects' costs are highly uncertain, that projects impose externalities on each other, and that these externalities are important for projects' queue decisions. This subsection and Figure 7 shows how these forces may lead to a cycle that generates inefficient entry and backlog, even with cluster analyses. Uncertainty means that projects need to pay to enter and learn costs. The potential to obtain low interconnection costs or TPD creates option value, which increases entry and delay. But, the externalities between projects are exacerbated by the extra projects and backlog. This, in turn, increases uncertainty. This intuition is intrinsic to our more detailed structural model of CAISO's interconnection process, which allows us to understand the impact of alternative market designs on equilibrium queue outcomes.

3.2 Empirical Model

Our empirical model builds on the conceptual model in Section 3.1 by incorporating the salient features of CAISO's cluster interconnection process. As above, potential entrants first decide whether to enter the interconnection queue and then receive a cost quote. However, as in the real world, we model the queue stage as more lengthy. Projects in the queue learn about their interconnection cost quotes and TPD over time and periodically make decisions about whether to exit the queue, continue without connecting, or connect to the grid.

Also different from our conceptual model is the presence of multiple clusters, or entry cohorts. We index each cluster by $y \in \{1,\dots,14\}$. We index annual periods by $\tau = 0,1,\dots,T_{max}$, where $T_{max} = 15$ is the maximum number of periods a project can remain in the queue. 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. To simplify notation, we focus this section on a single local area. Our empirical approach considers all all areas, but treats them as independent.

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 build and *connect* to the grid (Figure 2). Because interconnection costs depend on which other projects are in the queue, these decisions generate the interdependencies described in Section 3.1. The role of CAISO is to study the state of the queue and, in specific periods, to provide projects with new estimates of cost or an allocation of TPD.

3.3 Queue Stage

Focusing on the 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 per-MW 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, CAISO 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. Second, with updated information, projects act. We discuss these two phases starting with the ISO phase.

ISO Phase.—CAISO runs studies to assess costs and assigns TPD to projects. We model CAISO as conducting these tasks with fixed rules rather than modeling CAISO's underlying objective function. CAISO performs different tasks across periods. At the beginning of period $\tau = 1$, it assigns projects initial estimates of interconnection costs, $c^{POI}_{j\tau}$ and $c^{Net}_{j\tau}$. At the beginning of period $\tau = 2$, CAISO updates these costs with refined estimates. At this point, CAISO also assigns TPD to projects. In later periods, $\tau \geq 3$, CAISO may update the network costs of projects again—depending on the cluster and the period—in what is called a cost reassessment. We model network cost reassessments as occurring probabilistically, with a probability that varies with cluster and period. POI costs and TPD are fixed after $\tau = 2$.

For each cluster, CAISO uses engineering models to perform joint cost studies. We represent this study process as a stochastic mapping from the project's characteristics, including its own prior cost estimate and capacity in the local area, to an assessment 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. CAISO also assigns transmission plan deliverability to projects based on their characteristics. We similarly model 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 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 make decisions. The potential actions for project j for $\tau \geq 2$ are $a_{j\tau} \in \{\text{exit}, \text{continue}, \text{connect}\}$. At $\tau = 1$ projects are not allowed to connect, since the first cost estimates are seen as too uncertain.

At the time projects act, we assume they know the characteristics and cost quotes for themselves and all other projects in the queue, including those from previous clusters. Project j also has private information on its own persistent profit shock ω_j and idiosyncratic cost shocks, $\varepsilon_{j\tau}^{Exit}$, $\varepsilon_{j\tau}^{Continue}$, and $\varepsilon_{j\tau}^{Connect}$. We assume that the idiosyncratic shocks are distributed type 1 extreme value, and that the only potential actions at the terminal period $\tau = T_{max}$ are *exit* or *connect*.

We let the idiosyncratic shocks, $\varepsilon_{j\tau}$ and ω_j, be shocks 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 in dollars per MW, which we denote σ_ε.[7]

Payoffs.—If a project *continues* in the queue in periods one or two it needs to post a deposit, which CAISO calls an Interconnection Financial Security (IFS) payment. The decision to proceed in the queue has real stakes because projects only get part of their deposit back if they later exit. In general terms, we can write the per MW flow payoff from continuing as

$$-Deposit(MW_j, c_{j\tau}^{POI}, c_{j\tau}^{POI,min}, c_{j\tau}^{Net}, c_{j\tau}^{Net,min}, \tau) + \sigma_\varepsilon \varepsilon_{j\tau}^{Continue}, \tag{1}$$

where *Deposit* is the additional per MW deposit required to be posted to continue as of period τ. CAISO policy sets IFS payments. 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.

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

$$Refund(MW_j, c_{j\tau}^{POI}, c_{j\tau}^{POI,min}, c_{j\tau}^{Net}, c_{j\tau}^{Net,min}, \tau) + \sigma_\varepsilon \varepsilon_{j\tau}^{Exit}. \tag{2}$$

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

[7] 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.

high network costs. Appendix D.1 provides additional details on CAISO's IFS policies.

Finally, if a project *connects*, 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 per MW, now scaled in $1/\sigma_\varepsilon$ dollars rather than dollars, as

$$\frac{1}{\sigma_\varepsilon}\left(\overline{\pi}(x_j, TPD_j, \tau) - NetFees(MW_j, c_{j\tau}^{POI,min}, c_{j\tau}^{Net}, c_{j\tau}^{Net,min}, \tau)\right) + \frac{\sigma_\omega}{\sigma_\varepsilon}\omega_j + \varepsilon_{j\tau}^{Connect}, \quad (3)$$

where $\overline{\pi}$ is the mean per MW profit in dollars for a project based on its observables, $NetFees$ are the remaining payments the project owes CAISO to be able to connect, in dollars, and the next term is the contribution of the persistent project shock to profit.

3.4 Entry Stage

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

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 simultaneously receives a one-time opportunity to enter with its energy project. We assume that there is an idiosyncratic shock to the cost of entry and its distribution varies by type, $\eta_k \sim N(\mu_k, \sigma_{\eta k})$.

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.

3.5 Equilibrium

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

We assume that project j's fixed characteristics, x_j, are publicly observable. Some components of x_j enter profits directly, notably technology and size. Others, such as voltage, affect payoffs indirectly because they enter interconnection cost assessments. At any period τ, project j also has a set of publicly observable time-varying characteristics that can affect its future payoffs in at least some periods. Beside τ itself, we refer to these as: $s_{j\tau} \equiv (c_{j\tau}^{Net}, c_{j\tau}^{Net,min}, c_{j\tau}^{POI}, c_{j\tau}^{POI,min}, TPD_{j\tau})$. Current assessed costs affect future cost reassessments. Lagged minimum cost estimates enter into the determination of refunds in the case of exit and eventual interconnection fees in the case of connection. The project's also knows its privately observable persistent and transitory payoff shocks, ω_j and $(\varepsilon_{jt}^{Continue}, \varepsilon_{jt}^{Exit}, \varepsilon_{jt}^{Contract})$, respectively.

In principle, the information set for project j at time τ includes $s_{k\tau}$ and x_k for all projects k as well as the project's privately observed payoff shocks. Public information allows project j to form equilibrium beliefs regarding other projects' ω values. However, there can be hundreds of projects in the local CAISO queue at any point in time, reflecting the large number of potential renewable energy projects. Modeling the impact of every potential project would imply an intractably huge number of states and is probably more complex than how projects' actual decisions are being made.

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) as extended by 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 in the area (or area-gigawatts AGW_y)—either in the interconnection queue or connected to the grid—as the aggregate state. AGW_y includes capacity from the current cluster and capacity from prior clusters that is either still in the queue or has connected.

We assume that projects make decisions at period τ given their x_j, $s_{j\tau}$, their private information, and this aggregate state. Following Ifrach and Weintraub (2017), an MME will provide a reasonable approximation of beliefs if most of the project's dynamic game information is encapsulated in the aggregate and individual states. In our case, the idea of using AGW_y is that it captures cost externalities as well as expectations regarding other projects' ω values. We model this aggregate capacity measure as affecting CAISO's cost study models and thus affecting project costs. Higher AGW_y can result in both congestion that necessitates additional infrastructure and the potential to share infrastructure costs across projects.

The ABOE extension of MME allows all projects to act as large players who understand that their decisions influence the aggregate state. In an ABOE, if a project chooses to continue in the queue, it understands that AGW_y will be higher than it would be in expectation, where there is some chance that that it would have exited. Similarly, if a potential project enters the queue, it understands that AGW_y 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 $(\tau,\ s_{j\tau},\ AGW_y)$. 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 expected

MWs or MWs to correct for the fact that the AR(1) regression included some probability of the project exiting.

4 Estimation

We estimate the model in three parts. In the first part we estimate the CAISO policy functions and endogenous regressions with panel regressions. These include the structural engineering cost assessments and TPD policy function as well as 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.

4.1 Parameters estimated outside the dynamic game

We first estimate regressions of cost assessments 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.

4.2 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

$$L_j(\omega_j,\theta) = \prod_{\tau=1}^{\overline{\tau}_j} P(a_{j\tau} = a_{j\tau}^{obs} | \mathbf{s}_{j\tau} \backslash \omega_j, \omega_j, \theta), \tag{4}$$

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

$$\log L(\theta,\Omega) = \sum_{j=1}^{J} \log\left(\sum_{l=1}^{L}(1/L)\prod_{\tau=0}^{\overline{\tau}_j} P(a_{j\tau} = a_{j\tau}^{obs} | \mathbf{s}_{j\tau} \backslash \omega_j, \omega_j = \sigma_\omega \omega^l, \theta)\right). \tag{5}$$

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

5 Model Results [Preliminary and subject to change]

5.1 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 3.4, 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. [8] 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 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 **??**, which showed that projects that receive TPD 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.

[8] 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 $1/3.054 = 0.327$

5.2 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.

5.3 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 **??**).

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 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.

6 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).

References

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.

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

Federal Energy Regulatory Commission. 2023. "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.

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.

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

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

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)."

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. "California state energy profile."

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."

7 Figures

Figure 1: Active Interconnection Queue Capacity by Year

A. All projects in the United States

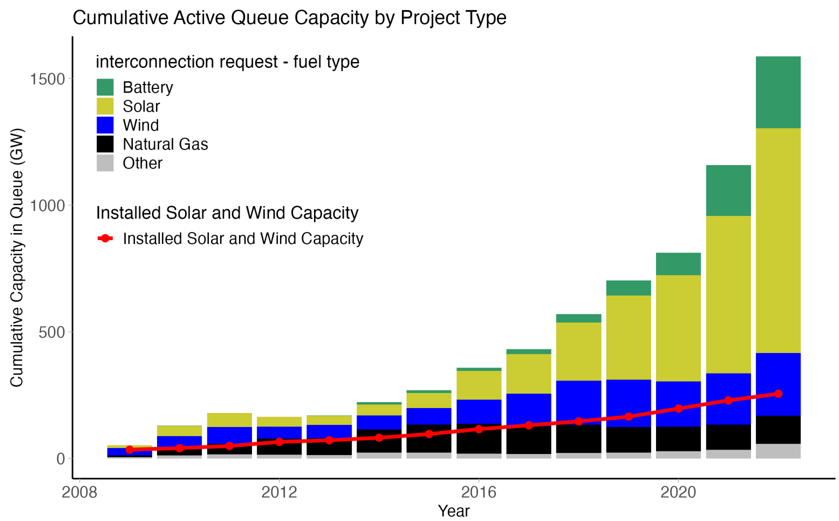

B. Projects in CAISO

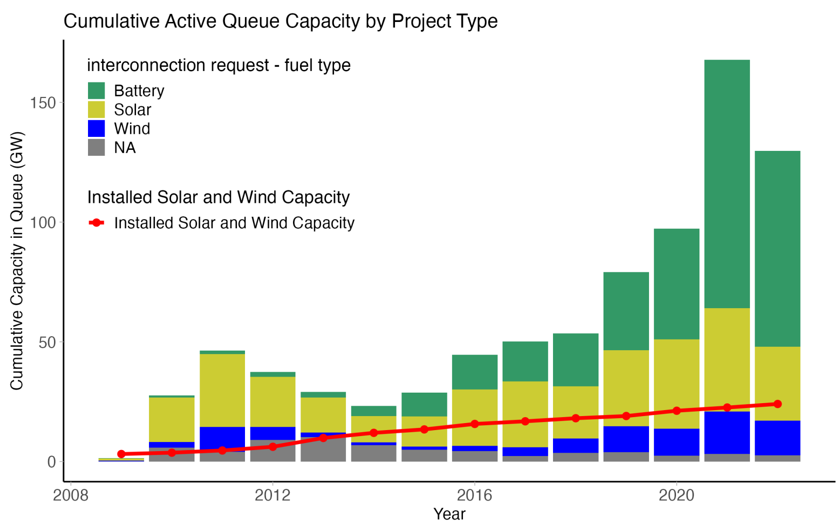

Notes: The top figure shows the total cumulative capacity (GW) of projects in interconnection queues across the U.S. based on data from the LBNL as published in 2023. The red line plots the cumulative capacity of wind and solar projects ($\geq$ 1 MW) across the U.S. for each year. The data come from 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

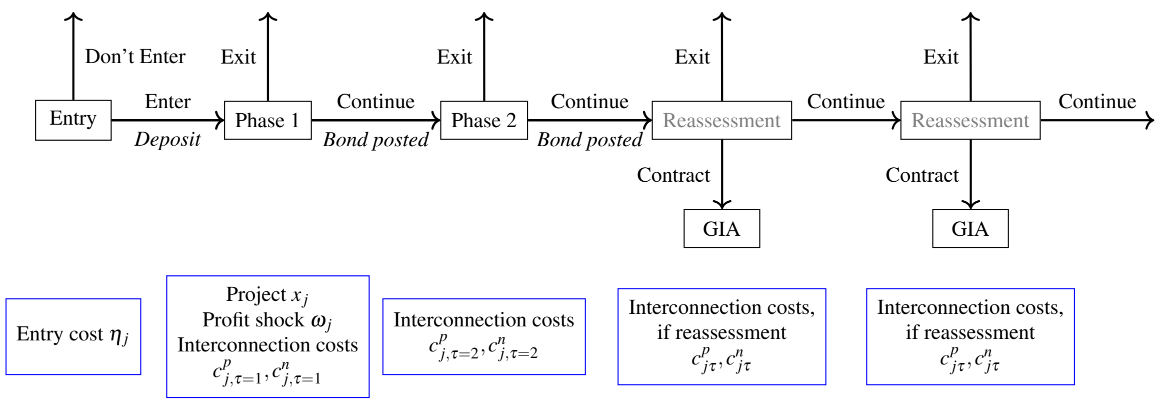

Notes: The figure shows the extensive form of the interconnection game for a single project. Potential entrants decide whether to enter the queue (i.e., file an interconnection request) or not. In Phase 1 they receive an initial cost estimates and can continue in the queue or exit. In Phase 2 they receive 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)

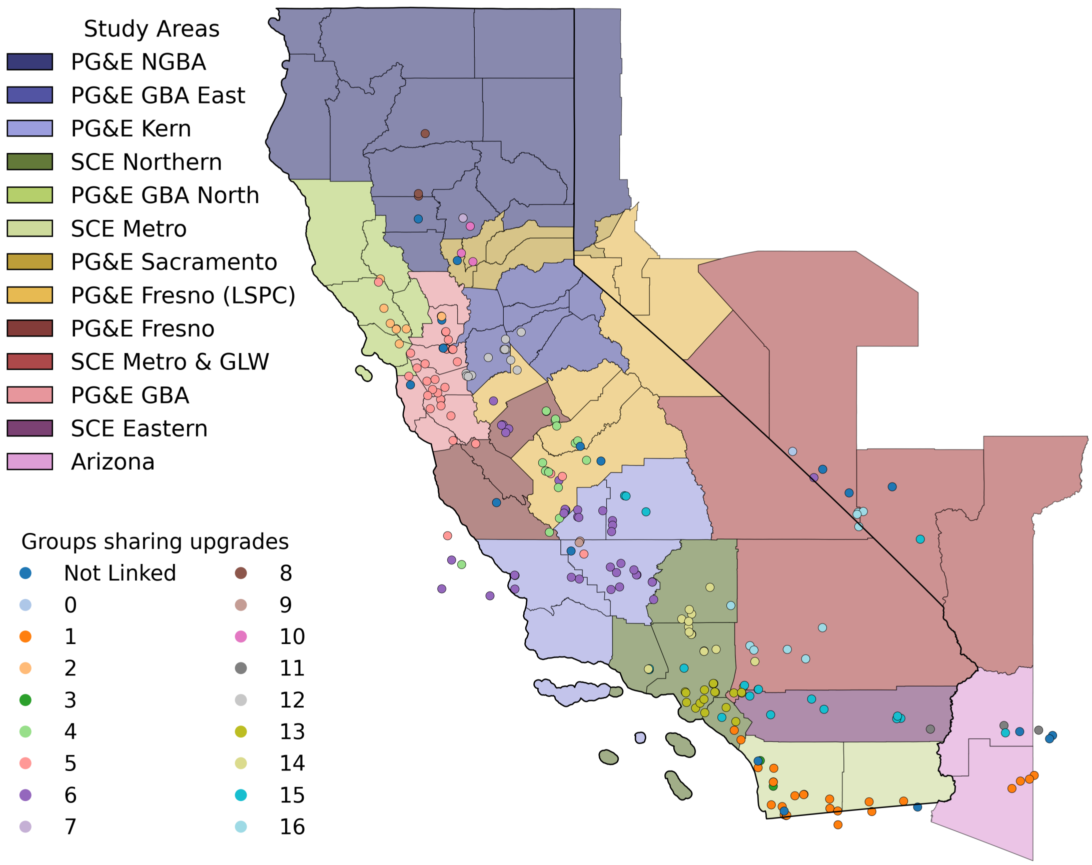

Notes: The figure shows each cluster of projects as a result of using Louvain clustering using points. Points which have no common upgrades with others are considered in the Not Linked cluster. The black outline highlights California's state boundary. Each shaded portion constitutes an area in the model.

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

A. Projects remaining in the queue

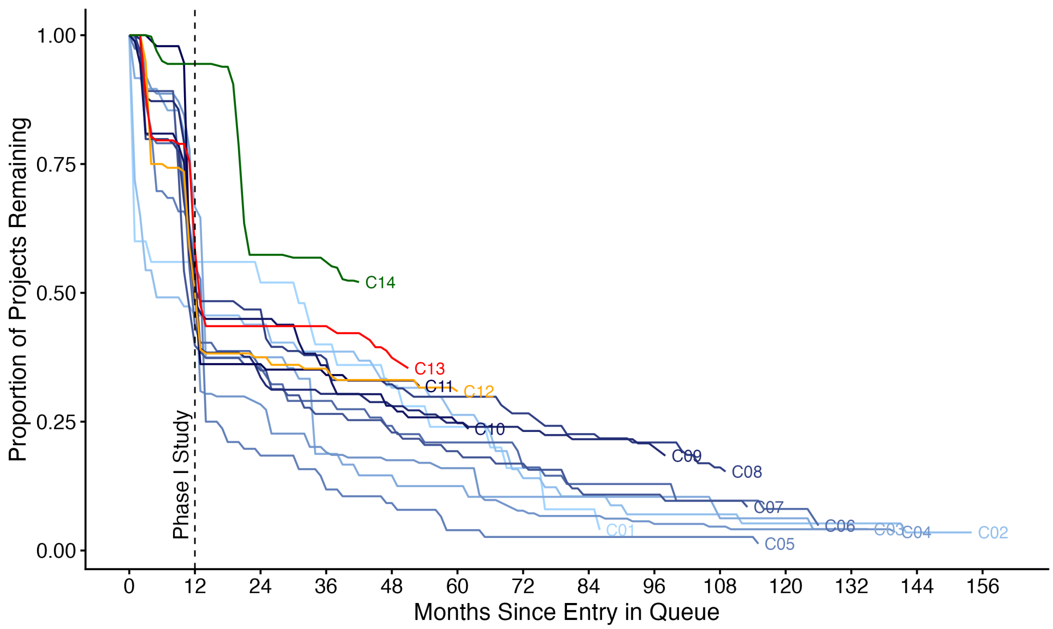

B. Projects signing a Generator Interconnection Agreement (GIA)

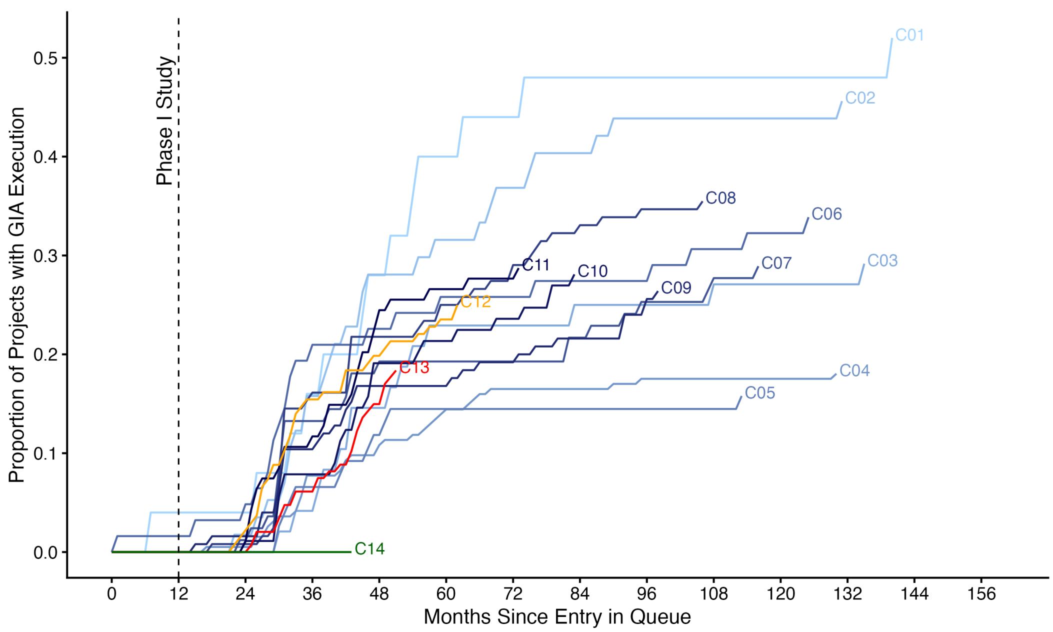

The figure shows the share of projects remaining in the queue (panel A) or signing a Generator Interconnection Agreement (GIA) (panel B) over time. In each panel, each line indicates a separate cohort (which CAISO calls "clusters") of queue entrants, from cluster 1 (entry year 2009) to cluster 14 (entry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

A. POI costs, initial distribution (Phase I)

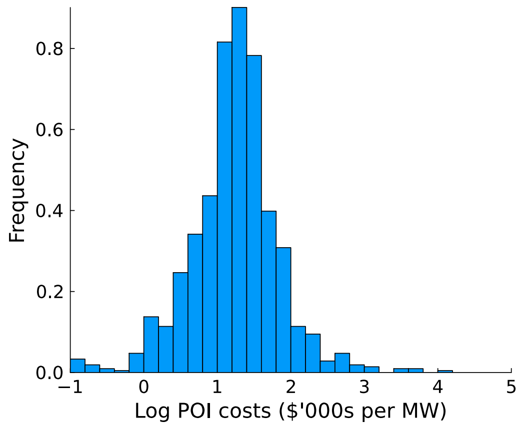

B. Network costs, initial distribution (Phase I)

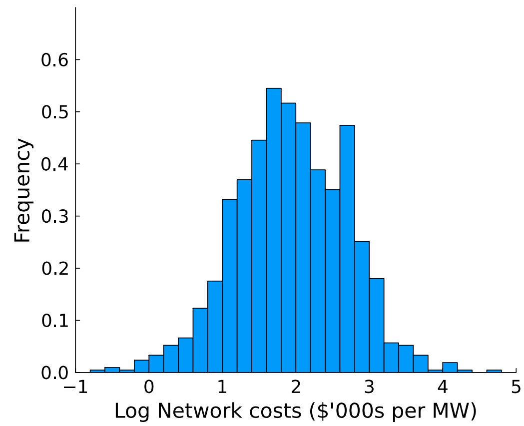

C. POI costs, transition from Phase I to Phase II

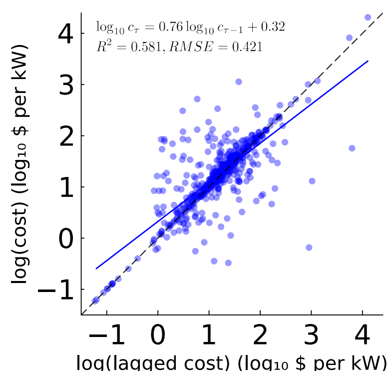

D. Network costs, transition from Phase I to Phase II

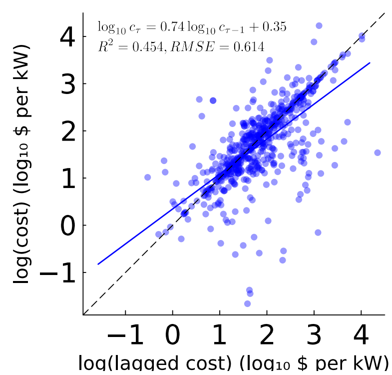

E. POI costs, transition from Phase II to Reassessment

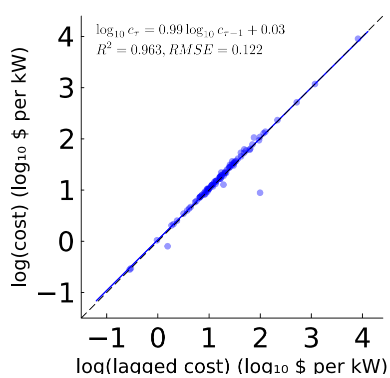

F. Network costs, transition from Phase II to Reassessment

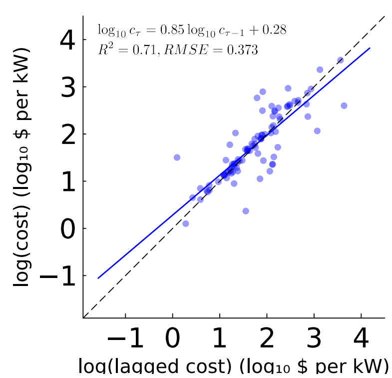

The figure shows the initial cost distribution and its evolution over time for the two different components of interconnection costs. The panels in the left column show Point of Interconnection (POI) costs, the costs paid by one project alone to plug-in to the grid. The panels in the right column show 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

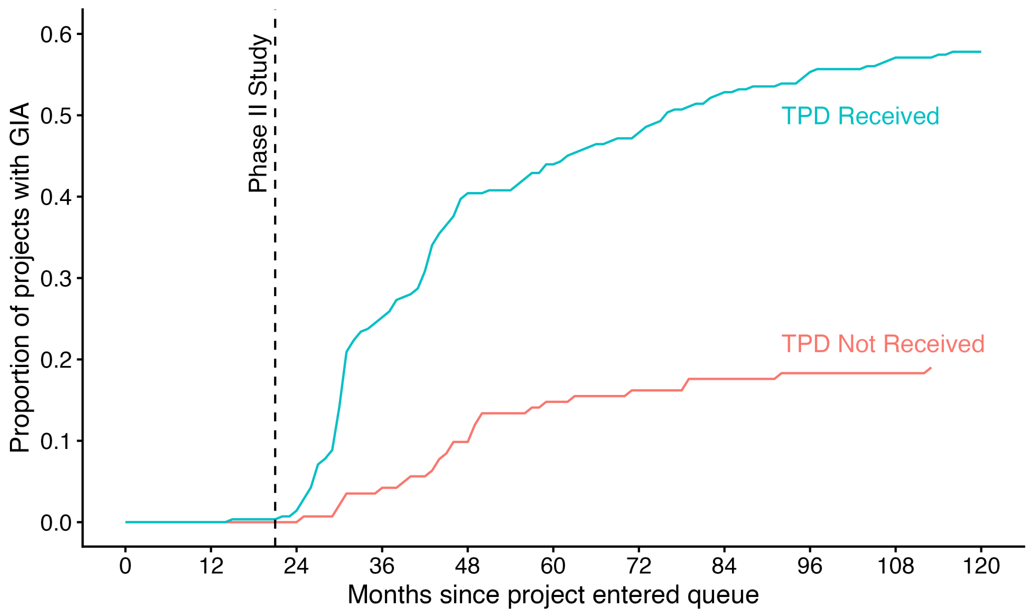

The figure plots the cumulative proportion of projects that have executed a Generator Interconnection Agreement (GIA) against the number of months since the project entered the interconnection queue. A GIA is a contract to connect to the power grid. Separate curves are shown for projects that received 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: Sources of Equilibrium Backlog

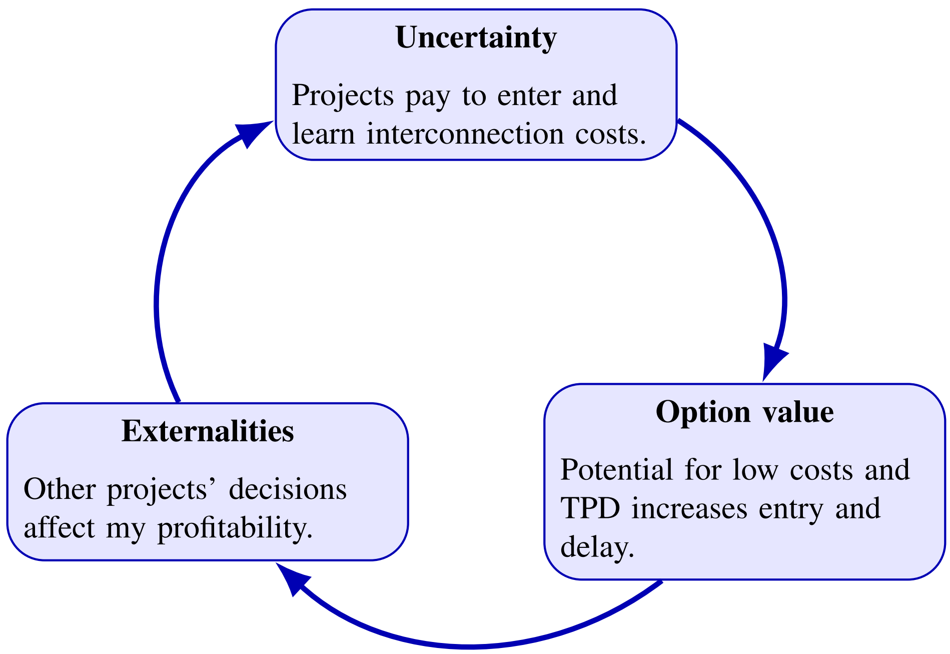

The figure illustrates the forces behind equilibrium queue congestion. Projects enter the queue to learn costs. Entering the queue increases the volume of projects studied. The higher volume of projects increases costs for other projects. These increases in costs further raise the option value of entering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

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

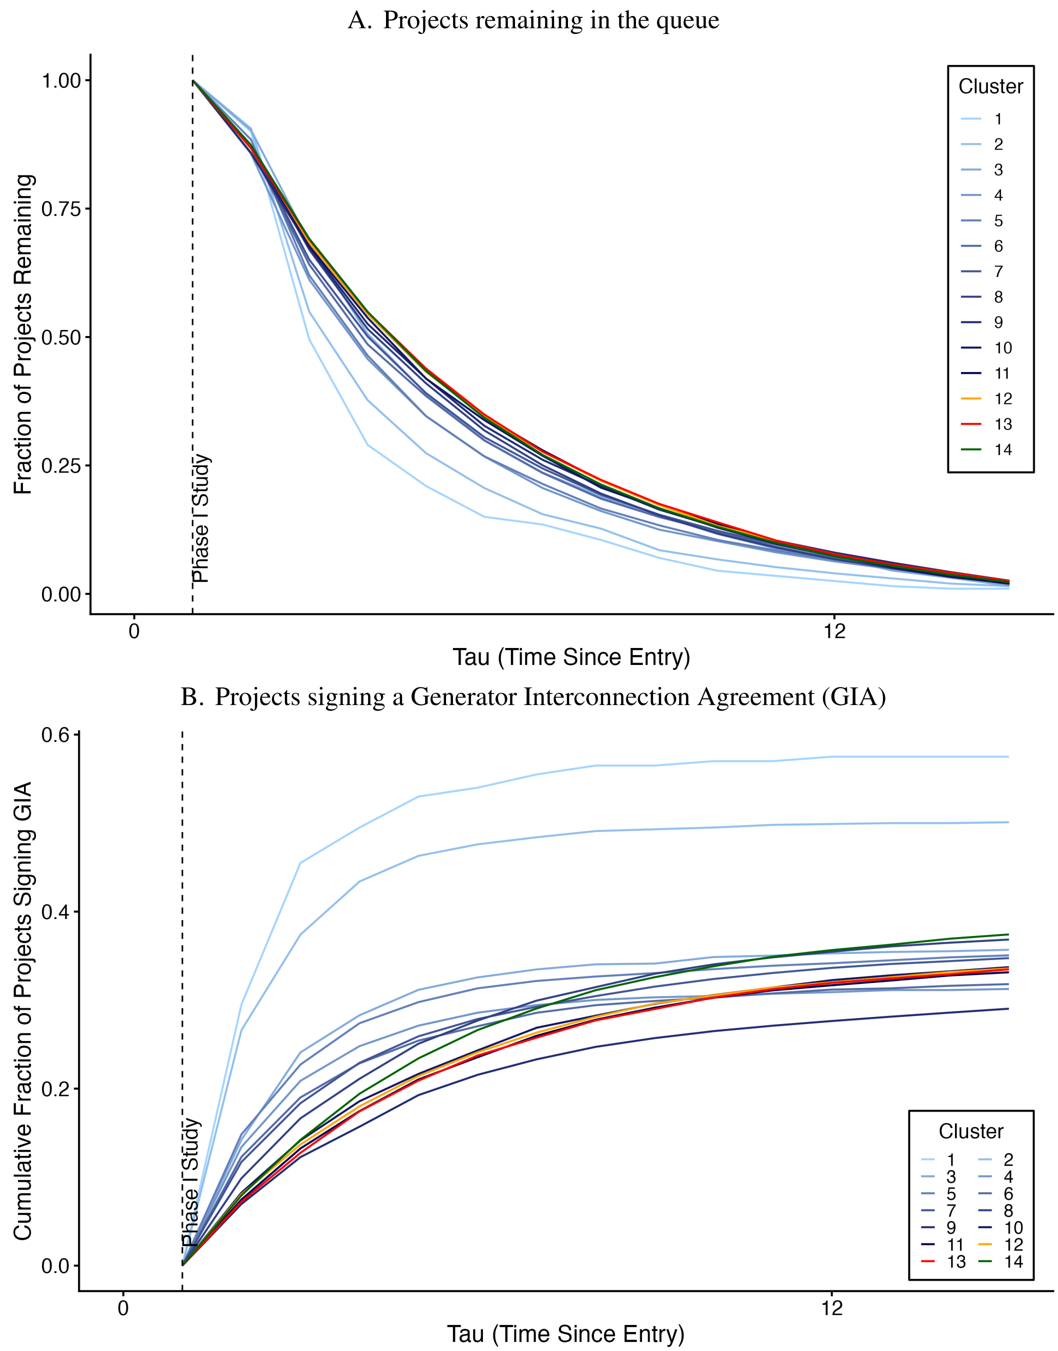

This figure emulates figure 4 with simulated data, based on preliminary estimates. The sample of unique projects is 99,700 comprised of a single queue of 997 projects, simulated 100 times. It shows the share of projects remaining in the queue (panel A) or signing a Generator Interconnection Agreement (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 Receipt, Simulated Data

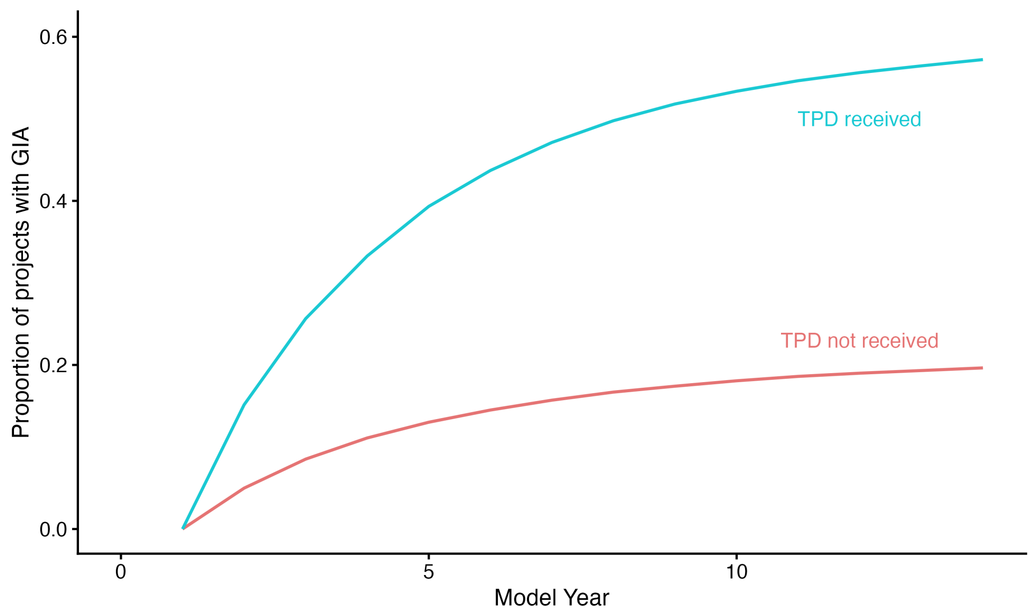

This figure emulates figure 6 with simulated data, based on preliminary estimates. The sample of unique projects is 99,700 comprised of a single queue of 997 projects, simulated 100 times. The model It plots the cumulative proportion of projects that have executed a Generator Interconnection Agreement (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.

8 Tables

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

	Initial Log POI Cost($/kW)	Initial Log Network Cost($/kW)	Reassessed Log POI Cost($/kW)	Reassessed Log Network Cost($/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) × Network cost Quartile 1			0.288***
			(0.088)
TPD received (=1) × Network cost Quartile 2			0.418***
			(0.084)
TPD received (=1) × Network cost Quartile 3			0.283***
			(0.089)
TPD received (=1) × 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	$\delta^{net,init}, \sigma_{\nu}^{net,init}$	Initial network cost estimate
	$\delta^{poi,init}, \sigma_{\nu}^{poi,init}$	Initial POI cost quote
	$\delta^{net}, \sigma_{\nu}^{net}$	Network cost transitions
	$\delta^{poi}, \sigma_{\nu}^{poi}$	POI cost transitions
	δ^{tpd}	Transmission plan deliverability allocation
Aggregate state	$\delta^{agw}, \sigma_{\nu}^{agw}$	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	1.715***	0.055
Inverse solar entry cost scale	6.742***	0.598
Wind mean entry cost	2.381***	0.190
Inverse wind entry cost scale	3.391***	0.775
Battery mean entry cost	2.188***	0.056
Inverse battery entry cost scale	11.649***	0.630

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	3.695^{***}	0.616
Scale of persistent profit shock	-0.027	21.392
Common payoffs		
Transmission Plan Deliverability (0/1)	1.267^{***}	0.278
Land cost (1/acre)	2.419	3.000
Capital cost	-0.657^{***}	0.125
Technology-specific profits		
Solar	-2.552^{***}	0.535
Wind	-1.946^{***}	0.622
Battery	-0.739^{**}	0.349

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: CAISO 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.[9] 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

[9]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).

to depend on the presence of other nearby projects. CAISO attributes all network costs to 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 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.[10] 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

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

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 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 **??**, after introducing our data.

B Appendix: Data

B.1 Data from CAISO

B.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

$$\log(\text{CapEx}_t) = a \cdot e^{-b(t-2035)} + c. \tag{6}$$

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 B1 shows these projections alongside the original cost series. The fit of this simple functional form is remarkably good.

B.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.

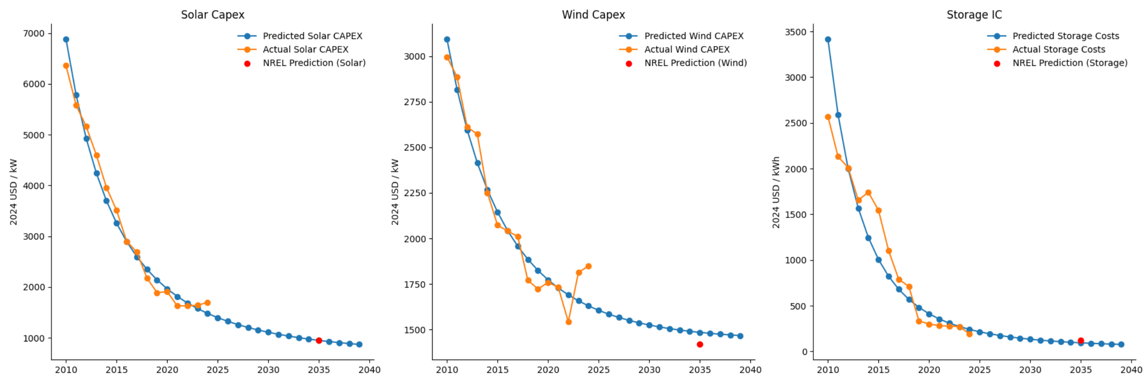

Figure B1: Capital Cost Time Series

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 adjustments where the clustering and geography did not fully align.

Figure B2: Correspondence of Study Areas to Shared Network Upgrades (Cluster 14)

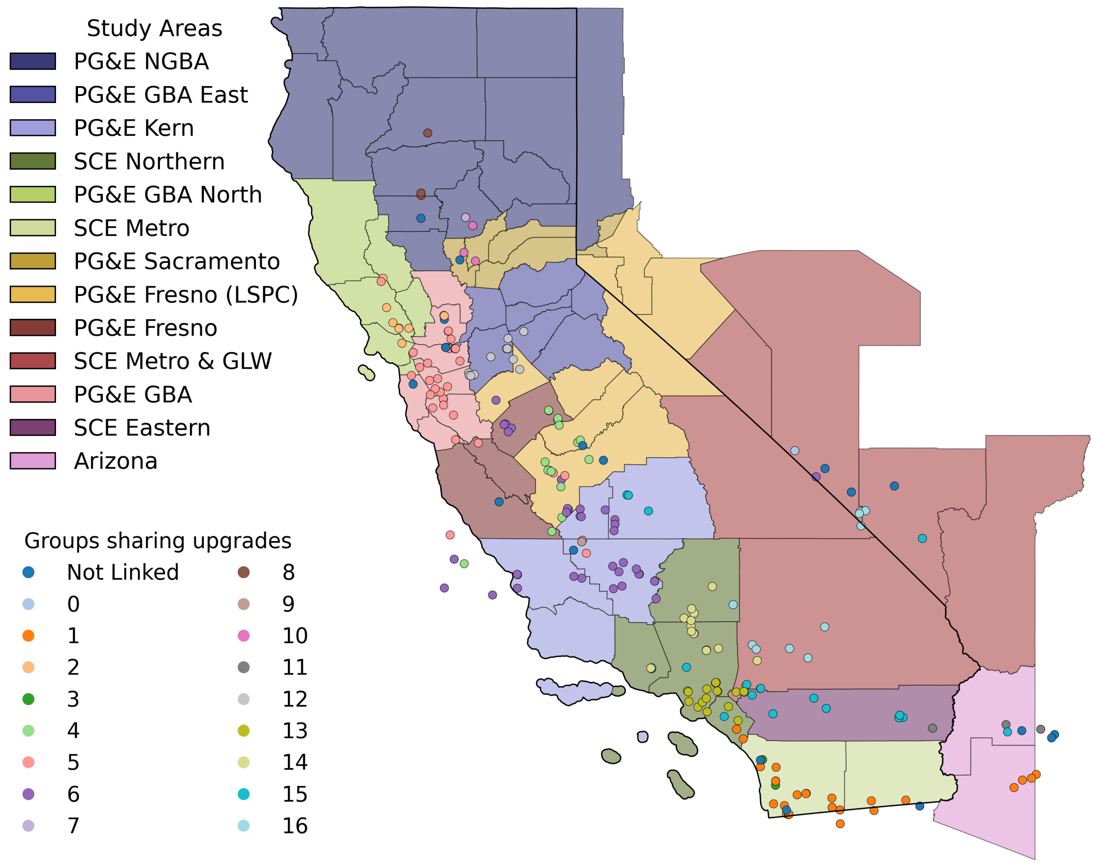

Notes: The figure shows each cluster of projects as a result of using Louvain clustering using points. Points which have no common upgrades with others are considered in the Not Linked cluster. The black outline highlights California's state boundary. Each shaded portion constitutes an area in the model.

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 B3: Probability of Projects Sharing Upgrade Conditional on Areas

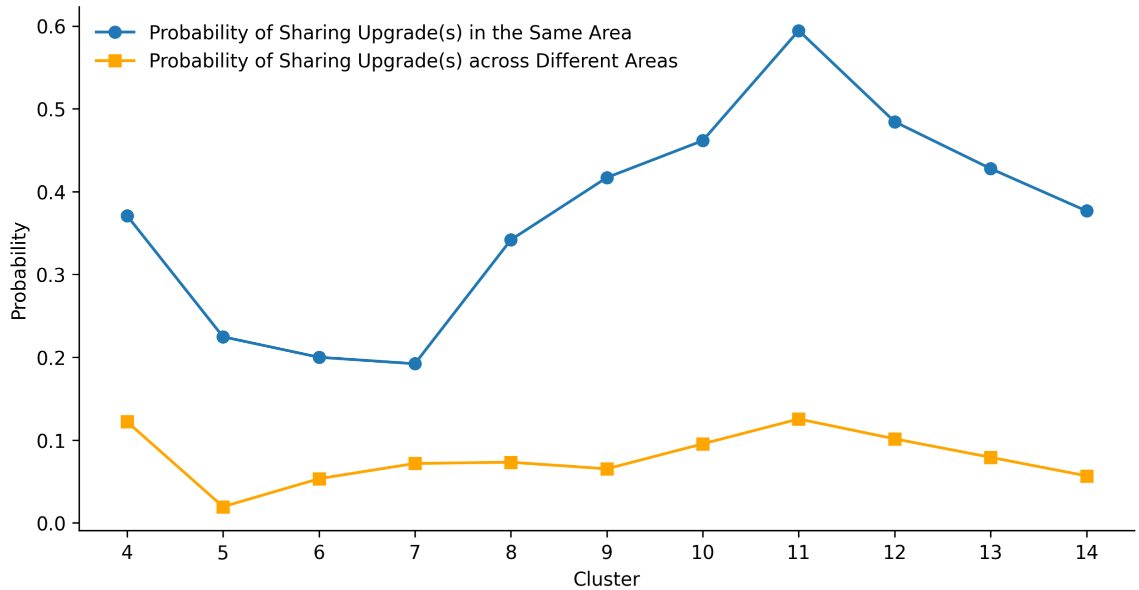

Notes: This figure plots the probability that two energy projects share a transmission upgrade, conditional on geographic area. The blue line reports the within-area linkage probability, while the orange line reports the cross-area linkage probability. Two projects are defined as sharing a link if either 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 B3 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 B4: AMW by Area

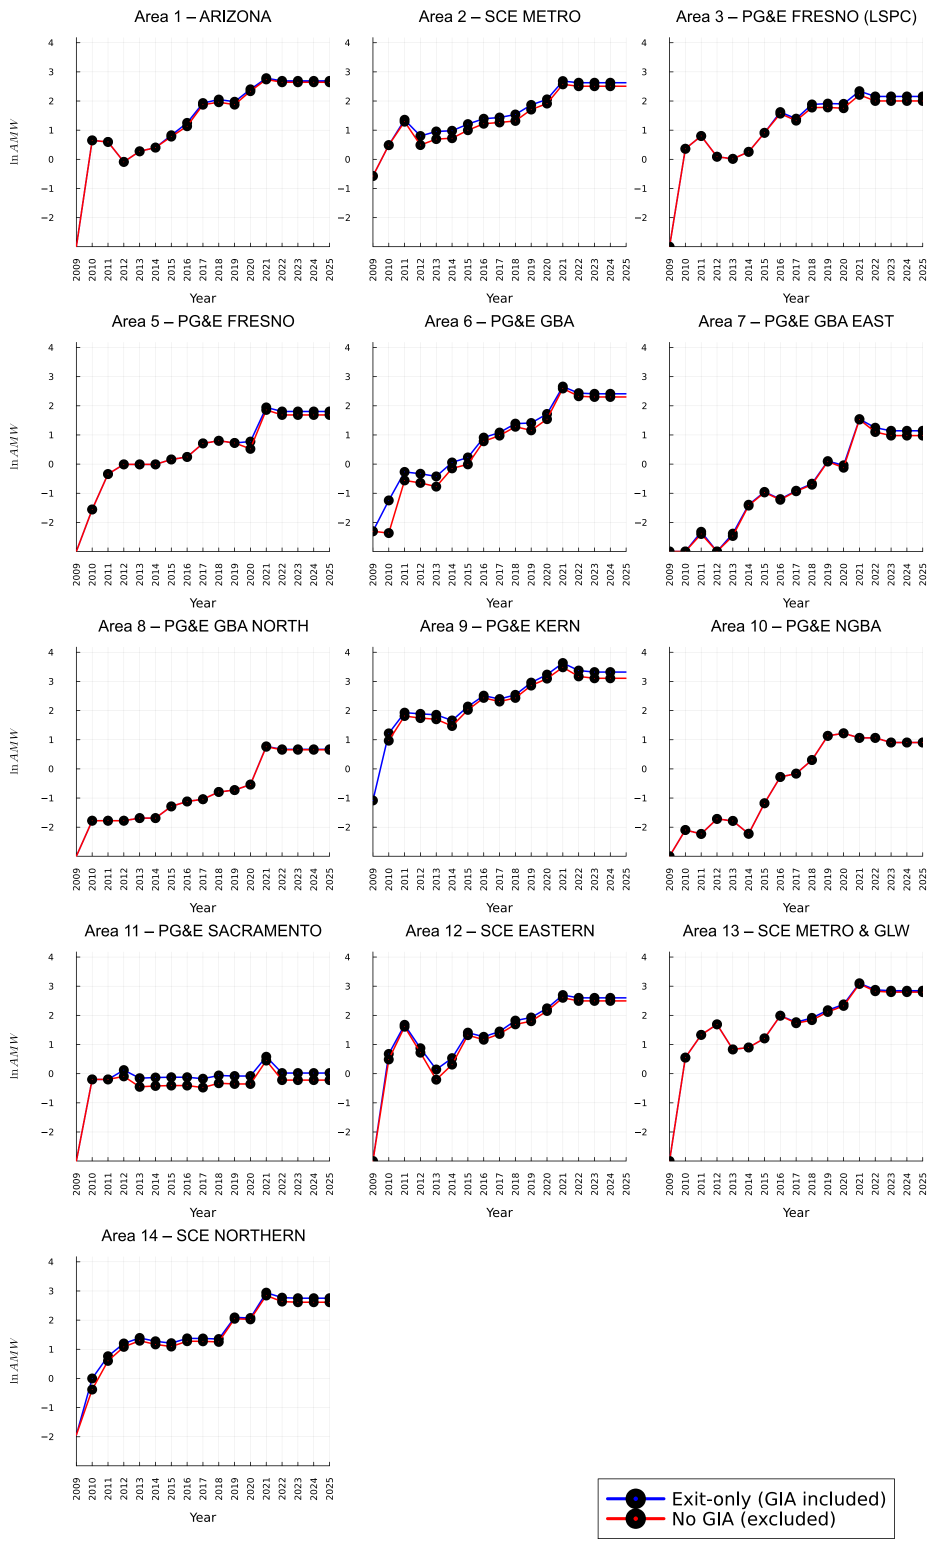

Notes: AMW is the active stock of MW in an area in the queue up to calendar year t. The figure plots the trends in queue congestion ($\ln AMW$) evolution as the queue progresses over time.

C Appendix: Supplementary Descriptive Analysis

Figure C5: Locations of Projects in the Interconnection Queue, Clusters 1 to 14

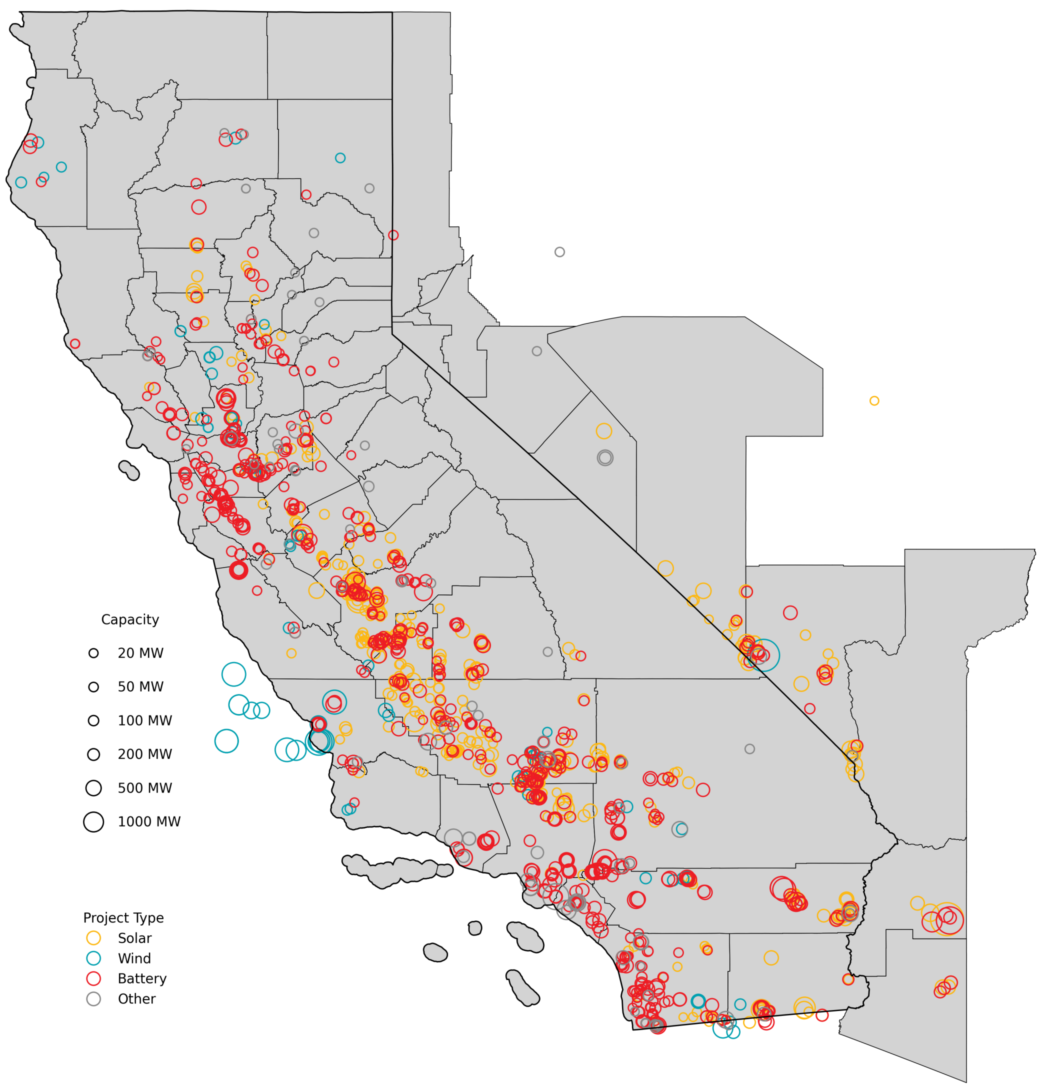

Notes: The figure shows capacity, type, and location for projects proposed to be linked to the CAISO grid. The black outline highlights California's state boundary. Each point is scaled by their capacity at the point of interconnection. The sample includes all valid coordinates of the project from CAISO's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

Figure C6: AR(1)

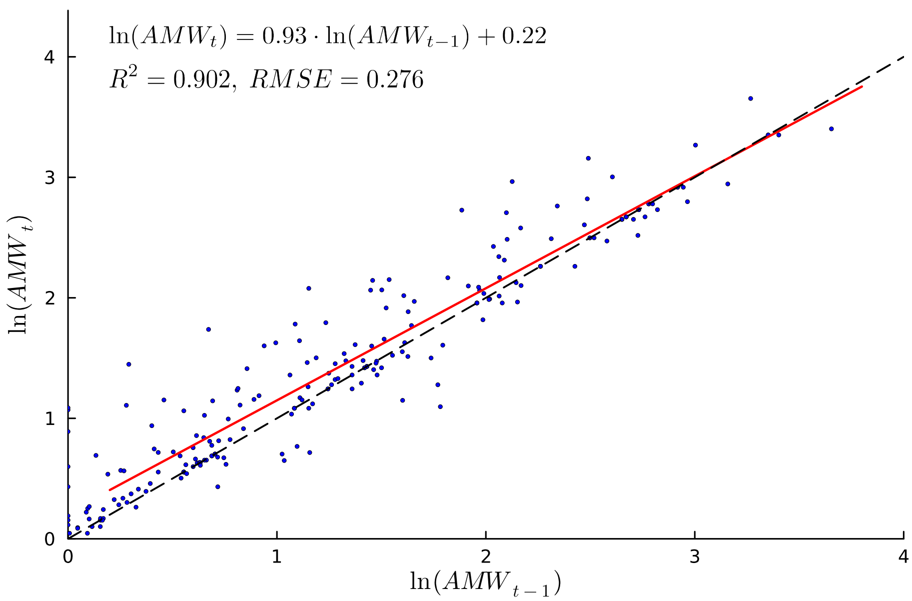

Notes: AMW_t is the active stock of MW in the queue up to calender year t. The figure shows the persistence in AMW_t

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

$$IFS(\mathbf{x}_j, \tau) = \begin{cases} 0.15 \cdot MW_j \cdot c_{j\tau} & \text{for } \tau = 1, \text{ to } Continue \\ 0.30 \cdot MW_j \cdot c_{j\tau} & \text{for } \tau = 2, \text{ to } Continue \\ 1.00 \cdot MW_j \cdot c_{j\tau} & \text{for } \tau \geq 2, \text{ to } Contract. \end{cases} \tag{7}$$

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 (7) 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.[11]

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.

[11]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

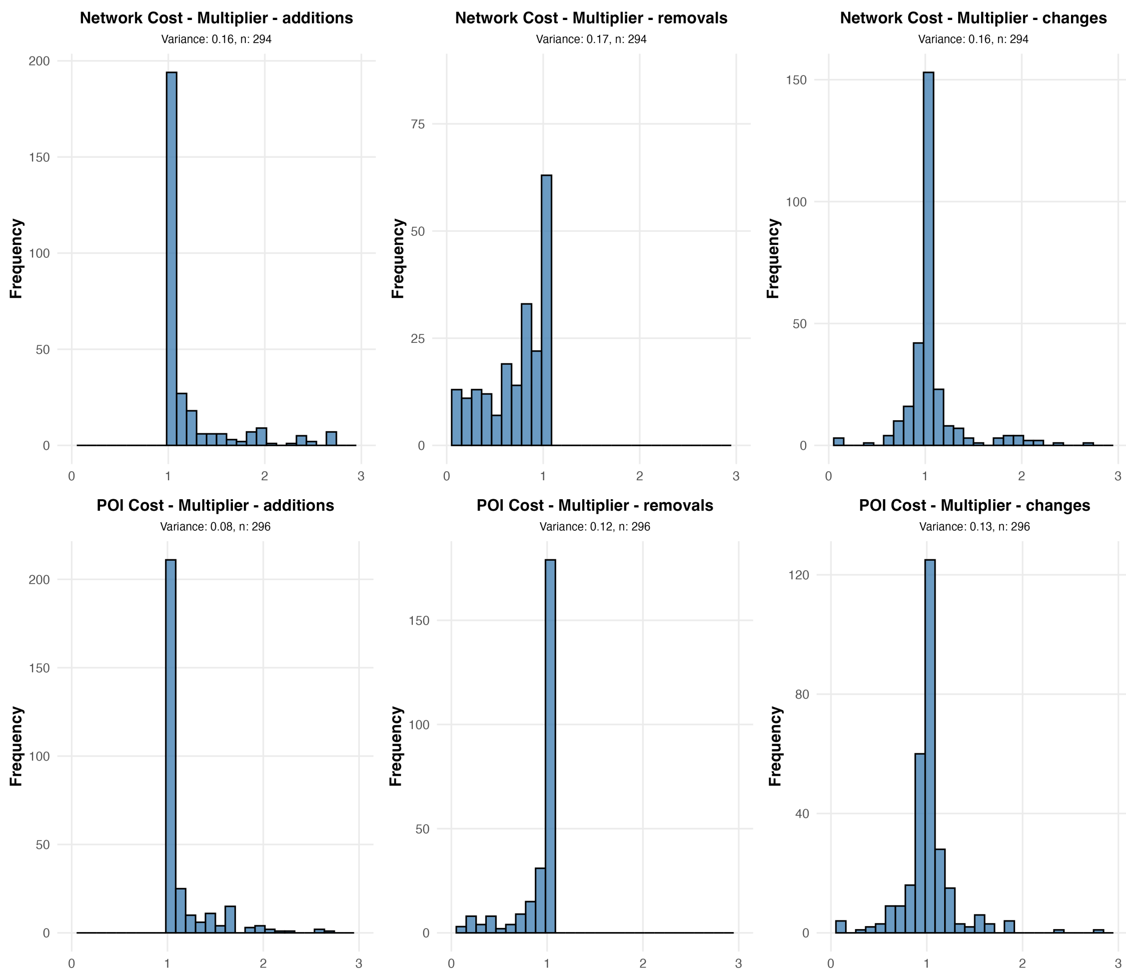

Notes: This figure shows the distribution of the cost multipliers of the upgrade matched sample

Table F2: Regression Results for Network Cost Multipliers

	Additions Multiplier (1)	Removals Multiplier (2)	Changes Multiplier (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 (1)	Removals Multiplier (2)	Changes Multiplier (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.