

An Empirical Analysis of the Interconnection Queue*

Sarah Johnston[†]

Yifei Liu[‡]

Chenyu Yang[§]

University of Calgary

University of Wisconsin-Madison

University of Maryland

August 21, 2026

Abstract

Electricity generation projects applying to connect to the U.S. power grid go through an interconnection queue. Most projects that begin the process do not complete it. Using new data, we find that a long queue delays the necessary engineering studies, and high interconnection costs cause projects to withdraw from the queue. We develop and estimate a dynamic model of the queue and quantify the effects of policy reforms. We find that speeding up the delivery of studies can significantly increase completed capacity, with about 30% of the gain coming from projects re-optimizing their entry, waiting, and completion strategies. Imposing entry fees reduces entry into the queue but not average waiting time, because changes in queue composition and project strategies weaken the direct effect of having fewer entrants. Fee designs we consider mainly deter marginal entrants that contribute little to congestion, and effects on completed capacity are modest.

*We thank Joseph Aldy, Jim Bushnell, Steve Cicala, Jackson Dorsey, Ken Gillingham, Will Gorman, Kate Ho, Charles Hodgson, Koichiro Ito, Ashley Langer, David Popp, Mar Reguant, Andrew Sweeting, Frank Wolak, and participants at many seminars and conferences for helpful comments and discussions. We also thank Christian McDowell and four industry experts for insight into the electricity market. We also thank the numerous undergraduate research assistants at UW-Madison and the University of Maryland who helped with data collection. We gratefully acknowledge financial support from the National Science Foundation (SES-2215063), the NBER-Sloan Economics of Innovation in the Energy Sector Initiative, the University of Maryland Grand Challenges Grant Program, and the University of Wisconsin - Madison Office of the Vice Chancellor for Research and Graduate Education with funding from the Wisconsin Alumni Research Foundation.

[†]Department of Economics, University of Calgary, Calgary, AB T2N 1N4; sarah.johnston2@ucalgary.ca.

[‡]Department of Agricultural and Applied Economics, University of Wisconsin-Madison, Madison, WI 53706; yifei.violet.liu@wisc.edu.

[§]Department of Economics, University of Maryland, College Park, MD 20740; cyang111@umd.edu.

1 Introduction

Electricity production accounts for over a quarter of both U.S. and global carbon emissions (EPA, 2026; IEA, 2026, 2025). Many countries’ climate goals center on transitioning to a low-carbon grid while electrifying heating and transportation. Meeting these goals will require large-scale investment in wind- and solar-powered generators.

Yet, connecting an electricity generation project to the U.S. power grid is difficult. The process, known as interconnection, takes several years and must be completed before the developer can build the project. It can also be costly: connecting projects must often pay to upgrade the transmission system because the local grid is at capacity (Plumer, 2023). Renewable energy developers cite interconnection as the single biggest hurdle they face (Driscoll, 2022), and less than a quarter of the wind and solar projects that start the process complete it (Rand et al., 2025). We study the design of this interconnection process.

A project (a proposed generating facility at a specific site) seeking to connect to the transmission grid joins a waitlist known as the interconnection queue. The grid operator then conducts a series of engineering studies to determine the project’s “interconnection cost.” This cost covers both the equipment required to connect the project to the grid and the transmission system upgrades (if any) needed to accommodate it on the broader network. The date the project joined the queue determines priority for these studies. A project usually needs three studies to arrive at a final interconnection cost, after which it either pays this cost and connects, or withdraws. Projects can withdraw at any time.

Each additional project in the queue can delay the studies of projects behind it. The cost of being in the queue is low, however, and may be insufficient to offset this congestion externality. Grid operators devote substantial resources to processing requests that are ultimately withdrawn. These concerns have motivated reforms such as FERC Order 2023, which raises study deposits and introduces penalties for withdrawing from the queue.

To study this process, we hand-collect novel data from PJM, the largest of the U.S. regional grid operators. PJM serves 67 million people in parts of the Mid-Atlantic, Midwest, and Southern United States (PJM, 2026). Our data cover all generator interconnection requests from 2008 through 2020: 4,174 requests and their 7,492 engineering studies.

Project developer complaints about delayed studies and high interconnection costs are borne out in the data. We find that the studies often take far longer than PJM’s official

timeline, with the later studies being the most delayed. For example, PJM expects the third and final interconnection study to take 6 months, but the median wait in the data is 15 months, and the 90th percentile is 36 months. While the median second study interconnection cost is \$0.05 million per megawatt (MW), the 90th percentile is \$0.40 million per megawatt, roughly a quarter of the installation cost for wind and solar projects. We also find that observably similar projects can have very different interconnection costs, a finding consistent with developers' complaints that these costs are unpredictable (Caspary et al., 2021).

We also show that high interconnection costs cause projects to withdraw from the queue. Withdrawal rates rise with cost at each study, but cost may be correlated with unobserved project quality. We therefore instrument for cost with the change in cost across studies, which is difficult to predict. A second-study interconnection cost above \$0.1 million per megawatt increases the probability of withdrawal before the third study by 38 percentage points, 81 percent of the mean withdrawal rate.

We next quantify three potential externalities in the queue. The first is a congestion externality: are the necessary engineering studies returned more slowly when more projects are ahead in the queue? The second is a geographic cost externality: do contemporaneous projects in the queue or completed interconnections affect the interconnection costs of other nearby projects? The third is a local output externality: do completed interconnections reduce local wholesale electricity prices and thus the expected profits of nearby projects in the queue?

We find the congestion externality is quantitatively important. We estimate how much a project's probability of receiving a new study in a given quarter declines with the number of higher-priority projects ahead of it. This effect is identified because PJM, not the project, controls when studies arrive. A 10 percent increase in the number of higher-priority projects reduces a project's probability of receiving the third and final study by 12.5 percent, on average.

In contrast, we do not find that the other two externalities are quantitatively important in our setting. We test for cost externalities via two channels: whether a project's cost is affected by others still in the queue, either through nearby capacity ahead of it or withdrawals from its cost-share group; and whether its cost is lowered by spare transmission capacity left by costly past interconnections. In both cases, the estimates are small and mostly statistically

insignificant. This finding is likely due to PJM’s cost-sharing rules: cost-sharing is extensive but confined to small groups, so most projects do not affect one another’s costs. Even within a cost-share group, a withdrawal does not systematically raise or lower remaining projects’ costs. Similarly, using a staggered difference-in-differences design, we find that a new project beginning operation does not reduce wholesale prices in its location relative to other locations in PJM. The likely explanation is PJM’s deliverability requirement, under which connecting projects must upgrade the transmission system to mitigate local grid congestion. These findings motivate our focus on the congestion externality and policies that target waiting times.

We then develop a structural model of the queue to quantify how policy reforms affect completed capacity and welfare. We model withdrawal decisions as an optimal stopping problem. A project waits in the queue for studies and forms beliefs about when the next study will arrive and what its interconnection cost estimate will be. The continuation value of a project depends on its characteristics and the status of the queue, which we measure both PJM-wide and at the level of the local transmission owner, whose engineering input the studies rely on.

To handle the non-stationarity in the data, we develop a tractable queuing equilibrium concept. In this equilibrium, the composition and size of the queue are jointly determined by optimal project strategies and the PJM study process. Project beliefs are consistent with the state transitions implied by the queue in equilibrium; projects still face uncertainty over their own study arrivals and cost realizations. We prove that such an equilibrium exists.

We estimate the model in three steps. First, we estimate models of new study arrival rates, interconnection costs, and additional information revealed by each study. These estimates allow us to construct project beliefs. Second, we embed these beliefs in the project’s optimal stopping problem. We then use the observed withdrawal and completion decisions to recover the waiting cost and payoff for completing interconnection. We prove that project payoff heterogeneity and waiting costs are nonparametrically identified by variation in interconnection costs across studies. Intuitively, because these cost changes are difficult to predict, the cost a project learns at each study acts as a shock to its continuation value, and how its withdrawal and completion decisions respond to these shocks across successive studies pins down payoff heterogeneity and the cost of waiting. Third, we use the implied value of entering

the queue to estimate the entry cost.

We use the estimated model to simulate the equilibrium effects of policy reforms: speeding up study delivery and increasing the cost to enter or wait in the queue. Faster study delivery is the goal of recent efforts to deploy AI for engineering studies, such as PJM’s 2025 partnership with Google’s Tapestry (PJM Interconnection, 2025*b*). FERC Order 2023 takes both approaches: it penalizes transmission providers for study delays, and it raises study deposits and introduces withdrawal penalties (FERC, 2023). The withdrawal penalties function like fees collected at study stages and refunded upon completion.

We find that faster engineering studies substantially increase completed capacity. For 2012–2020 queue entrants, speeding up study delivery by 50% increases completed capacity by about 32%, or 18.1 GW, and shortens average waiting time by 19%. Of the added capacity, 63% (11.4 GW) is renewable. Given a social cost of carbon of \$185 per metric ton, this increase in renewable capacity would create \$2.4 billion in annual benefits.

About 30% of the capacity gain comes from projects re-optimizing entry, waiting, and completion strategies in anticipation of the faster queue. A partial-equilibrium calculation that holds these strategies fixed captures only 12.6 GW of the 18.1 GW gain.

In comparison, the benefits of increasing fees are more modest. Our main counterfactual exercise evaluates three types of payments to enter the queue that differ in refundability: a nonrefundable fee, a deposit returned only upon completion (deposit), and a deposit that is partially returned to projects that withdraw early (staged deposit). We evaluate a range of entry payment levels; here, we describe the effects at \$20,000 per megawatt.¹ The nonrefundable fee reduces entry by 1.7%, but mainly deters projects that wait less, shifting the composition of the queue. Completed capacity falls by 0.8% and average waiting time is essentially unchanged in equilibrium. The deposit reduces entry less and increases completed capacity by 0.2%. Entering projects wait longer, so average waiting time does not fall.

The third design, the staged deposit, whose refund schedule mirrors PJM’s ongoing reform, increases completed capacity the most, by 0.8%. Like the deposit, the staged deposit is refunded in full upon completion, but differs in that a project that withdraws before requesting the second study can recover half its deposit. Relative to the deposit, the staged deposit

¹This level is high enough to distinguish the designs and comparable to the fee PJM adopted in 2026 for a temporary program to accelerate some projects (PJM Interconnection, 2026). The level in PJM’s wider reforms is \$4,000 per megawatt, at which the effects are smaller but in the same direction.

shortens waiting times because more projects withdraw after the first study, thinning the queue. It also retains the least revenue: \$3.6 billion, compared with \$4.0 billion under the deposit and \$5.2 billion under the nonrefundable fee.

Effects on completed capacity remain modest when PJM rebates the forfeited deposits to completers or charges flat rather than per megawatt fees. Policies that rebate all forfeited deposits to completing projects retain no revenue for PJM and provide a useful benchmark for what a budget-neutral fee policy might achieve. We find that completed capacity rises by 3.6% under the budget-neutral deposit and 4.0% under the budget-neutral staged deposit. Finally, flat-fee versions of the same three designs substantially reduce entry by small projects but have similarly modest effects on completed capacity.

The key intuition for the small effects is that the projects creating the most congestion are also the most valuable. Higher-value projects wait longer in the queue, slowing studies for others; these projects are also the most likely to complete. Entry fees therefore screen out marginal entrants that contribute little to congestion. Under the deposit, the gains come from the completion subsidy created by the refund rather than from reduced congestion. The staged deposit is more effective at reducing congestion because it screens on the first study’s results rather than on expected value at entry. Screening later in the queue, however, increasingly burdens projects likely to complete. All three designs nonetheless retain revenue that could, in principle, fund policies to speed up study delivery.

Related Literature

The interconnection queue is a critical but understudied bottleneck to electricity market entry. While it has received attention in the energy policy literature (Gergen, Cannon Jr and Torgerson, 2008; Alagappan, Orans and Woo, 2011; Gorman et al., 2024), it has been largely absent from formal economic analysis, primarily due to data constraints. We collect novel data on the PJM queue and make them publicly available. In a companion paper (Johnston, Liu and Yang, 2026), we use these data to study the determinants and consequences of interconnection costs. In concurrent work, Gowrisankaran, Langer and Ryan (2026) study cost externalities under a different queue design in California. Another closely related paper is Gonzales, Ito and Reguant (2023) which studies how transmission grid expansion enables renewable entry; our paper instead focuses on the interconnection queue as a barrier to entry.

We also contribute to a broader literature on how transmission constraints affect competition, emissions, and investment (Wolak, 2015; Ryan, 2021; Davis and Hausman, 2016; Fell, Kaffine and Novan, 2021; LaRiviere and Lyu, 2022; Doshi, 2026), as well as on the institutional barriers to expanding transmission infrastructure (Davis, Hausman and Rose, 2023; Hausman, 2025).

Our structural approach contributes to the energy and environmental economics literature studying investment and industry dynamics (Ryan, 2012; Gowrisankaran, Reynolds and Samano, 2016; Fowlie, Reguant and Ryan, 2016; Blundell, Gowrisankaran and Langer, 2020; Butters, Dorsey and Gowrisankaran, 2021; Elliott, 2021; Gowrisankaran, Langer and Zhang, 2022; Leisten and Vreugdenhil, 2023; Chen, 2025). By quantifying how interconnection queue policy reforms can improve renewable entry, we also contribute to the literature on public policy and renewable energy investment (Metcalf, 2010; Hitaj, 2013; Johnston, 2019; Aldy, Gerarden and Sweeney, 2023; Deschenes, Malloy and McDonald, 2023).

We also contribute to the empirical literature on dynamic assignment mechanisms (Agarwal et al., 2021; Waldinger, 2021; Verdier and Reeling, 2022; Liu, Wan and Yang, 2024; Huitfeldt, Marone and Waldinger, 2024) by studying a queuing problem in a novel and consequential market. Our central methodological contribution is a tractable equilibrium concept for a non-stationary queue with time-varying unobserved heterogeneity. The size and composition of the queue shifted substantially over our sample period, making standard oblivious equilibrium (Weintraub, Benkard and Van Roy, 2008) inapplicable. We build on the nonstationary oblivious equilibrium of Weintraub et al. (2010). The non-stationarity in the data is driven by factors that projects may reasonably foresee, such as the continued decrease in wind, solar, and battery costs and the increase in queue congestion. We adopt a finite-horizon assumption, as in Igami (2017) and Yang (2020), where agent beliefs in equilibrium are consistent with the trend of non-stationarity. We provide theoretical results on both the existence of the equilibrium and identification of agent preferences in this environment.

2 PJM Interconnection Process

The grid operator, PJM, runs the interconnection process, which involves two other parties: developers and transmission owners (TOs). Developers (e.g., NextEra Energy) enter their projects (e.g., 2.5 MW Front Royal Solar Field) in the queue and pay the interconnection cost identified in the studies. These developers can be either independent power producers

(more common for renewables) or regulated utilities (more common for natural gas). PJM conducts the interconnection studies in conjunction with the local TO, which supplies the design standards and engineering input and constructs the upgrades (PJM Interconnection, L.L.C., 2020). We assume the cost developers pay reflects the competitive cost of these upgrades. PJM is a nonprofit grid operator, and TOs are regulated utilities with no incentive to inflate cost because, during our sample period, interconnection-funded upgrades did not enter their rate base.²

2.1 Queuing Rules

Projects can enter the queue at any time, but each year is typically divided into two windows. Projects that enter in the same 6-month window are put in the same cohort and receive up to three studies (feasibility, system impact, and facility study) sequentially. Through these studies, projects learn increasingly accurate information about the interconnection cost. At any point in the process, a project may withdraw from the queue. PJM may require just one or two studies if it determines that a project is required neither to make significant network upgrades nor to share costs with other projects. After the last study is issued, the project chooses to leave the queue or sign an interconnection service agreement in which it agrees to pay the final interconnection cost, thus completing the interconnection process.

Developers must pay to enter the queue and to advance through it. Before entry, a developer must secure land sufficient to build a project (PJM, 2023*b*), typically through exclusive options to lease the land. Developers must also pay a deposit at entry and to advance to each subsequent study. The deposit amount depends on the request size and stage, with larger requests typically requiring higher deposits. For the median request size of 20 MW, the initial deposit ranged from \$10,000 to \$15,000 depending on when within the 6-month window the request entered, while the deposits for the second and third studies were \$10,000 and \$50,000, respectively. Projects that withdraw have their deposit returned, less a 10 percent non-refundable portion and any study costs already incurred (PJM, 2023*a*).

The official timeline for the interconnection process is rigid. For projects that apply within the same cohort, PJM starts conducting the first studies (feasibility studies) one month after

²PJM also plans transmission separately from the interconnection process, through the Regional Transmission Expansion Plan (RTEP). This investment is driven by reliability needs and funded by consumers through transmission rates. Over our sample period, the RTEP did not consider projects waiting in the queue (PJM, 2020).

the closing of the time window. Within three months, projects are supposed to receive their first studies. At this point, projects have another month to decide whether to advance to the second study (system impact study). The second study then takes four months, at which point projects have one month to decide whether to request the third and final study (facility study). The third study takes six months. Finally, projects withdraw or agree with PJM on final details and sign the interconnection service agreement, over a 6.5-month period (PJM, 2021).³

Despite this timeline, significant delays in delivering studies can occur due to the number of backlogged interconnection requests and a lack of staff capacity (Shoemaker, 2021; Petersen, Siegner and Coequyt, 2026). Projects have little control over when PJM delivers the study. These delays have been the subject of formal developer complaints; for example, a solar developer in PJM filed a complaint with the Federal Energy Regulatory Commission (FERC) after waiting more than two years for a second study (Hale, 2021).

2.2 Study Information

Each study reports an interconnection cost estimate along with engineering information. The interconnection cost has two components: the direct cost to connect the project to the grid, and the cost of network upgrades to relieve system overloads the new project triggers.⁴ The first study reports preliminary direct interconnection costs and flags whether the project may share network upgrade costs with others. The second study reports the network upgrade costs, including the project’s share, and updates the direct costs. The third study finalizes engineering specifications and costs.

The engineering tests reported in each study determine which network upgrades the project is responsible for and whether it shares that responsibility with other projects. Three tests—generator deliverability, multiple facility contingency, and short circuit analysis—jointly quantify how much the project will overload the grid; they are conducted in either the first or the second study, depending on project characteristics and local transmission network conditions. Two additional tests—short circuit dynamic analysis and system protection analysis—identify

³This agreement also sets a commercial operation date. After signing the interconnection service agreement, the project may suspend the process for up to 3 years (up to 1 year if the suspension has a negative impact on subsequent requests) (PJM, 2025), though these suspensions are rare.

⁴We combine the two components into a single cost measure for our analysis; see Johnston, Liu and Yang (2026) for more detail and descriptive statistics on each.

whether the project triggers the same network violation as other projects in the queue, in which case PJM assigns it to a cost-share group jointly responsible for the upgrade. Group assignments are typically first reported in the second study, alongside the network upgrade cost.

Studies report additional information that affects a developer’s assessment of the project. They specify the types of equipment required to connect the project, which can affect the costs of permitting and right-of-way access. PJM expressly states in studies that the interconnection costs do not consider these costs, which can be significant hurdles (Liu, 2025). Studies may also flag costly requirements even when no interconnection cost is assessed, like bringing facilities up to the transmission owner’s published standards.⁵ In the structural model, this additional information is part of the unobserved payoff heterogeneity, realized when the relevant study arrives.

2.3 Timing of Interconnection within Project Development

Projects apply for interconnection early in project development. The permitting and interconnection processes occur simultaneously, and developers perceive interconnection as the primary bottleneck (Nilson, Hoen and Rand, 2024). Projects signing long-term contracts to sell their power—nearly all renewable projects—do so while in the queue or after receiving the terminal study. After signing the interconnection service agreement, the project and the required interconnection infrastructure are constructed, typically in less than a year, far shorter than the time spent in the queue.⁶

The slow, unpredictable interconnection process can disrupt other steps in project development. Power buyers are willing to sign long-term contracts a few years before a project begins operation, but developers must apply for interconnection much earlier. Developers are thus left trying to forecast demand a few years in advance and enter projects into the queue accordingly. A long wait may also cause signed long-term contracts or a developer’s land lease option to expire, prompting the developer to renegotiate the lease or abandon the project.

⁵For example, the impact study of a gas power plant (AA2-030) does not assess an interconnection cost because the project is an expansion of an existing plant, but the project “is required to construct all connection facilities in accordance with the TO published standards. ComEd would review and approve all customer relay protection design drawings and relay settings.”

⁶For wind projects, construction typically takes six months to a year (AWEA, 2019); solar projects follow similar timelines. Network upgrades can take longer, but projects can sometimes begin operating before all upgrades are completed.

Projects also face risks unrelated to interconnection that can cause cancellation. A project may not be selected in a competitive solicitation for a power purchase agreement, fail to secure a necessary permit, or encounter local opposition that blocks siting. Each additional quarter in the queue is another quarter of exposure to these shocks, so a slower interconnection process increases cancellations from these channels.

2.4 Queue Design Reforms

Recent regulatory reforms aim to address the growing interconnection backlog. In 2023, FERC, the federal regulator overseeing U.S. transmission policy, approved Order No. 2023, which aimed to shift queue priority from “first-come, first-served” to “first-ready, first-served” (Hale and Christian, 2023). The Order has four main components: (a) raising the cost to enter and remain in the queue, (b) imposing penalties for withdrawing from the queue, (c) imposing penalties for delays in completing studies, and (d) transitioning from studying each request serially to cluster studies, under which requests are grouped into large clusters and receive a single study per cluster. The study-delay penalties in (c) fall on the transmission provider, which in PJM’s case is PJM itself rather than the transmission owners; the penalties cannot be recovered through transmission rates, though PJM may seek FERC approval to recover a penalty from a transmission owner at fault (FERC, 2023). Implementation is ongoing: PJM, for example, opened its first reformed cluster cycle in 2026, with applications due April 27 (PJM Interconnection, 2025*a*).

3 Interconnection Request Data and Descriptive Evidence

Our data cover every interconnection request to PJM from 2008 to 2020 and the engineering studies those requests received. Four facts in these data underpin the model: entry to the queue increases over time, waits for later studies are long and uncertain, changes in interconnection cost estimates across studies are difficult to predict, and withdrawals concentrate in the months right after a study arrives. Using cost changes across studies as an instrument, we show that high interconnection costs cause withdrawal.

3.1 Data Sources and Construction

Our main data cover the 4,174 interconnection requests to PJM from 2008 to 2020.⁷ Each request is to connect a single project at a specific point on one transmission owner’s network. A project is either a new plant or a capacity increase at an existing plant; the study process is the same for both. Requests are almost always distinct projects rather than resubmissions: only 6.8% of requests enter at the site of a previously withdrawn request, and 4.1% do so within two years of that withdrawal.⁸

We hand-collected data from the PDFs of 7,492 engineering studies. From each study we record the interconnection cost estimate, the engineering tests performed and their results, and whether the project is assigned a network allocation.⁹ We also record each study’s issue date and, from PJM, the queue entry date, entry cohort, and withdrawal or operation date. These dates let us reconstruct, for each quarter, the set of requests waiting for each study and each request’s priority relative to others, both PJM-wide and within its TO. Table 1 reports summary statistics.

3.2 Queue Composition and Growth

Solar, natural gas, and wind dominate the PJM queue, together making up 85 percent of requests from 2008 to 2020. Natural gas projects tend to be much larger: they make up 17 percent of requests but 42 percent of requested capacity. The composition of requests shifted over our sample, with solar gaining share after 2015 and batteries after 2017. We refer to wind, solar, and battery projects as “renewable” and all others as “non-renewable”.

The number of interconnection requests has increased dramatically over time, driven by falling installation costs for wind and especially solar (Seel et al., 2024; Wiser et al., 2024). Figure 1 plots, by year, new requests (dotted line) and the average number of projects in the queue (solid line). The increase began around 2015 and was concentrated in renewables.¹⁰ Because renewables are smaller on average, the rise in requested capacity was less pronounced than the rise in requests: 42 GW of new requests in 2008 compared to 69 GW in 2020. Con-

⁷Earlier studies use irregular formats that make it difficult to identify the relevant interconnection costs.

⁸We treat requests with coordinates within 1 km of each other as the same site.

⁹Network allocation indicates whether a project is explicitly assigned a network upgrade cost in the study; this cost may be zero.

¹⁰The spike around 2010 was due to a temporary federal program that gave solar and wind developers the option to claim a cash grant in lieu of the non-refundable investment tax credit (see Aldy, Gerarden and Sweeney (2023) for program details).

Table 1: Summary Statistics

	Study 1		Study 2		Study 3	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Cost per MW of capacity	0.12	0.40	0.17	0.47	0.11	0.21
Wait time (mos.)	4.9	1.5	11.3	8.1	15.1	9.8
Size in MW	97	196	105	190	157	257
Capacity increase at existing plant	0.21	0.41	0.23	0.42	0.15	0.36
Solar	0.59	0.49	0.60	0.49	0.61	0.49
Natural Gas	0.17	0.37	0.16	0.36	0.19	0.39
Wind	0.09	0.29	0.10	0.30	0.12	0.33
Battery	0.10	0.30	0.10	0.29	0.05	0.21
Coal, oil, diesel	0.02	0.14	0.02	0.12	0.01	0.08
Other	0.03	0.18	0.03	0.17	0.02	0.15
Network allocation	0.04	0.20	0.59	0.49	0.62	0.48
Receive engr. tests	0.81	0.39	0.88	0.32	0.04	0.19
Distance to substation (km)	5.13	7.00	5.26	7.64	5.70	8.54
Ordinance	0.28	0.45	0.31	0.46	0.32	0.47
Observations	4,168		2,495		829	


The sample is projects that entered the queue in 2008-2020. Cost per MW of capacity is interconnection cost estimate, in millions of 2020 dollars, divided by the project’s size in MW. Wait time for the 1st Study is wait in months for the first study after joining the queue. Wait time for the 2nd Study is wait in months for second study after receiving the first study. Wait time for the 3rd Study is wait in months for third study after receiving the second study. Reported wait times are capped at the 90th percentile (36 mos. for the second and third studies) to limit the influence of outliers. The corresponding uncapped means are 5.3, 14.9 and 18.0. Capacity increase at existing plant is an indicator for a request to increase maximum generating capacity at an existing power plant; the size for these requests is the incremental capacity, not the total plant capacity. Network allocation is an indicator of whether a project receives an explicit allocation of network upgrade costs, even if the allocated amount is \$0; network upgrade costs are typically first reported in the second study. Receive engr. tests is an indicator for receiving any of three engineering tests: generator deliverability, multiple facility contingency, and short circuit analysis. Distance to substation is the distance to the nearest connected substation in km. Ordinance is an indicator for a local ordinance restricting renewable energy development.

sistent with a technology-driven explanation, similar increases in renewable requests occurred across all U.S. regional grid operators (Gorman et al., 2024).

3.3 Developers and Transmission Owners

Developers. We can identify the developer for only a subset of projects. From the second study onward, PJM lists each request’s interconnection customer, often the limited liability company holding the project rather than the developer that controls it. We match these names to developers using news articles, regulatory documents, and developer websites. We

Figure 1: Queue Size Over Time



The figure plots the size of the PJM interconnection queue from 2008 to 2020, based on projects that entered from 1997 to 2020. The solid red line is the annual average number of projects in the queue, counting each project from its entry until withdrawal or final study receipt (or, for entries before 2008, its operation date). The dotted blue line is the total number of new entries each year.

identify the developer for 43.3% of projects overall, 57.5% of those that reach the second study, and 80.7% of those that reach the third.

Although developers can control several projects, treating each project as an independent decision maker, as we do in the model, is a reasonable approximation. First, no developer is large enough to internalize much of the congestion its projects create: the 1,809 matched projects map to 389 unique developers, the largest of which (Invenergy) accounts for 4.3%, and the fifteen largest together account for less than 32% (Appendix Table E.3). Second, a developer has little reason to manage its projects as a portfolio. Projects are small relative to output and input markets, so one project has little direct effect on the profitability of another. A developer should therefore advance every individually profitable project. Financing or staffing constraints could force developers to choose among their projects, but they can and do sell projects rather than withdraw them.¹¹

¹¹For example, Savion developed the TWE Myrtle solar project in Virginia (interconnection request AD1-144) and sold it to Dominion Energy in 2019 (Savion, 2022), before the project reached commercial operation in June 2020. Geenex began developing the Sweetleaf solar project in North Carolina (interconnection request AD1-056) in 2017 and sold it to CleanChoice Energy in February 2026 while it is still active in the queue (CleanChoice Energy, 2026).

Transmission Owners. PJM has 22 transmission owners whose engineering input the studies rely on.¹² Appendix Table E.4 shows that the average TO has 190 projects and 18.5 GW of queued capacity in our sample.

Projects are individually small relative to their TO. The average project is about 0.5% of its TO by both count and capacity, and under 3% even at a 25th-percentile TO. Developer concentration is also low within TOs: the average TO contains projects from 31 distinct developers.

3.4 Waiting Times

The first study arrives quickly and with little variation, but the later studies arrive much later and with more variation. Projects receive the first study after a mean of 5 months, with a standard deviation of just 1.5. The second and third studies take far longer: against PJM’s official targets of 4 and 6 months, the second arrives after a mean of 11 months (a standard deviation of 8 months) and the third after 15 months (a standard deviation of 10 months).

While priority is by entry cohort, study arrival is stochastic. Cohorts that entered the queue earlier are more likely to receive a study earlier than projects in later cohorts that are waiting for the same study, but this is not guaranteed. For example, among projects entering in 2012, the average second study arrival date was 2013Q2 for projects entering in the first half of the year and 2014Q4 for projects entering in the second half. Thus, the earlier cohort received their second study substantially earlier on average. However, 3.5% of projects that entered in the first-half 2012 cohort had still not received their second study by 2014Q4, the average arrival date for the later cohort.¹³

Time spent in the queue increased over our sample, though not as dramatically as the number of projects. On average, projects that reach the terminal study receive it 27 months after entering the queue. This duration rose from 22 months for projects entering in 2008-2012 to 29 months for those entering in 2013-2017 (see Appendix Figure E.3 for wait time statistics by entry quarter).

¹²TOs can own incumbent generation and projects in the queue, but we find no evidence this ownership correlates with interconnection wait times or costs.

¹³One reason wait times are variable is that these studies sometimes need to be revised. For example, 22 percent of projects in our sample had their second study revised. We will model the arrival of the final version of each study, not the intervening versions which are not always posted by PJM.

3.5 Interconnection Costs

Interconnection costs are low for most projects but very high for some. We use cost per megawatt as our measure of cost.¹⁴ For the second study, 29 percent of projects have interconnection costs less than 0.01 million per megawatt. The 75th percentile of the cost distribution is 0.14 million per megawatt and the 90th percentile is 0.40 million per megawatt. For comparison, installation costs for wind and solar projects are roughly 1.5 million per megawatt.


Developers face substantial uncertainty about what their interconnection costs will be. First, although a developer may be able to use engineering models to identify whether an interconnection will overload the grid, it is hard to predict what remedies a transmission owner will require and the cost of these remedies, according to our interview with an industry expert. Second, the withdrawal of other projects sharing the cost of a network upgrade changes a project’s cost. After withdrawals, PJM may still require the upgrade, so the remaining projects’ share of its cost increases. If PJM no longer requires the upgrade, the cost for the remaining projects will decrease. Consistent with these sources of uncertainty, observable project characteristics explain less than half the variation in interconnection costs, and projects in the same geographic area can have very different costs (Appendix Figure E.2, Johnston, Liu and Yang (2026)).

Changes in costs across studies within a project are also hard to predict. Figure 2 shows that interconnection costs for the same project do not systematically decrease across studies.¹⁵ This pattern suggests that cost changes are not anticipated. If developers were able to predict how a project’s interconnection cost would evolve across studies, we would expect selection on this difference, i.e., projects that predicted their cost to fall would be more likely to wait, and the distribution would have more mass below zero.

¹⁴There do not appear to be economies of scale in interconnection costs for moderately sized projects. For projects from the 10th to 90th percentile in size, a 1 standard deviation increase in capacity is associated with 0.05 standard deviation decrease in the second-study interconnection cost.

¹⁵Costs on average increase from the first to the second study because the second study includes the contribution to shared network upgrade costs, while the first study does not. For projects that do not share costs, the distribution of this cost difference is symmetric around zero.

Figure 2: Within-Project Changes in Interconnection Cost Estimates across Studies




Panel (a) shows a histogram of the difference in the first and second study interconnection cost estimates for projects that received the second study. $N = 2,495$; 151 projects with cost differences outside the plotted range are excluded (126 above and 25 below), as are 4 without a first study. Panel (b) shows a histogram of the difference in the second and third study interconnection cost estimates for projects that received the third study. $N = 829$; 19 projects with cost differences outside the plotted range are excluded (10 above and 9 below), as are 51 without a second study. The sample is projects queuing from 2008–2020. Y-axis is the percent of observations in each bin. Costs are in millions of 2020 dollars per MW.

3.6 Withdrawals

Withdrawals concentrate sharply in the quarter immediately following each study. Figure 3 plots the cumulative share of projects that have withdrawn from the queue by quarter relative to the arrival of each study. Each panel is restricted to projects that required a further study to complete the interconnection process but withdrew before receiving the next study. Following the first study arrival, 59 percent withdraw within the first quarter and 84 percent within two years. The pattern after the second study is similar, with 55 percent withdrawing within the first quarter. Both curves rise steeply at first and then flatten. Because the deposits required to request the next study are small (Section 2.1), withdrawals concentrating at these points are more consistent with developers acting on the cost and engineering information each study reveals.

We next test whether high interconnection costs explain these withdrawals. Specifically, we regress an indicator for withdrawing from the queue on an indicator for having a high (above 0.1 million per megawatt) interconnection cost. In this regression, we define withdrawing as

Figure 3: Withdrawals Relative to Study Arrival



Share of projects that withdrew within t quarters of receiving their first study (left panel) or second study (right panel). For each curve, the sample includes all projects that entered the queue from 2008 to 2020, received the relevant study, required a further study, but withdrew before receiving it, $N = 1,183$ and $N = 666$ respectively. The zero point on each x-axis marks the quarter in which the corresponding study arrived. The figure displays only the first eight quarters for visualization purposes. In the left panel, 59.3% withdrew within the first quarter and 13.9% withdrew more than eight quarters after receiving the first study; in the right panel, 55.3% withdrew within the first quarter and 21.3% withdrew more than eight quarters after receiving the study.

leaving the queue before the next study arrives, or before beginning operation for projects that have received their final study.

Across all three studies, projects with a high interconnection cost estimate are more likely to withdraw from the queue (Table 2). The OLS estimates imply that a high interconnection cost is associated with an increase in the probability of withdrawal of 12 percentage points after the first study (39% of the mean withdrawal rate), 21 percentage points after the second (45%), and 16 percentage points after the third (31%).¹⁶

The OLS estimates may not isolate the causal effect of costs. Interconnection costs tend to be higher in densely developed areas (Appendix Figure E.2), where projects may also face more effective local opposition, a cancellation risk unrelated to the assessed cost. If so, the OLS estimates would be biased upward. Conversely, if developers plan projects in expensive

¹⁶These results are robust to including additional controls for permitting difficulties—distance from project to grid connection point, an indicator for local solar and wind ordinances—that might be correlated with high interconnection costs.

Table 2: Effect of Interconnection Costs on the Probability of Withdrawing from the Queue

	1 st Study	2 nd Study		3 rd Study	
	OLS	OLS	IV	OLS	IV
High cost (>0.1M/MW)	0.116*** (0.019)	0.212*** (0.031)	0.379*** (0.052)	0.161** (0.077)	0.364** (0.154)
Mean of dependent var.	0.30	0.47	0.47	0.52	0.52
Share with high cost	0.26	0.33	0.33	0.31	0.31
First-stage F			153.5		23.2
Observations	3,983	1,974	1,974	471	471

The sample is projects that entered the queue in 2008–2020; projects still active at the end of the sample are excluded, and each column is limited to projects receiving the relevant study. The third-study columns further exclude the 46 projects without a second study, so the OLS and IV samples coincide. The dependent variables indicate withdrawal before the next study or, after the final study, before beginning operation. High cost indicates that the study’s interconnection cost estimate exceeds \$0.1 million per MW. All specifications control for project size (three bins), fuel type, whether the project increases capacity at an existing plant, transmission owner FE, and queue-entry-year FE; the second- and third-study columns also control for the prior study’s high-cost indicator. The IV instruments the high-cost indicator with the within-project change in the cost estimate across studies, in millions of dollars per MW, winsorized at the 1st and 99th percentiles. First-stage F is the Kleibergen-Paap rk Wald statistic. Ninety-five percent Anderson–Rubin confidence sets, robust to weak instruments, are [0.28, 0.48] and [0.08, 0.67] for the second- and third-study IV coefficients. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

locations only when those projects are otherwise especially profitable, the resulting positive selection would attenuate the OLS estimates. To address both concerns, we instrument for current interconnection costs with the change in costs across studies, which is difficult to predict (Section 3.5) and is the same variation that identifies the payoffs and waiting costs in the structural model (Section 6.2).

The IV estimates are larger: high costs increase the probability of withdrawal by 38 percentage points after the second study and 36 after the third. Both are statistically significant, and the weak-instrument-robust confidence sets in Table 2 exclude zero. The gap between IV and OLS is consistent with positive selection attenuating the OLS estimates. Taken together, the estimates imply that high interconnection costs substantially increase the probability of withdrawal.

Although most withdrawals occur immediately after a new study, withdrawals also occur

at a low and roughly constant rate between studies. As discussed in Section 2.3, projects face exit risks from outside the queue. These risks can interact with queue delays: each additional quarter spent waiting for the next study exposes a project to another period in which risks such as local opposition may materialize.

4 Externalities in the Queue

Projects in the queue interact with one another through three channels: PJM’s finite capacity to process engineering studies (congestion externality), the physical transmission infrastructure shared across interconnections (cost externality), and the wholesale electricity market (output externality). We find that the congestion externality is quantitatively important, while the geographic cost and local output externalities are not.

4.1 Congestion Externality

We first investigate how queue length affects how long projects wait for studies. The congestion externality is mechanical under strict first-come, first-served queuing with a fixed processing rate. PJM’s study arrival is probabilistic rather than strictly ordered, so the strength with which arrival probability responds to queue position is an empirical question. As noted in Section 3.4, the first study arrives soon after entry, but the arrival times of later studies are more variable.

We estimate a discrete-time hazard model, where the conditional probability that the second or third study arrives in a quarter is given by a probit model. We measure queue position using the natural log of the number of higher-priority projects (those with the same or earlier entry dates), both PJM-wide and within the project’s TO, and estimate how each measure affects the conditional arrival probability. Throughout, we condition flexibly on time already spent in the queue, thus comparing projects with the same waiting time. These controls separate the queue-position effect from pure duration dependence in study processing, i.e., a study becoming more or less likely to arrive simply because a project has already waited longer. We allow the effect of higher-priority projects to vary with whether a project is waiting for the second or third study.

Identification. The queue-position coefficients identify the congestion externality because PJM, not the project, determines study arrival. Selection affects which projects occupy a given queue position: those with high profits or low waiting costs (reflecting unobservables such as power purchase agreements or land costs) are more likely to survive to the front. Yet, as the PJM manual and industry interviews confirm, study priority depends solely on queue position rather than on the profitability driving this selection, and projects cannot act to expedite their studies. The selection that places high-profit projects at the front therefore does not bias the position effect. The remaining threat is that project characteristics affect both how long PJM takes to complete a study and queue position. We control directly for the main characteristics affecting how long a study takes: project size, fuel type, and the prior study’s information (cost estimate, engineering tests, and whether it is assigned a network allocation). Together, these arguments support the identifying assumption that the number of higher-priority projects is independent of the error term in the probit model.

Results. We find a significant congestion externality, which is largest for the third study. Table 3 reports the implied average marginal effects of higher priority projects on own study arrival. Our preferred specification, column (3), implies that a 10% increase in higher-priority projects reduces the second-study arrival probability by 2.9% and the third-study arrival probability by 12.5%.¹⁷

Withdrawals by nearby projects can trigger study revisions, slowing delivery of the final version and thus imposing a further congestion externality. Column (4) adds a control for recent number of withdrawals in the same TO along with the total number in the queue in the same TO, so that the withdrawal coefficient reflects withdrawals conditional on queue size. The estimated coefficient is close to zero and statistically insignificant. Measuring withdrawals and queue size in MW rather than project counts yields similar results (Appendix Table E.5, column (5)).

These estimates are robust to flexible time controls, addressing the concern that PJM’s study process may have evolved over this period. PJM did not make major changes to its process, but TOs may have learned to process studies more efficiently or hired more engineers in response to rising demand. To account for these effects, we re-estimate our preferred

¹⁷We assume the PJM-wide and TO queue sizes increase by the same percentage, so the effect of a 10% increase is the total queue-position effect in Table 3 multiplied by 0.10, divided by the study-specific mean arrival rate reported in the table note.

specification (column (4) of Table 3) controlling for a linear and a quadratic time trend, as well as their interactions with an indicator for renewable generation to allow for technology-specific learning. We also estimate a model with year-by-quarter fixed effects, a specification that highlights how the effect of queue position can be identified solely from cross-sectional differences in arrival probability for projects at different queue positions. Finally, we estimate a model with transmission owner-by-year fixed effects, which absorb shocks common to a transmission owner in a given year that could drive both queue length and study delays. Appendix Table E.6 shows that we consistently find a negative and statistically significant effect of the number of higher priority projects on study arrival for both the second and third studies. The estimates imply that a 10% increase in higher priority projects reduces study arrival probability by 1.9 to 2.9% for the second study and 8.6 to 12.5% for the third study.

A second concern is that the effect is driven by the end of the sample, when the queue was most congested. Restricting to projects that queued before 2019, which excludes 40% of the projects in our sample, we still find statistically significant and negative effects of queue position on study arrival: 3.6% for the second study, 8.3% for the third (Appendix Table E.6, column (6)).

Finally, in support of our claim that PJM’s study delivery does not respond to expected project profitability, we re-estimate the arrival model including two proxies for profitability: the average local wholesale electricity price and an indicator for binding local wind or solar ordinances (Appendix Table E.5). Both effects are small. The price effect is significant at the 5 percent level only without transmission owner and year fixed effects, and even this largest estimate implies that a one standard deviation increase in the wholesale price, \$18 per MWh, raises quarterly study arrival by 0.76 percentage points, a sixth of the 4.7 percentage point effect of a one standard deviation change in queue position. The ordinance effect is indistinguishable from zero and changes sign across specifications. The queue-position effects remain similar.

4.2 Cost Externality

Projects connecting in the same location can affect one another’s interconnection costs through the network infrastructure they share. Those queued ahead of a given project may consume available local capacity, leaving the project responsible for an upgrade. When several projects

trigger a common upgrade, PJM assigns them to share its cost; if one member of a cost-share group withdraws, the others' costs can rise or fall. A completed project that paid for a costly network upgrade can also lower later projects' costs, because that upgrade may leave spare capacity behind.

4.2.1 Spillovers from Other Projects in the Queue

Higher-Priority Local Capacity. PJM models higher-priority projects as in service when conducting engineering studies, so these projects can consume available local transmission capacity. More higher-priority local capacity could therefore increase the probability a project triggers costly network upgrades, raising its interconnection cost.

We regress the indicator for a high interconnection cost (above \$0.1 million per megawatt) on the amount of higher-priority capacity in a project's local area when its study is issued. We repeat the analysis for the network-upgrade component of cost. We define the local area four ways: within 10 km, within 20 km, within the same transmission owner, and within the same hierarchical cluster, where we group substations into 50 regions by location and similarity in locational marginal prices.¹⁸ We include geographic-region and queue-year fixed effects and control for recent transmission investment by the grid operator. The resulting estimates are nonetheless suggestive rather than causal, since the entry and withdrawal decisions that determine higher-priority capacity are themselves affected by the state of the network.

We find no evidence that more higher-priority capacity raises a project's cost under any of the four definitions, whether we measure total cost or its network-upgrade component alone. The point estimates in Appendix Table E.7 are uniformly small: even the largest in absolute value implies that a doubling of higher-priority capacity changes the probability of a high cost by about 2.4 percentage points, or roughly 6% of its sample mean of 0.37. This result is robust to alternative definitions of a high interconnection cost and to using the total capacity of active waiting projects rather than higher-priority capacity.

Withdrawals Within Cost-Share Groups. When multiple projects trigger the same network upgrade, PJM assigns them to a cost-share group and allocates the cost in proportion to each project's impact. More than half of projects (59%) share costs as of the second study. We observe complete group membership for only 29% of cost-sharing projects (153 groups),

¹⁸For details, see Johnston, Liu and Yang (2026).

and conditional on belonging to a fully observed group, 37% of projects see at least one member withdraw before the focal project reaches a terminal study.

We test whether these withdrawals systematically affect costs for remaining projects. Because both the total group cost and its allocation can change when a project withdraws, the effect on remaining group members is theoretically ambiguous: if PJM still requires the upgrade, the remaining members bear a larger share of its cost; if the withdrawal removes the need for it, their cost falls. We also test whether member withdrawals predict a focal project’s own withdrawal, since selective exit by projects facing high costs could mask a cost effect.

We find the estimated effects of member withdrawals on cost are small and statistically insignificant (Table 4, columns (1) and (2)). A one-standard-deviation increase in the withdrawn share (0.22) implies a \$0.02–\$0.03 million per megawatt reduction in cost between studies, about a tenth of a standard deviation. The corresponding 95% confidence intervals rule out cost increases larger than 6% of a standard deviation of cost changes or decreases larger than one-quarter of a standard deviation. Member withdrawals also have no significant effect on a project’s own decision to withdraw (column (3)), providing no evidence that the null cost estimates are driven by selective exit.

4.2.2 Spillovers from Completed Interconnections

A completed interconnection can lower a later project’s cost by leaving spare transmission capacity, but cost-sharing and the project-specific nature of upgrades may limit this spillover.¹⁹ We test for the spillover by regressing an indicator for a high cost estimate (above \$0.1 million per megawatt) on whether the most recently completed interconnection in the same area was costly, defined as a total interconnection cost above \$1 million; 22% of completed interconnections meet this threshold. We exclude prior interconnections that share costs with the focal project, so the regressor captures the effects from earlier completion instead of the effects from being in the same cost-share group.

The estimated effects on a project’s subsequent total cost are statistically insignificant across the 10 km, 20 km, and transmission-owner definitions (Table 5, Panel A). The only exception is the hierarchical-cluster specification, where a prior costly interconnection is asso-

¹⁹In FERC Order 2023, the Commission states that “absent cost sharing provisions among clusters, interconnection customers may significantly benefit from earlier-in-time network upgrades but not share in the cost of those network upgrades in a manner that is roughly commensurate with benefits” (p. 320), suggesting that this externality is more relevant in systems without cost sharing (FERC, 2023).

ciated with a 3.1 percentage point reduction in the probability of a high total cost estimate, or 8% of the sample mean, and the confidence interval barely excludes zero. This effect is not systematic across alternative geographic definitions or cost thresholds (Appendix Table E.8). Overall, the total-cost results provide at most limited evidence of a cost-reducing spillover. For the network-upgrade component of interconnection cost, where these spillovers are most likely to occur, we find insignificant effects across all four geographic definitions (Table 5, Panel B). We likewise find no stable evidence of network-upgrade-cost spillovers across alternative geographic definitions or cost thresholds (Table E.9).

Discussion. Across these three channels, the evidence points to limited cost externalities. The null effect of higher-priority capacity is likely explained by PJM’s cost-sharing rules, which spread the cost of an upgrade across the projects that jointly trigger it. Consistent with this explanation, total active capacity in the local queue predicts network allocation, and network allocation predicts lower subsequent costs (Appendix Tables E.10 and E.11). The largely null effect of completed interconnections likely reflects the project-specific nature of upgrades: an upgrade is prescribed to resolve the violations caused by the projects that trigger it, so the spare capacity it leaves behind is incidental. The within-group null effect follows a different logic: a member’s withdrawal raises remaining members’ shares when the upgrade is still required but eliminates the cost when it is not, and the slightly negative point estimates suggest these forces roughly offset.

4.3 Local Output Externality

By lowering local wholesale prices, a completed interconnection could reduce the expected revenues, and hence the entry and completion incentives, of nearby projects. This channel is most likely to be present at locations with binding transmission constraints, where a new project’s output cannot flow freely across the grid and instead depresses local prices.

We test for the effect of a new project beginning operation on local prices. We first construct a monthly panel with an observation for each substation, which is a potential grid connection point. Our measure of monthly wholesale prices is the average peak-hour locational marginal price at that substation. We use the event-study difference-in-differences estimator from Callaway and Sant’Anna (2021) to estimate the effect of a new project beginning operation on these prices. The comparison group is substation-months where no new projects

above the size threshold begin operation; we use thresholds of 20 and 100 MW.


We find little effect of a new project beginning operation on local prices. Figure 4 shows that the estimated price decrease is small and statistically indistinguishable from zero in each post-period month. The point estimates are consistently negative for the first seven to eight months of operation before turning positive, suggesting a small and transitory negative effect of roughly \$2 per MWh.

Local Price Effect at Locations of Transmission Constraints. To directly address the concern that the output externality may operate only at locations where projects face binding transmission constraints, we re-estimate the event study restricting the treatment group to substations where the new project was responsible for paying transmission network upgrade costs. These are precisely the projects connecting to an already-congested local grid, so any local effects should be most detectable in this subsample. We continue to find a null effect (Appendix Figure E.4).

Discussion. This null result is consistent with the institutional details of the interconnection process. PJM’s deliverability requirement obligates most new projects to upgrade the transmission infrastructure to the point that they can deliver their power to demand centers even during peak demand hours. These upgrades relieve local congestion, limiting the effect of added generation on local prices. A completed interconnection’s effect on the expected profits of future projects in the same location is therefore likely small, even where the grid was constrained prior to the upgrade.

As more projects complete and add to supply, PJM wholesale prices should fall. We incorporate this aggregate effect in the model by constructing project revenues from the wholesale-market equilibrium.

Figure 4: Effects of Interconnections on Local Wholesale Prices



Panel (a) plots event-study estimates of the effect of a new project with capacity of at least 20 MW beginning operation at a substation on local wholesale electricity prices. Panel (b) uses a 100 MW threshold. Treated observations are substation-months in which a project above the corresponding capacity threshold begins operation; control observations are substation-months without such entry. The dependent variable is the monthly average peak-hour locational marginal price at the substation level (mean = \$34.89/MWh), constructed from PJM locational marginal price data over 2008-2020. Estimates are obtained using the Callaway and Sant’Anna (2021) difference-in-differences estimator with doubly robust inverse probability weighting. Standard errors are clustered at the substation level.

Table 3: Effect of Queue Position on Study Arrival

	(1) Probit	(2) Probit	(3) Probit	(4) Probit	(5) Probit
ln(# Higher priority)	-0.043*** (0.007)	-0.031*** (0.006)	-0.031*** (0.006)	-0.022*** (0.006)	-0.020 (0.012)
ln(# Higher priority) X Waiting for third study	-0.026*** (0.007)	-0.017** (0.008)	-0.016** (0.008)	-0.014* (0.007)	-0.011 (0.007)
ln(# Higher priority in TO)		-0.016*** (0.004)	-0.016*** (0.003)	-0.043*** (0.007)	-0.020*** (0.007)
ln(# Higher priority in TO) X Waiting for third study		-0.013* (0.007)	-0.012* (0.007)	-0.009 (0.007)	-0.011* (0.007)
Waiting for third study	-0.209*** (0.022)	-0.218*** (0.021)	-0.216*** (0.023)	-0.231*** (0.024)	-0.243*** (0.023)
ln(# of withdrawals in queue within TO)				-0.000 (0.005)	
ln(# Projects in queue within TO)				0.030*** (0.007)	
<i>Total queue-position effect (linear combinations)</i>					
At second study	-0.043*** (0.007)	-0.046*** (0.006)	-0.047*** (0.006)	-0.064*** (0.008)	-0.040*** (0.012)
At third study	-0.069*** (0.006)	-0.076*** (0.006)	-0.075*** (0.006)	-0.087*** (0.007)	-0.062*** (0.011)
Prior Study Controls			X	X	X
Transmission Owner FE					X
Year FE					X
Observations	28,917	28,917	28,917	28,917	28,917
Pseudo R ²	0.166	0.173	0.176	0.179	0.195

The sample includes projects queued from 2008-2020, observed over time at the quarterly level as they wait in the queue for study arrival. Marginal effects from probit regressions are reported. Standard errors in parentheses, clustered by transmission owner - entry cohort (416 clusters). The dependent variable is an indicator for study arrival (overall mean = 0.11; mean for second study is 0.16 and for third study is 0.06). All specifications control for project size (three bins), fuel type (wind, solar, storage), if the project is a capacity increase at an existing plant, and time spent waiting in the queue (a separate indicator for each quarter since queue entry as well as their interactions with an indicator for whether the project is awaiting its third study). Specifications (3)-(5) additionally include controls from the immediately prior study: the project's own cost estimate (e.g., the first study costs for projects awaiting the second study; three bins), indicators for having ever received each of the three engineering tests and indicator for whether the project receives an explicit network upgrade cost. Specification (4) further controls for the log number of withdrawn projects waiting for the same or a later study within the same transmission owner over the past two quarters, and the log number of projects in the queue within the same transmission owner. The total-effect rows report linear combinations of the marginal queue-position effects. *At second study* sums the system-wide and transmission-owner queue effects, while *At third study* additionally includes their interactions with the indicator for awaiting the third study. *** p<0.01, ** p<0.05, * p<0.10.

Table 4: Effects of Cost-Share Group Member Withdrawal on Own Cost and Withdrawal

	Change in cost per MW		Own withdrawal
	Total (1)	Network upgrade (2)	(3)
Share of group MW withdrawn	-0.132 (0.093)	-0.100 (0.089)	-0.009 (0.010)
Transmission Owner FE	X	X	X
Queue Year FE	X	X	X
Mean of dep. var [SD]	-0.017 [0.289]	-0.042 [0.274]	0.023 [0.149]
Mean of indep. var. [SD]	0.077 [0.224]	0.077 [0.224]	0.068 [0.197]
Observations	206	206	8,928

The sample is projects that entered the queue in 2008–2020 that are in a cost-share group whose full membership we observe. Standard errors are in parentheses and clustered by cost-share group. Columns (1)–(2) restrict the sample to projects that have received the third study, since projects typically begin to observe their cost-share group in the second study; the dependent variable is the change in cost per MW from the second to the third study. Column (1) uses total interconnection cost per MW, and column (2) uses network upgrade cost per MW. Column (3) reports linear-probability estimates from a project-quarter discrete-time hazard panel; the dependent variable is an indicator for project withdrawal in the current quarter. The independent variable, Share of group MW withdrawn, is the share of cost-share group peers withdrawn during the relevant exposure window (the inter-study window for columns (1)–(2); since the focal project’s most recent study for column (3)). All specifications control for project size (three bins), fuel type (wind, solar, storage), whether the project is a capacity increase at an existing plant, and the log of total cost-share group size in MW. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 5: Effect of Prior Costly Interconnection on Probability of High Cost

	10km (1)	20km (2)	Transmission owner (3)	Hierarchical cluster (4)
Panel A: High Total Cost Estimate (mean = 0.37)				
Costly prior interconnection	0.056 [-0.045, 0.156]	-0.005 [-0.091, 0.082]	-0.012 [-0.078, 0.056]	-0.031** [-0.063, -0.000]
Geographic Region FE	X	X	X	X
Queue Year FE	X	X	X	X
Observations	2,328	2,328	2,328	2,328
R^2	0.525	0.523	0.239	0.242
Panel B: High Network Upgrade Cost Estimate (mean = 0.13)				
Costly prior interconnection	0.111 [-0.217, 0.475]	0.060 [-0.129, 0.270]	-0.044 [-0.106, 0.090]	0.037 [-0.032, 0.110]
Geographic Region FE	X	X	X	X
Queue Year FE	X	X	X	X
Observations	2,328	2,328	2,328	2,328
R^2	0.475	0.480	0.178	0.204

Projects queuing from 2009–2020. Wild cluster bootstrap (9,999 replications) clustered at the substation level in the 10 km and 20 km columns, the transmission owner level in the transmission owner column, and the hierarchical cluster level in the hierarchical cluster column; 95% confidence intervals in brackets. In Panel A, the dependent variable is an indicator for whether the total cost estimate in the second or third study is high, defined as exceeding \$0.1 million per MW. In Panel B, the dependent variable is an indicator for whether the network upgrade cost estimate in the second or third study is high, defined as exceeding \$0.1 million per MW. The regressor of interest, Costly prior interconnection, is defined differently across panels. In Panel A, it is an indicator for a prior completed interconnection whose total interconnection cost exceeds \$1 million within the local geographic region, varying across columns (10km radius, 20km radius, transmission owner, or hierarchical cluster). In Panel B, it is an indicator for a prior completed interconnection whose network upgrade cost exceeds \$1 million within the local geographic region. All specifications control for project size (three bins), fuel type (wind, solar, storage), whether the project is a capacity increase at an existing plant; whether PJM Regional Transmission Expansion Plan investment was above the 75th percentile of the RTEP distribution within the same transmission owner, measured over a four-year window spanning three years before and one year after the study issue date; the log of total entrant capacity (MW) within the previous three years in the local geographic region; and the log capacity (MW) of higher-queued projects in the local geographic region. The geographic region fixed effects are substation fixed effects in the 10 km and 20 km columns, transmission owner fixed effects in the transmission owner column, and hierarchical cluster fixed effects in the hierarchical cluster column. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

5 Model

Each project solves a finite-horizon optimal stopping problem. Time is discrete, indexed by $t = 1, 2, \dots$ in quarters. Potential projects arrive at the beginning of a period and decide whether to pay an entry cost to join the queue or leave immediately. Projects that join receive up to three studies. Between studies, projects incur waiting costs and can withdraw at any time. Each project also faces a per-period probability of exogenous exit due to events outside the interconnection process. After the final study, a project decides whether to complete interconnection or withdraw.

Projects do not directly interact with each other. Each project responds to a vector of variables called the *queue status*, which determines its study-arrival probabilities, distributions of interconnection costs, and operating revenues. The transitions of the queue status, aggregated from strategies of atomistic projects, are deterministic. The queue status has both PJM-wide and TO-specific components, since PJM conducts the studies within each TO's territory in conjunction with that TO.

Below, we first specify the project decision problem. Then we provide an equilibrium concept where project strategies are functions of their beliefs about the equilibrium path of the queue statuses, and their strategies generate the path consistent with their beliefs.

5.1 Decision Problem

We describe a project's decision problem recursively, beginning with the completion decision and working backwards. PJM may require fewer than three studies if the project needs no significant network upgrades and does not share costs. A project's last required study is its *terminal study*, after which it decides whether to complete or withdraw. The first or second study can be terminal, while the third study always is. After a non-terminal first study, the project enters *Stage 1*, waiting for its second study; after a non-terminal second study, it advances to *Stage 2*, waiting for its third study. A project waits at most $T = 32$ quarters after entry; one that reaches T without a terminal study exits with zero payoff.

Terminal Study Decision. Upon receiving its terminal study in period t , project i chooses to *complete* or *withdraw* at zero payoff. The net payoff from completing, on a per megawatt

basis, is

$$\begin{aligned}\pi_{it} &= r_{it} - \omega_{it} - c_{it} + \xi_i \\ &= \tilde{r}_{it} + \chi_t(x_i)\theta - \omega_{it} - c_{it} + \xi_i\end{aligned}\tag{1}$$

where r_{it} is the present value of expected operating profit, ω_{it} is the expected cost of construction, and c_{it} is the interconnection cost in the terminal study. We decompose r_{it} into $\tilde{r}_{it}$, the present value of revenue from power production, and $\chi_t(x_i)\theta$, which captures other contributions to profit through a vector $\chi_t(x_i)$ that varies with project characteristics x_i and completion time t , e.g., operating costs and technology-specific subsidies. This decomposition lets revenues from power production $\tilde{r}_{it}$ respond endogenously to capacity changes in our counterfactual simulations. We describe $\tilde{r}_{it}$ and ω_{it} in more detail in Section 6.2.

The term ξ_i captures payoff information unobserved to the researcher that has accumulated by the terminal study. The component ξ_i^{S1} is revealed at the first study. If that study is terminal, $\xi_i = \xi_i^{S1}$; otherwise, ξ_i^{S2} is revealed when the second study is received, and $\xi_i = \xi_i^{S1} + \xi_i^{S2}$. Thus the composition of ξ_i depends on the stage reached, even though the accumulated information itself does not carry a time subscript. As described in Section 2.2, studies can signal additional costs during the development process. More broadly, ξ^{S1} and ξ^{S2} also reflect information a project learns about conditions outside the interconnection process while waiting, such as the contract terms it can expect from power buyers or the cost of renegotiating its land lease.

Project i also learns a final payoff shock ξ_i^f at the terminal study and completes if $\pi_{it} + \xi_i^f > 0$. Conditional on terminal-study receipt in period t , we assume ξ_i^f is normally distributed with mean 0 and standard deviation $\sigma \cdot (t - e_i)^\kappa$. We also assume ξ_i^f is conditionally independent of ξ_i given project characteristics x_i and the interconnection cost c_{it} . Here, e_i is the quarter in which the potential project arrives and makes its entry decision; a project that enters joins the queue in $e_i + 1$. The dependence on $t - e_i$ accommodates larger payoff dispersion among projects that have waited longer. The distribution of ξ_i^f thus changes over time, but for notational simplicity, we suppress the time subscript.

Integrating over ξ_i^f , the expected value of receiving the terminal study given the interconnection cost c is

$$\bar{V}_{it}(c, \xi_i) = E_{\xi_i^f} \left[\max\{\pi_{it} + \xi_i^f, 0\} \right].\tag{2}$$

Because x_i and e_i do not change over time, we absorb them into the i subscript and condition on them throughout.

Stage-2 Continuation. After receiving a non-terminal second study, project i decides each period whether to wait or withdraw. The project's state consists of the current time t , the second study cost c , study information z , and the accumulated unobserved heterogeneity $\xi = \xi^{S1} + \xi^{S2}$. The variables (c, z, ξ) are specific to each project i , but for notational simplicity, we suppress their i subscript. We also assume that the costs have finite support. A waiting decision is made at the start of period t : $q_i^{(\cdot)}$ and $Q_i^{(\cdot)}$ describe study receipt during period t , while the no-arrival state is carried into $t + 1$. The continuation value satisfies

$$W_{it}^{S2}(c, z, \xi) = \max \left\{ -\gamma^{S2} + \beta \delta_i^{S2} \left[\sum_{c'} q_i^{S2}(c'|t, c, z) \bar{V}_{it}(c', \xi) + (1 - Q_i^{S2}(t, c, z)) W_{i,t+1}^{S2}(c, z, \xi) \right], 0 \right\}, \quad (3)$$

where γ^{S2} is the per-megawatt waiting cost,²⁰ and $\beta = 0.987$ is the quarterly discount factor, implying an annual discount factor of 0.95. The term $1 - \delta_i^{S2} \equiv 1 - \delta^{S2}(x_i)$ is the per-period probability of exogenous exit, capturing events outside the queue process such as failed permitting or expired land leases and varying with project characteristics x_i . The probability that the terminal study arrives with cost c' during period t is $q_i^{S2}(c'|t, c, z)$, and $Q_i^{S2} = \sum_{c'} q_i^{S2}(\cdot)$ is the total probability of study arrival. The study information z consists of two indicators: whether engineering tests have been performed, and whether the project is assigned a network allocation, as motivated in Sections 2.2, 4.1 and 4.2. They are useful signals of how fast future studies arrive and how costly the project may be.

Information Assumptions. A project has perfect foresight over the probabilities $q_i^{(\cdot)}$ and $Q_i^{(\cdot)}$ governing study arrivals and costs in each period t . In the equilibrium defined in Section 5.3, the deterministic transitions of the queue status imply these probabilities. At the start of period t , however, projects do not know whether a new study will arrive that period or, if one does, the realization of its cost, c' . We maintain these assumptions below.

²⁰We add the study fee to request the next study to the waiting cost of the period after the study is received, with fees based on the PJM manual (PJM, 2023b).

Stage-1 Continuation. After receiving a non-terminal first study, project i decides each period to wait or withdraw. The state consists of time t , the first study cost c , study information z , and the unobserved heterogeneity ξ^{S1} from stage 1. If the second study arrives during period t , a terminal study leads to the completion decision in t , while a non-terminal study moves the project into Stage 2 in $t + 1$. The continuation value satisfies

$$W_{it}^{S1}(c, z, \xi^{S1}) = \max \left\{ -\gamma^{S1} + \beta \delta_i^{S1} \left[\sum_{c'} q_i^{S1,f}(c'|t, c, z) E_{\xi^{S2}}[\bar{V}_{it}(c', \xi^{S1} + \xi^{S2}) | t] \right. \right. \\ \left. \left. + \sum_{c', z'} q_i^{S1}(c', z'|t, c, z) E_{\xi^{S2}}[W_{i,t+1}^{S2}(c', z', \xi^{S1} + \xi^{S2}) | t] \right. \right. \\ \left. \left. + (1 - Q_i^{S1}(t, c, z)) W_{i,t+1}^{S1}(c, z, \xi^{S1}) \right], 0 \right\}, \quad (4)$$

where γ^{S1} is the per-megawatt waiting cost and $1 - \delta_i^{S1}$ is the per-period probability of exogenous exit, $q_i^{S1,f}(c'|t, c, z)$ is the probability that a terminal second study arrives with cost c' during period t , $q_i^{S1}(c', z'|t, c, z)$ is the probability that a non-terminal second study arrives with updated state (c', z') during period t , and $Q_i^{S1} = \sum_{c'} q_i^{S1,f}(\cdot) + \sum_{c', z'} q_i^{S1}(\cdot)$ is the total probability of study arrival during period t . We assume that, like ξ_i^f , ξ^{S2} is normally distributed with mean 0 and standard deviation $\sigma^{S2} \cdot (t - e_i)^{\kappa^{S2}}$. We also assume that ξ^{S2} is conditionally independent of ξ^{S1} given project characteristics x_i , the current interconnection cost c , and study information z . We again suppress the time subscript on ξ^{S2} for notational simplicity.

Entry. A potential project arrives in period t and decides whether to enter or leave immediately.²¹ We set $e_i = t$. If it chooses to enter, it joins the queue and receives its first study in $t + 1$. The project's expected value of entry is

$$EV_i = E_{\xi^{S1}, c, z} \left[q_i^f \bar{V}_{i,t+1}(c, \xi^{S1}) + (1 - q_i^f) W_{i,t+1}^{S1}(c, z, \xi^{S1}) \right], \quad (5)$$

where q_i^f is the probability that the first study is terminal, and the integration is over the joint distribution of (ξ^{S1}, c, z) realized at the first study in $t + 1$. The entry value is evaluated

²¹A potential entrant thus receives a one-time entry opportunity. This abstracts from the option to delay entry. Delay is undesirable for two reasons: waiting consumes the option period on the leases used for site control, and a later request receives lower priority in a lengthening queue, since PJM assigns priority by request date.

at the moment of first-study receipt, abstracting from variation in waiting time up to that point, which is small relative to the variation in later study arrivals. We assume that ξ^{S1} is normally distributed with mean 0 and standard deviation σ^{S1} . We also assume that ξ^{S1} is conditionally independent of c and z given project characteristics x_i . Project i pays C_i and enters the queue if $EV_i \geq C_i$, where the entry cost C_i follows a distribution that depends on a vector of project characteristics x_i^{ent} . We specify the distribution in Section 6.3.

5.2 State Transitions and Initial State Distribution

We parameterize the state-transition process that governs how PJM manages the queue using probit and ordered-probit models. These models specify the probability of new study arrival and, conditional on the arrival, the distribution of the final interconnection cost c if the new study is terminal, or the distribution of the new study information (c, z) if the study is non-terminal. We also specify the distribution of a project’s state at the first study. We provide details in Appendix D.

5.3 Equilibrium

We assume that in each period t , an exogenous mass of potential projects $N_t(x)$ decides whether to enter the queue. A portion of the mass enters based on the entry decisions consistent with (5), and this mass then decreases deterministically because of endogenous withdrawals, exogenous exits, and completions.

Queue Status. Each project’s state-transition probabilities depend on a queue-status vector $\mathcal{K}_t(e, x)$, where e is the project’s arrival quarter and x collects its time-invariant characteristics, including fuel type, capacity, and TO location. In the empirical model, $\mathcal{K}_t(e, x)$ consists of four components that vary with t : (1) the number of queued projects with weakly higher priority, i.e., those with arrival quarter $\tilde{e} \leq e$, which determines the project’s position in the PJM-wide queue in period t ; (2) the number of such higher-priority projects located in the same TO; (3) the total capacity of queued projects in the same TO; and (4) the cumulative completed capacity of each fuel type up to period t . The first two components affect study-arrival probabilities, the third affects the distribution of interconnection costs through the likelihood of being assigned a network allocation (Section 4.2), and the fourth

affects the present value of revenue through wholesale electricity prices. Projects take the path $\{\mathcal{K}_t(e, x)\}_t$ as given when solving their decision problems.

Equilibrium Definition. We define a queuing equilibrium as a path of queue statuses $\{\mathcal{K}_t(e, x)\}_t$ for every (e, x) , together with entry probabilities and project strategies, satisfying three conditions: (i) project entry, waiting, and completion strategies are optimal given beliefs, as defined in Section 5.1; (ii) project beliefs about study-arrival probabilities and the joint distribution of costs and study information are consistent with those implied by the queue statuses and the prior study’s information (Section 5.2); and (iii) the path of queue statuses is generated by the entry, withdrawal, and completion decisions of projects.

Because we model projects as a continuum, the entry, withdrawal, and completion probabilities, together with exogenous exit, move mass deterministically. The resulting masses in turn imply the path of $\mathcal{K}_t$ that projects take as given. The equilibrium is therefore a fixed point where the queue status is consistent with the individual decisions it induces. In Appendix A, we formalize this definition, characterize the optimal strategies as cutoff rules, and prove equilibrium existence.

Under conditions we provide in Appendix A, endogenous withdrawals occur only immediately after projects receive a study. In the other periods between studies, projects do not endogenously withdraw, and only exogenous exits occur with probability $1 - \delta^{(\cdot)}$. This behavior is consistent with Section 3.6, where we document that withdrawals concentrate in a short period after a study arrives. This characterization separates the endogenous withdrawals from exogenous exits and simplifies the estimation of the model.

This equilibrium definition emphasizes the congestion externality, which we find to be quantitatively important in Section 4. A project’s strategy is, in equilibrium, a function of its belief about its queue position, which reflects queue congestion and affects study arrival, and its belief about whether it is assigned a network allocation, which affects the next study’s cost. Given the empirically small effect of withdrawals by other projects within a cost-share group and the weak correlation between study costs and local queue congestion, we do not model projects as tracking the decisions of individual projects.

Relationship to Similar Equilibrium Concepts. This equilibrium concept is a large-market rational-expectations equilibrium for a non-stationary environment. Section 3.3 shows

that, even at the TO level, an individual project is small relative to the total queued capacity in its TO. The equilibrium concept is related to dynamic industry and search models in which agents know equilibrium aggregate objects but still face idiosyncratic uncertainty. For example, firms in Hopenhayn (1992) know the stationary industry environment but not their future productivity shocks, and workers and firms in search models know market tightness and matching rates but not the date of a match. A Markov-perfect equilibrium (Ericson and Pakes, 1995) would require each project to track the individual states of all other projects. Instead, as in oblivious equilibrium (Weintraub, Benkard and Van Roy, 2008) and nonstationary oblivious equilibrium (Weintraub et al., 2010), each project’s strategy conditions on its own state and on the equilibrium path of aggregate states.²²

6 Identification and Estimation

We estimate the model in three steps, discussing identification at each step. First, we estimate state transition and exogenous exit probabilities directly from the data; these serve as beliefs in the project’s decision problem. Second, taking these beliefs as given, we use simulated maximum likelihood to recover payoffs, waiting costs, and shock distributions from observed withdrawal and completion decisions. Third, we construct a pool of potential entrants, compute their model-implied entry values, and estimate the distribution of entry costs from observed entry decisions. Our estimation sample is projects entering the queue from 2012–2020, with pre-2012 entrants’ actions held fixed; the 2012 start avoids the entry spike from a temporary federal program (Section 3). We use project actions through the second quarter of 2023,²³ before PJM began transitioning from serial to cluster studies.

²²Projects in our model correctly anticipate the non-stationary changes in equilibrium objects that affect payoffs and state-transition probabilities. In particular, we assume that projects have rational expectations over both technology costs, i.e., the time variation in ω , and interconnection-queue congestion. The anticipation of falling renewable costs is consistent with contemporaneous policy and industry forecasts, such as the DOE’s 2012 SunShot scenario and IRENA’s 2016 cost-reduction projections (DOE, 2012; International Renewable Energy Agency, 2016). Projects also understood that increasing renewable development would congest the queue and slow queue processing, as described in AWEA’s petition to FERC and in FERC’s subsequent interconnection-reform Notice of Proposed Rulemaking (AWEA, 2015; FERC, 2016).

²³For projects that received a terminal study in 2023, we supplement these data with completion decisions through the third quarter of 2024.

6.1 Transition and Exogenous Exit Probabilities

We discuss identification and give an overview of estimation of the state transitions and exogenous exit probabilities. Detailed specifications are in Appendix D, with estimates in Appendix Tables D.1 and D.2.

We estimate the study arrival process using a hazard model with a probit link. The specification is a simplified version of the hazard models in Section 4.1.²⁴ It relies on the same identification argument: in PJM’s study process, priority for studies is based only on entry date, and PJM cannot expedite or delay studies in response to the project profitability that drives selection into the surviving sample. Although the estimation sample conditions on survival, withdrawal depends on payoffs while study arrival does not, so survivor selection does not bias the arrival estimates.

We separately estimate the cost and study-information transitions. We use an ordered probit for the cost transition across the four cost bins and probit models for changes in the two study information indicators in z . The identifying assumption is that the error terms in these models are independent of the regressors in the estimation sample. A concern is that developers might selectively withdraw based on private information about future cost changes, biasing the survivor sample. We find no empirical or anecdotal evidence suggesting that developers can correctly predict these cost changes (Section 3.5).

For the distribution of first-study states (c, z) , we take a nonparametric approach. We group projects by observable characteristics, estimate the within-group marginals of c and z , and combine them, assuming that c and z are independent (see Appendix D.4 for details).

Because endogenous withdrawals only occur immediately after a study arrives (Section 5.3), we treat all other withdrawals as exogenous. These exits account for less than half of withdrawals, and they are distributed fairly evenly over the wait window (Section 3.6). We therefore model these exits as a discrete-time hazard that is constant within each study stage, with the stage-specific hazard depending on project characteristics: fuel type, an indicator for a capacity increase at an existing plant, and log capacity. This specification predicts an average per-period exogenous exit probability of 0.029 with little difference between the two stages.

²⁴The transition model does not include calendar time, a simplification supported by the finding that queue-position effects are not sensitive to the inclusion of time controls (Appendix Table E.6).

6.2 Project Payoffs and Waiting Costs

Measured Payoff Components. Estimation takes the revenue from power production $\tilde{r}_{it}$ and construction cost ω_{it} as known inputs to the project’s payoff. We construct $\tilde{r}_{it}$ from hourly PJM wholesale electricity prices over the sample period, summed over a 25-year operating horizon and discounted at 5% (Appendix B). We do not observe the long-term contracts through which many projects sell power. We therefore approximate revenues using observed wholesale prices, which also let revenues respond endogenously to entry in the counterfactuals. For ω_{it} , we use EIA installation costs adjusted for tax credits, multiplying by 0.7 for solar and storage (investment tax credit) (Appendix Table E.12). For wind, we reduce installation costs by the value of the production tax credit (PTC), which we compute from predicted production.²⁵ In counterfactual simulations, $\tilde{r}_{it}$ is recomputed under each candidate equilibrium by reconstructing the supply curve with the updated capacity mix.

Identification. Taking the transition and exogenous exit probabilities, present value of revenue from power production $\tilde{r}_{it}$, and construction cost ω_{it} as given, we use the projects’ withdrawal and completion decisions to identify the following parameters: (1) the payoff parameters θ ; (2) the distributions of the unobserved heterogeneity ξ^{S1} , ξ^{S2} , and ξ^f ; and (3) the waiting costs γ^{S1} and γ^{S2} . Because interconnection costs enter the payoff in dollars per megawatt, alongside measured revenues and construction costs, all parameters are denominated in dollars and no scale normalization is required.

We formally prove the nonparametric identification of (1)–(3) in Appendix C under the conditions that the interconnection costs are continuous.²⁶ The proof proceeds backwards from the terminal study. Projects that reach the third study are selected: they previously chose to remain in the queue, and their unobserved heterogeneity may be correlated with earlier study costs. The identifying variation is therefore not the level of the third-study cost in isolation. Rather, conditional on the state at the second study and the elapsed waiting time, the realized third-study cost provides exogenous variation in the completion probability.²⁷

²⁵We predict the wind production using the capacity factors calculated from EIA data. The PTC rate is based on the in-service date and lasts for 10 years. We assume a 2-year lag between the date of the terminal study and the in-service date, roughly the observed lag in the data. Given the PTC schedule, we assign rates by terminal-study year: \$25/MWh for 2018 or earlier, \$20/MWh for 2019, \$15/MWh for 2020–2022, and \$0 thereafter.

²⁶We assume a finite support for the costs in the existence proof in Appendix A. Therefore, the identification assumptions are satisfied only approximately with discretized costs in the model.

²⁷Some projects may form beliefs about the change in the third-study costs, such as expecting the costs to

The second- and third-study costs play distinct roles in the proof: the second-study cost isolates a marginal group of projects, and the third-study cost traces out the distribution of ξ^f . Conditional on observables, a marginal increase in the second-study cost tips this group into withdrawing: projects exactly at the threshold of indifference after the second study, which share the same accumulated heterogeneity $\xi^{S1} + \xi^{S2}$. The key step compares two derivatives with respect to the second-study cost. The first, the derivative of the probability of withdrawing immediately after the second study, measures the size of the marginal group. The second, the derivative of the joint probability of continuing, receiving a third-study cost, and withdrawing at the terminal study, equals that size times the group's terminal withdrawal probability. Their ratio therefore nets out the density of $\xi^{S1} + \xi^{S2}$ and delivers the terminal withdrawal probability of the marginal project at each third-study cost. Varying the third-study cost, which enters the payoff dollar for dollar, traces out the distribution of ξ^f , with its location pinned by the mean-zero normalization. A second application of the argument, with respect to the first-study cost, identifies the distribution of ξ^{S2} ; the distribution of ξ^{S1} and the payoff parameters then follow from the withdrawal probabilities after the first study.

The waiting costs are identified last: a project at the continuation cutoff is indifferent between withdrawing and waiting, and once the payoffs, distributions of unobserved heterogeneity, and beliefs are identified, the waiting cost is the only remaining unknown in this indifference condition.²⁸ Because the waiting costs are identified from withdrawal behavior, they measure the per-period opportunity cost of remaining an active project, not only direct expenditures.

be revised downwards if the second-study costs are high, but we assume those beliefs are conditional on the information we can observe from the first and the second studies. The change in the interconnection costs of the third study is thus conditionally independent of the selection on ξ s given the information from prior studies.

²⁸Unobserved heterogeneity could also enter through waiting costs, but it is challenging to separately identify heterogeneity in both payoffs and waiting costs. If waiting costs were heterogeneous, continuation cutoffs would vary with them, so the projects tipped by a marginal increase in a study cost would no longer share the same accumulated heterogeneity, and the ratio argument could not recover the distribution of the terminal shock. Nor does the data pattern that motivates unobserved heterogeneity – later-stage projects are less likely to withdraw than recent entrants – distinguish the two sources. Heterogeneous waiting costs generate it through dynamic selection: later-stage projects are increasingly selected on low waiting costs, which, combined with the growing dispersion of ξ^f , lowers withdrawal at the terminal study. However, payoff heterogeneity alone generates the same pattern: projects with higher payoff realizations are more likely both to remain in the queue and to complete at the terminal study. We therefore allow for heterogeneity only in payoffs.

Estimation. Given the estimates of transition and exogenous exit probabilities in Section 6.1, we use simulated maximum likelihood to estimate the payoff parameters, the scale parameters in the parametric distributions of the shocks, and the waiting costs. For the latent (ξ_i^{S1}, ξ_i^{S2}) in each project’s state space, we discretize the support of unobserved heterogeneity to a 30-point grid.²⁹ We then solve the project’s optimal decision problem for each combination of a grid value and observed characteristics, allowing beliefs about state transitions and payoffs to vary across these combinations. For each observed project that enters the queue, we integrate over the unobserved heterogeneity to compute the likelihood of its observed withdrawal and completion decisions.³⁰

Table 6 reports the simulated maximum likelihood estimates. All units are in million dollars per megawatt. In Panel A, we find that project payoffs vary systematically with technology and size. The omitted group is new natural gas generation. In the payoff function, we include the present value of wholesale revenue over the service life for all fuel types except batteries, whose operation involves buying and selling power. We also include construction costs. Therefore, the fuel-type coefficients should be interpreted as the payoff of a fuel type relative to gas after netting out these two components. We estimate higher coefficients for wind and solar, which is not surprising: $\tilde{r}_{it}$ measures revenue rather than operating profit, so the fuel-type coefficients absorb production costs, which are near zero for these technologies and substantial for gas. Batteries have the highest coefficient, likely because we set the present value of their wholesale revenues to 0, so the battery coefficient also absorbs operating profit. Payoffs also increase with project size. Renewable payoffs net of measured revenues and construction costs rise after 2018 relative to gas, consistent with the expansion of state renewable portfolio standards and broader wholesale market access for batteries, policies we do not explicitly model.

Panel B shows substantial waiting costs: \$2,500 per megawatt per quarter while a project waits for the second study and \$1,100 per megawatt per quarter while it waits for the third.

²⁹In the dynamic program, beliefs about future latent values are represented as probability distributions over these grid points. For a project after a non-terminal first study, the belief is over $\xi^{S1} + \xi^{S2}$. Thus, when a project has not yet received a later study realization, its continuation value integrates over the possible future realizations by summing the value function over the grid, weighted by the corresponding probabilities. The grid spans $[-10, 10]$, which is wide relative to the sum of estimated scales of $\xi^{S1} + \xi^{S2}$ in Table 6 even when waiting time is long. The large support makes the approximation insensitive to boundary points in practice.

³⁰Consistent with the model timing, ξ_i^{S2} realizes whenever a second study is received, including when that study is terminal. For projects receiving a terminal first study, the baseline specification sets $\xi_i^{S2} = 0$; an alternative specification allowing an additional realization in this case has a small impact on the results.

We interpret these as the carrying costs of keeping a project in the queue and on schedule. The equivalent annual cost while waiting for the third study, \$4,400 per megawatt, is consistent with a project holding its position: maintaining site control and a minimal development team.³¹ The higher cost while waiting for the second study, \$10,000 per megawatt per year, is consistent with the spending of a project in active development, when developers advance permitting, engineering, and the agreements required for the next study. An itemized estimate of development spending for a utility-scale solar project in PJM implies roughly \$8,000 per megawatt per year excluding posted security and study deposits, or \$13,500–\$16,800 per megawatt per year including the financing cost of posted security (\$110,000 per megawatt in the same itemization) at a cost of capital between 5 and 8% (Leyline Renewable Capital, 2022).

Panel C shows that unobserved payoff heterogeneity is substantial and that its dispersion grows as a project waits. The standard deviation of ξ^{S1} , the heterogeneity revealed at the first study, is \$1.23 million per megawatt, similar to the construction cost of a new project. The dispersions of ξ^{S2} and ξ^f grow with waiting time ($\kappa^{S2} = 0.30$ and $\kappa = 0.23$), and both rates are precisely estimated. Evaluated at 4 quarters, one standard deviation of ξ^{S2} is \$0.81 million per megawatt and one standard deviation of ξ^f is \$0.76 million per megawatt. At 8 quarters, they increase to \$0.99 million per megawatt and \$0.90 million per megawatt, respectively.

6.3 Entry Costs

We construct a set of potential entrants and estimate entry costs. Each potential entrant differs by fuel type, size, whether it increases capacity at an existing plant, time of entry, and transmission owner. We provide details in Appendix D.5.

To ensure that the equilibrium simulation remains tractable, we include 514 non-entrants in the sample, which implies an average entry probability of 84%. As a robustness check, we also report the estimates using twice as many non-entrants.

Identification. Identification follows from the threshold-crossing structure of entry: a potential entrant enters if its entry value exceeds its entry cost draw. Using the payoff parameters

³¹Developers typically hold land under option during development, converting to a lease once the project proceeds; industry estimates of development-period site-control spending for a utility-scale solar project in PJM are roughly \$4,000 per megawatt per year (Leyline Renewable Capital, 2022).

and continuation values from Section 6.2, we compute the entry value in Equation (5) for observed entrants and for potential entrants that do not enter. The entry value is a known dollar amount implied by the prior step’s estimates, so the distribution of entry costs is identified by how entry decisions respond to variation in the entry value.

The key source of variation underlying the entry values is the decline in renewable construction costs ω over time. Construction costs enter the entry value with a known coefficient through the payoff (Equation (1)) but do not enter the entry cost distribution, which is fixed over time conditional on project characteristics (x^{ent} contains no calendar-time terms). The decline in construction costs therefore raises the entry values of later cohorts without changing their entry costs.³²

Because developers in our model anticipate the decline in renewable construction costs, a concern is that they time their decisions to exploit it, either by delaying entry or by entering early and waiting in the queue, which would compress this variation. Delaying entry is not possible in our model because a potential entrant’s opportunity to enter arises only once (Section 5.1). Nor can a project enter and wait inside the queue for costs to decline: studies arrive on PJM’s schedule, and the terminal study forces a completion decision, so early cohorts face lower entry values even though they anticipate the decline. The increase in renewable entry relative to the increase in equilibrium entry values thus identifies the distribution of entry costs.

Unobserved Heterogeneity. Our entry model treats potential projects as ex ante identical up to observed characteristics: the payoff heterogeneity ξ realizes only after entry, and non-entrants differ from entrants only through their entry-cost draws. This approach simplifies estimation and allows us to separately estimate project preferences and entry costs. The key assumption is that heterogeneity known to developers at entry, such as the site quality beyond what observables capture, is limited compared with the post-entry realizations of interconnection costs or unobserved payoff heterogeneity ξ . This assumption is consistent with developers submitting exploratory interconnection requests to learn about interconnection costs (Gorman et al., 2024). Furthermore, if unobserved site quality known to developers

³²Using the prior step’s estimates, we indeed find that the capacity-weighted average value of entering the queue for renewables rises from \$0.10 million per megawatt for the 2012 cohort to \$0.24 million per megawatt for the 2020 cohort. For gas, construction costs are roughly flat over this period, and queue congestion and falling wholesale electricity prices push the value down from \$0.39 to \$0.16 million per megawatt.

were the dominant source of heterogeneity, good sites would attract repeated requests after a withdrawal. Instead, repeated entry is rare (Section 3.1).

Estimation. We estimate the entry cost distribution by maximum likelihood, assuming that the entry cost C is supported on $[0, \bar{C}]$ and that $C/\bar{C}$ has a beta distribution.³³ The entry cost has mean $\exp(x^{\text{ent}}\zeta^{\text{ent}})$, where x^{ent} is the vector of project characteristics reported in Table 7. The implied distribution of $C/\bar{C}$ is $\text{Beta}((\exp(x^{\text{ent}}\zeta^{\text{ent}})/\bar{C})\psi, (1 - \exp(x^{\text{ent}}\zeta^{\text{ent}})/\bar{C})\psi)$, determined by the mean and a free precision parameter ψ ; the corresponding variance of the entry cost is $\frac{\exp(x^{\text{ent}}\zeta^{\text{ent}})(\bar{C} - \exp(x^{\text{ent}}\zeta^{\text{ent}}))}{\psi + 1}$. We set $\bar{C}$ to \$0.6 million per megawatt, just above the largest entry value implied by the prior step’s estimates (\$0.58 million per megawatt), so that non-entry has positive probability at every entry value.

Table 7 shows that entry costs are low for most projects. The baseline estimates in column (1) imply that entry costs are nearly zero at the median and \$0.21 million per megawatt at the 85th percentile, relative to a median entry value of \$0.36 million per megawatt (\$0.32 million per megawatt for renewables). Costs are lower for renewables. These patterns are consistent with entry costs being driven mainly by site-control requirements. Renewable projects may face lower entry costs because they have a larger set of potentially suitable sites than gas plants, which require access to adequate natural gas infrastructure.

Estimates with twice as many non-entrants (column (2)) imply higher entry costs but a similar relationship between these costs and covariates. The increase is especially large at lower quantiles and follows mechanically from the construction: a larger pool implies a lower average entry probability, which the estimation matches by raising entry costs. Both columns imply expected numbers of entrants (2,708 and 2,709) that match the 2,709 observed entrants.

³³A lognormal entry cost yields a more dispersed cost distribution but similar counterfactual results. The costs from a lognormal and the beta distributions are similar at the 85th quantile (0.202 vs. 0.213 million dollars per megawatt), but the lognormal cost implies a higher cost at the 95th quantile (0.832 vs. 0.466 million dollars per megawatt).

Table 6: Estimates of Payoff and Waiting Cost Parameters

<i>Panel A. Payoff Estimates</i>	
Wind	0.713*** (0.076)
Solar	1.080*** (0.065)
Battery	1.367*** (0.054)
Capacity increase at existing plant	0.385*** (0.066)
ln(size)	0.049** (0.023)
Renewables $\times$ ln(size)	0.031 (0.020)
After 2018	-0.117 (0.105)
Renewables $\times$ after 2018	0.430*** (0.117)
Intercept	-1.732*** (0.095)
<i>Panel B. Waiting Costs $\times 1000$</i>	
Waiting for third study	-1.459*** (0.180)
Intercept	2.533*** (0.243)
<i>Panel C. Unobservables: Baseline Scale σ and Waiting Time Effect κs</i>	
Study 1: σ^{S1}	1.226*** (0.127)
Study 2: σ^{S2}	0.532*** (0.034)
Study 2: κ^{S2}	0.300*** (0.013)
Terminal Study: σ	0.556*** (0.032)
Terminal Study: κ	0.230*** (0.008)

Parameters in panels (A) and (C) are in units of \$ million per MW. Panel B's units are in \$ thousand per MW per quarter. Panel A uses gas projects as the omitted group. The after 2018 indicator is based on the terminal study year. The κ parameters are unitless. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors account for the statistical errors in the first-stage transition-function estimates using a Murphy–Topel correction (Murphy and Topel, 2002).

Table 7: Entry Cost Estimates

	(1)	(2)
	Baseline	More Entrants
<i>Panel A. Entry Cost Distribution Parameters</i>		
Renewables	-0.506*** (0.119)	-0.418*** (0.081)
Capacity increase at existing plant	-0.357 (0.217)	-0.245* (0.133)
ln(size)	0.068** (0.034)	0.059*** (0.022)
Intercept	-2.354*** (0.200)	-1.879*** (0.144)
ψ	0.920*** (0.230)	0.863*** (0.168)
<i>Panel B. Implied Entry Cost Quantiles (\$M/MW)</i>		
50th percentile	0.003	0.030
75th percentile	0.079	0.227
80th percentile	0.133	0.302
85th percentile	0.213	0.386
90th percentile	0.325	0.474
95th percentile	0.466	0.553

The entry cost distribution is estimated as a beta distribution on $[0, \bar{C}]$. Let the mean cost $\mu_i = \exp(x_i^{ent} \zeta^{ent})$ and let ψ denote the precision parameter. Then $C_i/\bar{C} \sim \text{Beta}(\alpha_i, \beta_i)$ with $\alpha_i = (\mu_i/\bar{C})\psi$ and $\beta_i = (1 - \mu_i/\bar{C})\psi$. Columns (1) and (2) use 514 and 1,028 potential entrants that do not enter, respectively, along with 2,709 entrants observed in the data. Quantile rows report the implied quantiles of the entry cost distribution in \$M/MW, based on the estimated mixture over observed entry covariates. Standard errors are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors account for the statistical errors in the first-stage transition-function estimates and second-stage model estimates using a Murphy–Topel correction (Murphy and Topel, 2002).

7 Equilibrium Simulation and Model Fit

To simulate the equilibrium in Section 5.3, we start from the observed queue and iteratively update project strategies, queue statuses, project beliefs, and wholesale prices until convergence.³⁴ We use this procedure for the model-fit analysis below and for all counterfactual simulations.

We assume that, in the no-policy (as-is) equilibrium, projects expect wholesale prices after 2020 to remain at the 2021 level. Specifically, the 2021 quarter-hourly prices repeat in each subsequent year; as described in Appendix B, we predict prices for 20 representative hours per quarter, so revenues reflect the prices in the hours when each project produces. The generation-weighted average price is \$31.7/MWh in 2021. This assumption reflects the relatively stable prices in PJM over the decade before 2021. It ignores the surge in demand and prices in later years, driven largely by data centers, which projects entering in 2020 and earlier would not have predicted.

Wholesale prices adjust endogenously in the counterfactuals. The stable-price assumption above implies demand growth, because more projects complete over time and add to PJM supply. We therefore compute the implied demand for 2021–2028 that makes the as-is equilibrium consistent with the stable-price assumption, where 2028 corresponds to the last quarter of action for projects that entered in 2020 and wait the maximum of 32 quarters. In policy counterfactuals, we hold demand fixed — observed for 2012–2020 (Appendix B) and implied for 2021–2028 — and recompute wholesale prices given each counterfactual’s completions. These prices enter the wholesale revenues in projects’ payoff functions. The new equilibrium is achieved when project strategies, queue statuses, the wholesale prices, and project beliefs are consistent with each other. We assume that wholesale prices after 2028 remain at the 2028 equilibrium level. We also report counterfactual policies’ effects on this long-run wholesale price, averaged over the quarter-hourly prices and weighted by generation.


Model Fit. The model fits the data well in the estimation sample (Figure 5). Panels A and B plot predicted and observed queue sizes, which are the numbers of waiting projects at a given time, for all projects and renewables; both predictions track the data closely, slightly under-predicting congestion at the end of the sample. Panels C and D plot cumulative eventual

³⁴We cannot rule out the presence of multiple equilibria, but starting from queue statuses perturbed around the observed values yields highly similar equilibrium outcomes.

completion by entry time, defined as the sum of eventually completed capacity by projects entering in or before a given quarter. The blue line is the model’s prediction; the solid black line shows capacity completed by 2023 (a lower bound on eventual completion), and the dotted black line above it adds capacity still waiting in 2023 (an upper bound). The model prediction is close to the data in early years and stays within the bounds in later years, for both total and renewable capacity. Our model predicts that about 57 GW will eventually be in service in PJM from those entering between 2012 and 2020.

The model prediction is also consistent with the completions observed after the estimation sample. We collected 2026 data on the 2012–2020 entry cohorts; by then, more projects had received terminal studies, and more had completed. Our prediction falls within the implied tighter bounds (indicated by stars).

Figure 5: Model Fit: Queue Size and Completions



Panels A–B plot queue counts by calendar quarter (2012–2020) for all projects and renewable projects, respectively. Panels C–D plot cumulative eventual completed capacity by entry quarter. The solid black line is observed completed capacity by end of sample; the dotted line adds projects still unresolved at the cutoff; the dashed blue line is model-predicted eventual completions. The solid star shows observed completions of those entering between 2012 and 2020 as of 2026; the hollow star adds projects still in queue as of 2026.

8 Counterfactuals

Using the simulation procedure in Section 7, we evaluate two reforms: speeding up engineering studies and raising the cost of entering or remaining in the queue. The counterfactuals allow for equilibrium responses in entry, waiting, completion and wholesale prices.

8.1 Study Speedup

Recent initiatives aim to accelerate engineering studies: grid operators are beginning to use AI to produce them, and FERC Order 2023 penalizes transmission providers for study delays. In the model, we uniformly scale up each component of study arrival probabilities, Q_i^{S1} and Q_i^{S2} in Equations (4) and (3), by 10 to 50%, subject to the constraint that each probability is capped at one.

We find that speedup substantially increases completed capacity. Table 8 reports levels under the as-is equilibrium and changes under each speedup for completion, waiting time (the time between a project’s entry and its completion or withdrawal), entry, entry surplus, and the long-run wholesale price. A 50% speedup increases total completion by 32%. This speedup also reduces waiting time by 19%.³⁵ Entry rises by only 3.3%, so the increase in completed capacity comes mostly from higher completion among projects that enter the queue.

The speedup raises entry surplus and lowers the long-run wholesale price. Entry surplus, defined as the sum of $E \max\{EV_i - C_i, 0\}$ across all potential projects, increases steadily, though less than proportionally at the largest speedup. These surplus figures cover only project developers. The added capacity under a 50% speedup decreases the long-run wholesale electricity price, as described in Section 7, by 4.1%. The lower wholesale price also benefits consumers, though with inelastic demand this is largely a transfer from incumbent generators rather than an additional welfare gain.

Option Value of Waiting. The muted entry response reflects two countervailing effects of the speedup on the value of entering the queue. A waiting project holds an option: it completes only on favorable payoff realizations and withdraws otherwise, so the value of waiting rises

³⁵The reduction in waiting time is proportionally lower than the increase in study arrival probabilities for two reasons. First, our waiting-time measure includes both withdrawals and completions. Speedup increases the incentive to wait in the queue, which can lengthen waiting time for projects that eventually withdraw. Second, about 15% of projects receive a terminal first study, so later-stage speedups do not affect them.

Table 8: Study-Speedup Counterfactuals

	(1)	(2)	(3)	(4)	(5)	(6)
	As-is	10%	20%	30%	40%	50%
Completed (GW)	57.4	+4.4	+8.6	+12.1	+15.1	+18.1
Renewable	27.5	+2.7	+5.3	+7.5	+9.5	+11.4
Non-renewable	29.9	+1.8	+3.3	+4.6	+5.7	+6.7
Waiting time (quarters)	9.9	-0.4	-0.8	-1.1	-1.5	-1.9
Renewable	10.9	-0.4	-0.8	-1.2	-1.6	-2.0
Non-renewable	6.7	-0.3	-0.6	-0.9	-1.2	-1.5
Entry (GW)	262.9	+2.4	+4.6	+6.1	+7.4	+8.6
Renewable	151.8	+1.4	+2.7	+3.6	+4.3	+4.9
Non-renewable	111.1	+1.0	+1.9	+2.6	+3.2	+3.7
Entry surplus (\$Bn)	56.3	+3.2	+6.3	+8.6	+10.5	+12.4
Renewable	29.1	+2.1	+4.2	+5.8	+7.0	+8.3
Non-renewable	27.3	+1.1	+2.1	+2.8	+3.5	+4.1
Wholesale price (\$/MWh)						
Long run	31.7	-0.3	-0.6	-0.9	-1.1	-1.3

Column (1) reports the as-is equilibrium in levels; columns (2)–(6) report changes relative to as-is. Entries and completions are in GW. Waiting time is in quarters. Surplus rows are entry surplus defined as $E \max\{EV_i - C_i, 0\}$ in \$Bn. Long-run price is the 2028 equilibrium wholesale price: the average of the quarter-hourly prices weighted by generation. Demand in 2021–2028 is chosen so that, given as-is capacity, equilibrium prices remain at the 2021 level (Section 7), and the calculated demand is held fixed across counterfactuals.

with the dispersion of those realizations. Because the standard deviations of ξ^{S2} and ξ^f increase with waiting time, a faster queue compresses this dispersion and reduces the option value of remaining in the queue. Offsetting this, faster study arrivals reduce accumulated waiting costs, discounting, and exposure to exogenous exits, raising the value of entering the queue.

To isolate the option-value effect, we compress the dispersion of payoff realizations to reflect the shorter waits under the 50% speedup. Specifically, for each waiting time τ , we replace the standard deviation of ξ^f , $\sigma \cdot \tau^\kappa$, with $\sigma \cdot (\max\{0, \tau - \Delta t\})^\kappa$, where $\Delta t = 1.9$ quarters is the average reduction in waiting time under the speedup. We do the same for ξ^{S2} . We hold study arrival probabilities and queue statuses at their as-is equilibrium values and allow projects to re-optimize once given the smaller standard deviations.

We find that compressing the dispersion of payoff realizations alone would lower completion, entry, and surplus relative to the as-is equilibrium (column (3) of Table 9). Projects also withdraw faster, leading to a slight reduction in waiting time.³⁶ This decomposition provides an additional explanation, beyond the high entry costs of remaining non-entrants, for why the entry response to a 50% speedup is smaller than the completion response. The reduction in option value would decrease total entry by 2.4 GW. The offsetting increase in entry value from lower waiting costs and avoided exogenous exits accounts for the remaining 11.0 GW. The net effect is the 8.6 GW increase.

Effects of Dynamic Strategies. To quantify the role of equilibrium strategy adjustments, we increase the study arrival probabilities while holding project strategies fixed at the as-is equilibrium. Starting from the as-is equilibrium queue, we simulate new queue statuses and compute completion, waiting time, and entry surplus under the higher arrival probabilities and the updated queue statuses.

At the 50% speedup, ignoring the equilibrium strategy adjustments understates the gains in total completion by 30% (12.6 GW rather than 18.1 GW; column (4) of Table 9). Because entry is mechanically unchanged, the difference in completions comes entirely from projects waiting more and receiving more terminal studies. The entry surplus gain is 18% smaller when strategies are fixed.

³⁶Queue statuses and study arrival probabilities are held fixed, so this reduction in waiting time is entirely from faster withdrawals.

Number of Potential Entrants. The increases in completed capacity from the speedup are not sensitive to the assumed size of the potential entrant pool. The simulations in Table 8 include 514 potential entrants that do not enter in the as-is equilibrium; we re-simulate the speedup counterfactuals with twice as many non-entrants, using the corresponding entry cost estimates in column (2) of Table 7. Under a 50% speedup, entry increases by 11.8 GW (6.7 GW renewable), 37% more than in Table 8, while completed capacity increases by 17.9 GW (11.3 GW renewable), only 1% less. The additional entrants congest the queue and crowd out completions among the original pool, which explains why the increases in completed capacity barely change.³⁷

Capacity Market Revenue. The additional capacity from a study speedup would lower prices in PJM’s capacity market, but the effect is small for the renewable projects that account for most of the additional completed capacity. PJM’s capacity market pays resources for being available to meet peak demand, separately from their energy market earnings. Our specification captures these payments with the time-invariant fuel-type coefficients, but it does not capture how new capacity would lower the capacity price. We provide an upper bound on the capacity price decline using the published demand and supply curves from the 2021/2022 Base Residual Auction (PJM Interconnection, 2018*b,a*). We shift the supply curve outward by the capacity the 50% speedup adds, scaling each technology’s capacity by its Effective Load Carrying Capability. The resulting 9.66 GW shift³⁸ lowers the price by 17%, from \$140 to \$116 per megawatt-day.³⁹

Even the upper-bound 17% price decline has a limited effect on entry values and completion. For renewable projects, which supply most of the added capacity and earn little capacity revenue per megawatt, the price decline lowers lifetime revenue by an amount equal to about

³⁷Compared with the additional entrants in Table 8, the additional entrants in this simulation tend to be larger projects in later cohorts, but their average completion probabilities are similar (28% vs. 27%).

³⁸We take the additional completed capacity, 6.67 GW gas, 6.33 GW solar, 4.12 GW wind, and 0.99 GW battery, and convert each to accredited capacity using class-average Effective Load Carrying Capability factors (0.92 for gas, 0.38 for solar, 0.13 for wind, 0.50 for battery, (PJM Interconnection, 2021)), yielding 9.66 GW of accredited capacity. We shift the supply curve rightward by this amount, offering the added capacity at a price of zero, and find the new intersection with demand (Variable Resource Requirement curve). New renewables offer at or near zero in practice. For gas, the new projects likely have lower costs than the existing capacity (about 90 GW), and the capacity supply is relatively flat in the region 90 GW below the clearing capacity level. Therefore, shifting the supply curve to the right by 9.66 GW is a reasonable approximation for the upper bound of the counterfactual’s impact on the capacity market revenue.

³⁹We use the system-wide price. PJM also clears higher prices in transmission-constrained zones. The drop in the system-wide price passes through to unconstrained zones, but the price decrease in constrained zones may be less. Therefore, our calculation likely provides an upper bound on the revenue loss.

15% of their average surplus from entering the queue. The effect is larger for natural gas projects, whose capacity payments are a larger share of revenue: the lost capacity revenue equals about 35% of their average entry surplus. These revenue losses have a more limited effect on the expected value of entering the queue because a project realizes them only if it completes, and most projects do not. Relative to the 50% speedup result without this price decline, the expected entry values decrease by 2.8% for gas and 1.1% for renewables, resulting in 1.3 GW lower total completion. The decrease in gas completion accounts for 74% of the total decrease. Our results are thus robust to the capacity market price response.

Environmental Benefits. The additional renewable capacity under the 50% speedup avoids 13.2 million metric tons of CO₂ annually. We compute these benefits using EPA’s Avoided Emissions and geneRation Tool (AVERT v4.3, April 2024), which estimates region-specific marginal emissions avoided when new renewable generation displaces fossil generation. For PJM, AVERT’s capacity factors imply that the counterfactual’s additional 4.12 GW of wind and 6.33 GW of solar generate 10.6 and 12.3 TWh per year, respectively; applying AVERT’s 2023 PJM marginal CO₂ emissions rates gives the 13.2 million metric tons.⁴⁰ At a social cost of carbon of \$185 per metric ton (Rennert et al., 2022), these avoided emissions are worth \$2.4 billion per year.

8.2 Fees

Another key component of the ongoing interconnection queue reforms is raising the cost of entering and waiting in the queue to reduce congestion. We focus on per megawatt fees, paid at queue entry, and the designs we consider differ in how much of the fee is refunded and when. We then compare their effects with those of flat fees.

Per Megawatt Entry Fees. The first design we consider is a nonrefundable fee, which the project pays upon entry and does not recover regardless of its outcome. Second, we consider a deposit that is fully at risk: the deposit is returned only if the project completes interconnection, so projects that withdraw forfeit the entirety of the deposit. Third, more similar to the PJM reform, projects pay a deposit upon entry, and, after the first study,

⁴⁰The avoided air pollutants, such as PM_{2.5}, would increase the benefits, while the 6.7 GW increase in gas capacity has an ambiguous effect, since its emissions may be offset by displacing coal-fired generation.

Table 9: Decomposition of 50% Speed Up Effects

	(1)	(2)	(3)	(4)
	As-is	50% speedup	Option value	Strategies fixed
Completed (GW)	57.4	+18.1	-1.1	+12.6
Renewable	27.5	+11.4	-0.5	+7.5
Non-renewable	29.9	+6.7	-0.6	+5.2
Waiting time (quarters)	9.9	-1.9	-0.1	-1.5
Renewable	10.9	-2.0	-0.1	-1.6
Non-renewable	6.7	-1.5	-0.1	-1.2
Entry (GW)	262.9	+8.6	-2.4	+0.0
Renewable	151.8	+4.9	-1.0	+0.0
Non-renewable	111.1	+3.7	-1.4	+0.0
Entry surplus (\$Bn)	56.3	+12.4	-3.0	+10.2
Renewable	29.1	+8.3	-1.4	+5.8
Non-renewable	27.3	+4.1	-1.6	+4.4

Column (1) reports the as-is equilibrium in levels; columns (2)–(4) report changes relative to as-is. Entries and completions are in GW. Waiting time is in quarters. Surplus rows are entry surplus defined as $E \max\{0, EV_i - C_i\}$ in \$Bn. Δ is the change relative to as-is. The Option value columns take the as-is equilibrium queue statuses and rescale the standard deviations of ξ^f to $\max(0, t - \Delta t)^\kappa$ (and similarly for ξ^{S2}) for each value of waiting time τ , where $\Delta t = 1.9$ quarters is the average reduction in waiting time from the speedup. We then allow the projects to re-optimize once, responding to the as-is equilibrium queue statuses but ξ^{S2} and ξ^f shocks with smaller standard deviations. In Strategies fixed columns, we take both the as-is equilibrium queue statuses as well as the project strategies, increase the study arrival probability and simulate new queue statuses. The completion, waiting time and the entry surplus are computed given the increased study arrival probability, the new queue statuses and the as-is strategies.

receive a 50% refund if they choose to withdraw rather than request the second study. Should they choose to continue the queuing process, they would not recover the deposit unless they complete interconnection. Therefore, the recoverability of the deposits is staged and decreases as projects advance in the queue.

We additionally consider budget-neutral versions of the two deposits. In these versions, the deposits forfeited by withdrawing projects are redistributed to completing projects on a per megawatt basis, so that PJM does not retain revenue. Completing projects therefore receive an amount greater than the deposit they paid at entry. Although such designs may be hard to implement in real time, since the rebate depends on future withdrawals and completions, they provide a useful benchmark for a fee design that fully recycles entrant payments to completers.⁴¹

Fees and deposits change completion by at most a few percent, with the sign determined by whether payments are refunded. Table 10 reports the effects of each design at \$20,000 per megawatt. We report effects at this level because it is high enough to distinguish the designs and plausible: PJM adopted a comparable fee in 2026 for a temporary program that accelerates some projects to address short-term reliability concerns (PJM Interconnection, 2026). For comparison, the level PJM is adopting in wider reforms is \$4,000 per megawatt. The nonrefundable fee reduces completion by 0.8%. The four refundable designs raise completion by 0.2% to 4.0%, with larger gains when the deposit is staged and, especially, when forfeitures are rebated to completing projects. Across the refundable designs, these gains arise from the incentive created by the refund rather than from reduced time in the queue.

Under the nonrefundable fee, average waiting time rises slightly because the fee shifts the composition of the queue, not because projects wait longer. When we group projects by fuel type, location, entry time, and size, average waiting time falls in most groups (Appendix Figure F.5). Yet the average waiting time among entrants does not fall: the fee disproportionately deters small projects, which often need fewer than three studies (Table F.14), so longer-waiting projects make up a larger share of the queue.

Waiting time rises under the deposit and the budget-neutral deposit but falls under the two staged deposits. These results reflect two opposing forces: refunds tied to completion

⁴¹Our budget-neutral designs are realistic in retaining no revenue for PJM but stylized in where forfeitures go. Under the ongoing reform, forfeited deposits fund restudies, with any remainder going to other projects in the cluster, whereas our designs direct all forfeitures to completing projects. The reform also requires additional deposits tied to network upgrade costs, which we abstract from.

Table 10: Fee and Deposit Counterfactuals at \$20,000 per MW

	As-is	Fee	Deposit		Budget neutral	
	(1)	(2)	(3)	(4)	(5)	(6)
Completed (GW)	57.4	-0.5	+0.1	+0.5	+2.1	+2.3
Renewable	27.5	-0.2	+0.1	+0.3	+1.0	+1.2
Non-renewable	29.9	-0.3	+0.0	+0.2	+1.0	+1.1
Waiting time (quarters)	9.9	+0.1	+0.1	-0.4	+0.3	-0.2
Renewable	10.9	+0.1	+0.2	-0.4	+0.4	-0.2
Non-renewable	6.7	-0.0	-0.0	-0.2	+0.1	-0.2
Entry (GW)	262.9	-4.5	-3.9	-2.8	-1.6	-0.9
Renewable	151.8	-2.6	-2.3	-1.6	-1.2	-0.7
Non-renewable	111.1	-1.9	-1.6	-1.2	-0.4	-0.2
Entry surplus (\$Bn)	56.3	-4.8	-4.0	-3.1	-1.2	-0.6
Renewable	29.1	-2.8	-2.4	-1.8	-0.9	-0.5
Non-renewable	27.3	-2.0	-1.6	-1.3	-0.3	-0.1
Total fees (\$Bn)						
Collected at entry	0.0	+5.2	+5.2	+5.2	+5.2	+5.2
Returned to projects	0.0	+0.0	+1.2	+1.6	+5.2	+5.2
Retained	0.0	+5.2	+4.0	+3.6	-0.0	-0.0
Wholesale price (\$/MWh)	31.7	+0.0	-0.0	-0.0	-0.2	-0.2
<i>Design: share of payment forfeited</i>						
If withdraw before requesting the 2nd study		100%	100%	50%	100%	50%
If withdraw later		100%	100%	100%	100%	100%
If complete		100%	0%	0%	0%	0%
Forfeitures rebated to completers		No	No	No	Yes	Yes

Column (1) reports the as-is equilibrium in levels; columns (2)–(6) report changes relative to as-is. All designs charge \$20,000 per MW of nameplate capacity at queue entry. The design rows report the share of that payment forfeited at each outcome. In column (2), the entry fee is not refundable. In column (3), the deposit is refundable when the project completes but forfeited if the project withdraws. In column (4), the deposit is 50% refundable if the project withdraws before requesting the second study; thereafter, the deposit is refunded if the project completes but forfeited if it withdraws. In columns (5) and (6), the forfeiture rules are the same as in columns (3) and (4), respectively, before a project completes, but all forfeited deposits are rebated to completing projects on a per megawatt basis. Entry surplus is $E \max\{0, EV_i - C_i\}$. Wholesale price is the long-run price.

induce longer waits, while staging induces earlier withdrawals. Under the deposit, 41% of entering projects wait longer than in the as-is equilibrium. Under the budget-neutral deposit, all entering projects wait longer because this design adds the expectation of a rebate upon completion. Staging works in the opposite direction: projects unlikely to complete withdraw after the first study, before the deposit becomes fully at risk.⁴² Under the staged deposit, 1,837 projects continue in the queue after the first study, compared with 1,980 in the as-is equilibrium and 1,956 under the deposit, and waiting time falls by 0.4 quarters. The two forces partially offset each other under the budget-neutral staged deposit, where waiting time falls by 0.2 quarters.

All fees reduce entry, by 0.9 to 4.5 GW, or at most 1.7%. Refunds weaken the deterrent: entry falls least under the budget-neutral deposits, which return every dollar to the queue, and most under the nonrefundable fee.

Because the completion effects are small, so are the wholesale price effects. The long-run wholesale price falls by \$0.2/MWh, or 0.6%, under the budget-neutral deposits and is essentially unchanged under the others.

The magnitude of each design’s effect on completed capacity increases with the fee level, while the sign remains unchanged for fees up to \$50,000 per megawatt. Appendix Figure F.6 plots completed capacity, in total and for renewables, across this range. Completion rises with the fee under the four refundable designs and falls under the nonrefundable fee. At still higher deposit levels not shown in the figure, completion eventually declines for all four. Under the staged deposit, for example, completion peaks near \$100,000 per megawatt at 59.5 GW, or 3.6% above as-is. At the reform level of \$4,000 per megawatt, the effects are smaller than at \$20,000 per megawatt but in the same direction.

Flat Entry Fees. We find that a flat fee, which charges the same dollar amount regardless of project size, has modest but positive effects on completion. Appendix Figure F.7 plots the effects up to \$400,000 per project, which is equivalent to a \$20,000 per megawatt fee for the median-sized project (20 MW). At this level, the nonrefundable fee and the deposit increase completed capacity by 0.7% and 0.6%, respectively, and renewable capacity by 0.4%

⁴²An alternative staged deposit refunds 50% until projects request the third study, moving the at-risk transition one study later. This design adds 0.48 GW more completed capacity than the staged deposit in Table 10 and reduces waiting time by a further 0.29 quarters. Paralleling the staged deposit’s effect after the first study, this design reduces the number of projects waiting for the third study, where congestion has a larger effect on study arrival (Table 3).

and 0.5%, relative to the as-is level. The staged deposit increases completed capacity slightly more, by 1.1%, and renewable capacity by 1.2%. The three designs have more similar effects under the flat fee than under the per megawatt fee: the flat fee mainly deters small projects, and for the larger projects that remain, the fee is small per megawatt, so the refund provisions matter little.

The flat fee reduces the entry and surplus of small projects, with little effect on large ones. In Appendix Tables F.13 and F.14, we compare a \$20,000 per megawatt fee with a \$400,000 flat fee; the two are equal for the median-sized project (20 MW), so the flat fee is higher on a per megawatt basis for smaller projects and lower for larger ones. The flat fee deters about three times as many renewable entrants as the per megawatt fee, nearly all below 20 MW, and reduces their surplus by more; projects above 100 MW are essentially unaffected. In contrast, the per megawatt fee reduces surplus roughly evenly across sizes. In practice, grid operators charge a mix of flat and per megawatt fees, and these distributional effects may explain why the per megawatt component tends to receive much more weight.

Discussion. Nonrefundable entry fees or deposits fully at risk at entry do little to relieve congestion. They deter the lowest-value projects that would have spent little time in the queue. Higher-value projects, which are more willing to wait and more likely to complete, contribute the most to congestion and still enter the queue.

The staged deposit, in contrast, does more to relieve congestion because it screens on the first study's results rather than on expected value at entry. Relative to the unstaged deposit, more projects enter but fewer remain after the first study, and those that remain are more likely to complete. Even so, staging improves completion only modestly because projects that remain later in the queue are positively selected: they have higher value and are more likely to complete. Putting deposits at risk (or equivalently, levying fees) later in the queue would increasingly burden projects that are likely to complete. The same positive correlation between project value and congestion limits the gains from fees throughout the queue.

Fees nonetheless raise revenue. For example, at the \$20,000 per megawatt level, the deposit design retains \$4.0 billion in forfeited deposits from withdrawing projects (Table 10). Our speedup counterfactuals abstract from the cost of delivering faster studies, which in practice requires real resources such as additional engineering capacity.⁴³ A bundled reform could

⁴³We do not model PJM's cost of producing studies and therefore cannot quantify how much acceleration

therefore use this revenue to fund study acceleration.

9 Conclusion

We use novel data from PJM, the largest U.S. regional grid operator, to study the interconnection queue. We find a congestion externality: if a project has more projects ahead of it in the queue, the probability of study arrival in a given quarter falls. We also find that interconnection costs are a key factor in projects' decisions to withdraw from the queue.

We next study policy reforms using a dynamic model that accounts for equilibrium effects. We find that speeding up study delivery can significantly reduce waiting time and increase completion. Projects of all fuel types and sizes benefit, even when entering projects adjust their strategies and wait longer. Speeding up study delivery by 50% adds more renewables to the grid, yielding \$2.4 billion per year in environmental benefits from avoided carbon emissions.

In contrast, imposing fees at levels comparable to current interconnection queue reforms has more limited effects on completion and does not reduce waiting time. This result reflects a positive correlation between a project's payoff and its impact on congestion: the highest-value projects are the most willing to wait and create the most congestion. A potential implication is that screening projects by their value may not relieve queue congestion.

Our analysis covers PJM's serial study process, which governed the queue throughout our sample period, but the economic mechanisms we identify extend beyond that institutional design. FERC Order 2023 requires a transition to cluster studies, in which interconnection requests are grouped and studied together. A cluster-study design still depends on engineering studies that resolve cost uncertainty. In both designs, faster studies reduce waiting costs and exposure to exogenous exits, and the projects most willing to wait are the most valuable. We therefore expect substantial returns to study acceleration under the cluster processes as well. At the same time, cluster-wide cost allocation likely makes interconnection costs more interdependent than they were under the small cost-share groups in our sample, a feature that may make entry screening more valuable under cluster designs. Evaluating these designs directly is an important avenue for future research.

this revenue would finance. Grid operators attribute study delays in part to constraints on engineering staff (Shoemaker, 2021).

The reforms underway across the U.S. bundle the two instruments we analyze: AI-assisted automation and penalties on transmission providers for delayed studies target study speed, while stricter deposits and readiness milestones screen entry. Our analysis suggests that returns to accelerating studies are large, and that, unlike entry fees, they raise the surplus of developers of every size and fuel type.

References

- Agarwal, Nikhil, Itai Ashlagi, Michael A Rees, Paulo Somaini, and Daniel Waldinger.** 2021. “Equilibrium allocations under alternative waitlist designs: Evidence from deceased donor kidneys.” *Econometrica*, 89(1): 37–76.
- Alagappan, Lakshmi, Ren Orans, and Chi-Keung Woo.** 2011. “What drives renewable energy development?” *Energy policy*, 39(9): 5099–5104.
- Aldy, Joe, Todd Gerarden, and Richard Sweeney.** 2023. “Investment versus Output Subsidies: Implications of Alternative Incentives for Wind Energy.” *Journal of the Association of Environmental and Resource Economists*, 10(4): 981–1018.
- AWEA.** 2015. “Petition for Rulemaking of the American Wind Energy Association to Revise Generator Interconnection Rules and Procedures.” *Federal Energy Regulatory Commission, Docket No. RM15-21-000*, States that developers submitted numerous interconnection applications, interconnection queues grew larger, and queue processing became slower.
- AWEA.** 2019. “Wind Powers America Annual Report 2019.” <https://www.awea.org/resources/publications-and-reports/market-reports>.
- Blundell, Wesley, Gautam Gowrisankaran, and Ashley Langer.** 2020. “Escalation of scrutiny: The gains from dynamic enforcement of environmental regulations.” *American Economic Review*, 110(8): 2558–2585.
- Butters, R. Andrew, Jackson Dorsey, and Gautam Gowrisankaran.** 2021. “Soaking up the sun: Battery investment, renewable energy, and market equilibrium.” NBER Working Paper 29133.
- Callaway, Brantly, and Pedro H.C. Sant’Anna.** 2021. “Difference-in-Differences with multiple time periods.” *Journal of Econometrics*, 225(2): 200–230.
- Caspary, Jay, Michael Goggin, Rob Gramlich, and Jesse Schneider.** 2021. “Disconnected: The Need for a New Generator Interconnection Policy.” Americans for a Clean Energy Grid Report.
- Chen, Luming.** 2025. “The Dynamic Efficiency of Policy Uncertainty: Evidence from the Wind Industry.” Working Paper.
- Cicala, Steve.** 2022. “Imperfect markets versus imperfect regulation in US electricity generation.” *American Economic Review*, 112(2): 409–441.

- CleanChoice Energy.** 2026. “CleanChoice Energy to Triple its Generation Capacity with the Acquisition of Two Solar Projects in North Carolina.” <https://www.globenewswire.com/news-release/2026/02/18/3240305/0/en/cleanchoice-energy-to-triple-its-generation-capacity-with-the-acquisition-of-two-solar-projects.html>, Press release, February 18, 2026. Accessed August 12, 2026.
- Davis, Lucas, and Catherine Hausman.** 2016. “Market Impacts of a Nuclear Power Plant Closure.” *American Economic Journal: Applied Economics*, 8(2): 92–122.
- Davis, Lucas W., Catherine Hausman, and Nancy L. Rose.** 2023. “Transmission Impossible? Prospects for Decarbonizing the US Grid.” NBER Working Paper 31377.
- Deschenes, Oliver, Christopher Malloy, and Gavin McDonald.** 2023. “Causal effects of Renewable Portfolio Standards on renewable investments and generation: The role of heterogeneity and dynamics.” *Resource and Energy Economics*, 75(101393).
- DOE.** 2012. “SunShot Vision Study: A Comprehensive Analysis of the Potential for U.S. Solar Electricity Generation.” U.S. Department of Energy DOE/GO-102012-3037; NREL/BK-5200-47927. Models a SunShot scenario in which solar technology costs decline by about 75 percent between 2010 and 2020.
- Doshi, Gaurav.** 2026. “Chicken and Egg: Electricity Transmission and Investment in Renewable Energy.” *Journal of Public Economics*, 256(105598).
- Driscoll, William.** 2022. “Interconnection delays and costs are the biggest barrier for utility-scale renewables, say developers.” <https://pv-magazine-usa.com/2022/02/14/interconnection-delays-and-costs-are-the-biggest-barrier-for-utility-scale-renewables-say-developers/>
- Elliott, Jonathan T.** 2021. “Investment, Emissions, and Reliability in Electricity Markets.”
- EPA.** 2026. “Electric Power Sector Emissions.” <https://www.epa.gov/ghgemissions/electric-power-sector-emissions>.
- Ericson, Richard, and Ariel Pakes.** 1995. “Markov-perfect industry dynamics: A framework for empirical work.” *The Review of economic studies*, 62(1): 53–82.
- Fell, Harrison, Daniel Kaffine, and Kevin Novan.** 2021. “Emissions, Transmission, and the Environmental Value of Renewable Energy.” *American Economic Journal: Economic Policy*, 13(2): 241–272.
- FERC.** 2016. “Reform of Generator Interconnection Procedures and Agreements: Notice of Proposed Rulemaking.” *157 FERC ¶ 61,212, Docket No. RM17-8-000*, Proposes reforms to improve certainty and transparency and to reduce withdrawals, restudies, delays, and multiple interconnection requests.
- FERC.** 2023. “Docket No. RM22-14-000; Order No. 2023: Improvements to Generator Interconnection Procedures and Agreements.” <https://www.ferc.gov/media/e-1-order-2023-rm22-14-000>.
- Fowlie, Meredith, Mar Reguant, and Stephen P Ryan.** 2016. “Market-based emissions regulation and industry dynamics.” *Journal of Political Economy*, 124(1): 249–302.

- Gergen, Michael J, George D Cannon Jr, and Shannon D Torgerson.** 2008. “A modest proposal: A market-based approach to generation interconnection process reform.” *The Electricity Journal*, 21(9): 8–18.
- 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.
- Gorman, Will, Julie Mulvaney Kemp, Joseph Rand, Joachim Seel, Ryan Wiser, Nick Manderlink, Fredrich Kahrl, Kevin Porter, and Will Cotton.** 2024. “Grid connection barriers to renewable energy deployment in the United States.” *Joule*.
- Gowrisankaran, Gautam, Ashley Langer, and Nicholas Ryan.** 2026. “Plugging in the Wind and Sun: Grid Integration with Cost Externalities and Uncertainty.” Working paper.
- Gowrisankaran, Gautam, Ashley Langer, and Wendan Zhang.** 2022. “Policy Uncertainty in the Market for Coal Electricity: The Case of Air Toxics Standards.” NBER Working Paper 30297.
- Gowrisankaran, Gautam, Stanley S Reynolds, and Mario Samano.** 2016. “Intermittency and the value of renewable energy.” *Journal of Political Economy*, 124(4): 1187–1234.
- Hale, Zach.** 2021. “Solar developer adds to pile of complaints targeting PJM study delays.”
- Hale, Zach, and Molly Christian.** 2023. “FERC approves ‘historic’ rule to clear backlog of US generation projects.”
- Hausman, Catherine.** 2025. “Power flows: transmission lines, allocative efficiency, and corporate profits.” *American Economic Review*, 115(8): 2574–2615.
- Hitaj, Claudia.** 2013. “Wind power development in the United States.” *Journal of Environmental Economics and Management*, 65: 394–410.
- Hopenhayn, Hugo A.** 1992. “Entry, exit, and firm dynamics in long run equilibrium.” *Econometrica: Journal of the Econometric Society*, 1127–1150.
- Huitfeldt, Ingrid, Victoria Marone, and Daniel C Waldinger.** 2024. “Designing dynamic reassignment mechanisms: Evidence from GP allocation.” National Bureau of Economic Research.
- IEA.** 2025. “CO2 Emissions – Global Energy Review 2025 – Analysis.” <https://www.iea.org/reports/global-energy-review-2025/co2-emissions>.
- IEA.** 2026. “Electricity 2026.” <https://www.iea.org/reports/electricity-2026>.
- Igami, Mitsuru.** 2017. “Estimating the innovator’s dilemma: Structural analysis of creative destruction in the hard disk drive industry, 1981–1998.” *Journal of Political Economy*, 125(3): 798–847.

- International Renewable Energy Agency.** 2016. “The Power to Change: Solar and Wind Cost Reduction Potential to 2025.” International Renewable Energy Agency, Abu Dhabi. Projects 2015–2025 LCOE reductions of 26 percent for onshore wind, 35 percent for offshore wind, at least 37 percent for CSP, and 59 percent for solar PV.
- Johnston, Sarah.** 2019. “Nonrefundable Tax Credits versus Grants: The Impact of Subsidy Form on the Effectiveness of Subsidies for Renewable Energy.” *Journal of the Association of Environmental and Resource Economists*, 6(3): 433–460.
- Johnston, Sarah, Yifei Liu, and Chenyu Yang.** 2026. “Grid connection costs as a barrier to new generation: Evidence from PJM and implications for transmission policy.” *Energy Economics*, 161: 109540.
- LaRiviere, Jacob, and Xueying Lyu.** 2022. “Transmission constraints, intermittent renewables and welfare.” *Journal of Environmental Economics and Management*, 112(102618).
- Leisten, Matthew, and Nicholas Vreugdenhil.** 2023. “Dynamic Regulation with Firm Linkages: Evidence from Texas.” Working Paper.
- Leyline Renewable Capital.** 2022. “The Growing Impact of Delays on Solar Development Costs across Different Regions.” <https://www.leylinecapital.com/news/the-growing-impact-of-delays-on-solar-development-costs-across-different-regions>, Reports an itemized development-cost estimate for a 100 MW utility-scale solar project in PJM over a 25-month timeline, including \$804,900 of site-control costs and \$11 million of posted upgrade security. Accessed June 10, 2026.
- Liu, Tracy, Zhixi Wan, and Chenyu Yang.** 2024. “Dynamic matching on a commuter carpooling platform.” Working Paper.
- Liu, Yifei.** 2025. “Permitting, Litigation Risk, and Energy Infrastructure Investment.”
- Metcalf, Gilbert.** 2010. “Investment in energy infrastructure and the tax code.” In *Tax policy and the economy*, Vol. 24., ed. Jeffrey Brown. Chicago:University of Chicago Press.
- Murphy, Kevin M, and Robert H Topel.** 2002. “Estimation and Inference in Two-step Econometric Models.” *Journal of Business & Economic Statistics*, 20(1): 88–97.
- Nilson, Robi, Ben Hoen, and Joseph Rand.** 2024. “Survey of Utility-Scale Wind and Solar Developers Report [Slides].” Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA (United States).
- Petersen, Alisa, Katie Siegner, and John Coequyt.** 2026. “The interconnection queue continues to be a barrier to American economic competitiveness.” *RMI*. <https://rmi.org/interconnection-reform-ai-data-centers-generator-queues/>.
- PJM.** 2020. “PJM Manual 14B: PJM Region Transmission Planning Process.” 48 ed., <https://www.pjm.com/-/media/DotCom/documents/manuals/archive/m14b/m14bv48-pjm-regional-transmission-planning-process-10-01-2020.pdf>.

- PJM.** 2021. “Connecting to the Grid FAQs: How long does the interconnection process take?” <https://learn.pjm.com/three-priorities/planning-for-the-future/connecting-grid/how-long-does-the-interconnection-process-take>.
- PJM.** 2023a. “PJM Manual 14A: New Services Request Process.” 30 ed., <https://www.pjm.com/-/media/DotCom/documents/manuals/m14a.ashx>.
- PJM.** 2023b. “PJM Manual 14G: Generation Interconnection Requests.” 8 ed., <https://www.pjm.com/-/media/DotCom/documents/manuals/m14g.pdf>.
- PJM.** 2025. “PJM Manual 14C: Generation and Transmission Interconnection Facility Construction.” 17 ed.
- PJM.** 2026. “State of the Market Report for PJM 2025.” https://www.monitoringanalytics.com/reports/PJM_State_of_the_Market/2025/2025-som-pjm-vol1.pdf.
- PJM Interconnection.** 2018a. “2021/2022 RPM Base Residual Auction Results.” <https://www.pjm.com/-/media/DotCom/markets-ops/rpm/rpm-auction-info/2021-2022/2021-2022-base-residual-auction-report.ashx>, Reports a Resource Clearing Price of \$140.00 per MW-day and 163,627.3 MW of unforced capacity cleared in the rest-of-RTO.
- PJM Interconnection.** 2018b. “2021/2022 RPM Base Residual Auction Supply Curves.” <https://www.pjm.com/-/media/DotCom/markets-ops/rpm/rpm-auction-info/2021-2022/2021-2022-bra-supply-curves.ashx>, Reliability Pricing Model Base Residual Auction for the 2021/2022 Delivery Year. Figure 1 plots the RTO capacity supply curve and the Variable Resource Requirement demand curve.
- PJM Interconnection.** 2021. “Effective Load Carrying Capability (ELCC) Report.” <https://www.pjm.com/-/media/DotCom/planning/res-adeq/elcc/elcc-report-for-july-2021-results.ashx>, Capacity accreditation ratings by resource class.
- PJM Interconnection.** 2025a. “PJM Announces Application Deadline for First Cycle of New Interconnection Process.” *PJM Inside Lines*.
- PJM Interconnection.** 2025b. “PJM, Google & Tapestry Join Forces to Apply AI to Enhance Regional Planning, Generation Interconnection.” *Press release, April 10, 2025*.
- PJM Interconnection.** 2026. “Expedited Interconnection Track Tariff Sheets.” <https://www.pjm.com/-/media/DotCom/committees-groups/cifp-lla/2026/20260203/20260203-expedited-interconnection-track-tariff-sheets.pdf>, Tariff, Part X [NEW]. Accessed May 13, 2026.
- PJM Interconnection, L.L.C.** 2020. “Open Access Transmission Tariff, Part VI (Administration and Study of New Service Requests), Attachment P (Interconnection Construction Service Agreement), Appendix 2.” *Version effective April 1, 2020 (FERC Docket No. ER19-1958)*, PJM designated as Transmission Provider; interconnected Transmission

- Owner builds the Transmission Owner Interconnection Facilities under the Standard Option, sec. 3.2. Available at <https://agreements.pjm.com/eTariff/>.
- Plumer, Brad.** 2023. “The U.S. Has Billions for Wind and Solar Projects. Good Luck Plugging Them In.” *The New York Times*. <https://www.nytimes.com/2023/02/23/climate/renewable-energy-us-electrical-grid.html>.
- Rand, Joseph, Nick Manderlink, Steven Zhang, Chris Talley, Will Gorman, Ryan H. Wiser, Joachim Seel, Julie Mulvaney Kemp, Seongeun Jeong, and Fredrich Kahrl.** 2025. “Queued Up: 2025 Edition, Characteristics of Power Plants Seeking Transmission Interconnection As of the End of 2024.” Lawrence Berkeley National Laboratory.
- Reguant, Mar.** 2019. “The efficiency and sectoral distributional impacts of large-scale renewable energy policies.” *Journal of the Association of Environmental and Resource Economists*, 6(S1): S129–S168.
- Rennert, Kevin, Frank Errickson, Brian C. Prest, Lisa Rennels, Richard G. Newell, William Pizer, Cora Kingdon, Jordan Wingenroth, Roger Cooke, Bryan Parthum, David Smith, Kevin Cromar, Delavane Diaz, Frances C. Moore, Ulrich K. Müller, Richard J. Plevin, Adrian E. Raftery, Hana Ševčíková, Hannah Sheets, James H. Stock, Tammy Tan, Mark Watson, Tony E. Wong, and David Anthoff.** 2022. “Comprehensive evidence implies a higher social cost of CO₂.” *Nature*, 609(7927): 687–692.
- Ryan, Nicholas.** 2021. “The Competitive Effects of Transmission Infrastructure in the Indian Electricity Market.” *American Economic Journal: Microeconomics*, 13(2): 202–242.
- Ryan, Stephen P.** 2012. “The costs of environmental regulation in a concentrated industry.” *Econometrica*, 80(3): 1019–1061.
- Savion.** 2022. “Savion Completes Agreement to Sell 200 MW Ohio Solar Energy Project.” <https://savionenergy.com/savion-completes-agreement-to-sell-200-mw-ohio-solar-energy-project/>, Accessed August 12, 2026.
- Seel, Joachim, Julie Mulvaney Kemp, Anna Cheyette, Dev Millstein, Will Gorman, Seongeun Jeong, Dana Robson, Rachman Setiawan, and Mark Bolinger.** 2024. “Utility-Scale Solar, 2024 Edition: Empirical Trends in Deployment, Technology, Cost, Performance, PPA Pricing, and Value in the United States.” Lawrence Berkeley National Laboratory.
- Shoemaker, Jason.** 2021. “Interconnection Queue Status Update.” <https://www.pjm.com/-/media/committees-groups/committees/pc/2021/20210831/20210831-item-07-queue-status-update.ashx>.
- Verdier, Valentin, and Carson Reeling.** 2022. “Welfare effects of dynamic matching: An empirical analysis.” *The Review of Economic Studies*, 89(2): 1008–1037.

- 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.
- Weintraub, Gabriel Y., C. Lanier Benkard, and Benjamin Van Roy.** 2008. “Markov Perfect Industry Dynamics With Many Firms.” *Econometrica*, 76(6): 1375–1411.
- Weintraub, Gabriel Y, C Lanier Benkard, Przemysław Jeziorski, and Benjamin Van Roy.** 2010. “Nonstationary oblivious equilibrium.” Working Paper.
- Wiser, Ryan H., Dev Millstein, Ben Hoen, Mark Bolinger, Will Gorman, Joseph Rand, Galen L. Barbose, Anna Cheyette, Naïm R. Darghouth, Seongeun Jeong, Julie Mulvaney Kemp, Eric O’Shaughnessy, Bentham Paulos, and Joachim Seel.** 2024. “Land-Based Wind Market Report: 2024 Edition.” Lawrence Berkeley National Laboratory.
- Wolak, Frank A.** 2015. “Measuring the competitiveness benefits of a transmission investment policy: The case of the Alberta electricity market.” *Energy Policy*, 85: 426–444.
- Yang, Chenyu.** 2020. “Vertical structure and innovation: A study of the SoC and smartphone industries.” *The RAND Journal of Economics*, 51(3): 739–785.

A Model Details

We first characterize optimal strategies in Section 5.1 as cutoff rules, derive the law of motion for project masses used in the equilibrium definition in Section 5.3, and prove equilibrium existence.

Notation. We focus on a particular project and suppress the project-specific subscript. Let x denote time-invariant project characteristics, e the quarter in which the potential project arrives and makes its entry decision, c the interconnection cost, and z study information. A project that enters joins the queue in $e + 1$. Let η denote accumulated payoff information: $\eta = \xi^{S1}$ in Stage 1 and $\eta = \xi^{S1} + \xi^{S2}$ in Stage 2, replacing ξ in the arguments of $\bar{V}$, W^{S1} and W^{S2} in (2), (3) and (4). The exogenous probabilities of exits, $1 - \delta^{S1}$ and $1 - \delta^{S2}$, are based on x , and we suppress their i subscript too. We also define the queue statuses $\mathcal{K}_t(e, x)$ as a finite vector of variables that determine study-arrival probabilities and the distribution of the next study's observed state (c, z) , conditional on the current observed state and (x, e) . The queue status does not affect the transition of η : a new second study adds the payoff information ξ^{S2} to ξ^{S1} , whether that study is terminal or moves the project into Stage 2. In the empirical model, $\mathcal{K}_t(e, x)$ consists of the number of queued projects with weakly higher priority, the number of such projects in the same TO, capacity of queued projects in the same TO, and accumulated capacity of each fuel type up to t . Finally, we assume that (x, e, c, z, t) have finite supports, the accumulated payoff state η has continuous support, and the queue-status $\mathcal{K}$ lies in bounded subsets of Euclidean space.

A.1 Cutoff Strategy

Define the *waiting surplus* $\mathcal{U}_t^s(x, e, c, z, \eta)$ as the value of waiting in each stage $s \in \{1, 2\}$ (waiting with the first or second study), so that the continuation values are $W_t^s(x, e, c, z, \eta) = \max\{\mathcal{U}_t^s(x, e, c, z, \eta), 0\}$. The waiting values in the two stages $\mathcal{U}_t^s$ are

$$\begin{aligned} \mathcal{U}_t^{S2}(x, e, c, z, \eta) = & -\gamma^{S2} + \beta\delta^{S2} \left[\sum_{c'} q^{S2}(c'|x, e, t, c, z) \bar{V}(x, e, t, c', \eta) \right. \\ & \left. + (1 - Q^{S2}(x, e, t, c, z)) W_{t+1}^{S2}(x, e, c, z, \eta) \right], \end{aligned} \quad (\text{A.1})$$

$$\begin{aligned}
 \mathcal{U}_t^{S1}(x, e, c, z, \eta) = & -\gamma^{S1} + \beta\delta^{S1} \left[\sum_{c'} q^{S1,f}(c'|x, e, t, c, z) E_{\xi^{S2}} [\bar{V}(x, e, t, c', \eta + \xi^{S2})] \right. \\
 & + \sum_{c', z'} q^{S1}(c', z'|x, e, t, c, z) E_{\xi^{S2}} [W_{t+1}^{S2}(x, e, c', z', \eta + \xi^{S2})] \\
 & \left. + (1 - Q^{S1}(x, e, t, c, z)) W_{t+1}^{S1}(x, e, c, z, \eta) \right], \tag{A.2}
 \end{aligned}$$

In these equations, the q^s and Q^s objects give the probabilities of study arrival during period t . If no study arrives, the project carries its state into period $t + 1$. We characterize the project strategies under the following assumptions:

Assumption 1 (Regularity). Waiting costs satisfy $\gamma^{S1}, \gamma^{S2} > 0$. The unobserved heterogeneity ξ^{S1} , ξ^{S2} , and ξ^f have absolutely continuous distributions with full support. The probabilities of surviving exogenous departures are strictly positive, $\delta^s > 0$.

We assume that waiting costs are positive because projects pay for the costs of site control and study management. We also assume that ξ^{S1} , ξ^{S2} , and ξ^f are normally distributed in the model in Section 5. Empirically, we find that the δ^s values are close to 1. Under these assumptions, we first show that $\mathcal{U}_t^s$ has the following properties:

Lemma 1. For fixed (x, e, t, c, z) , $t < T$, and s , $\mathcal{U}_t^s$ is continuous and strictly increasing in η , and there exists a unique value of η such that $\mathcal{U}_t^s = 0$.

Proof. Backward induction on t : the terminal payoff $\bar{V}(x, e, t, c, \eta) = E_{\xi^f} [\max(\pi_i(t) + \xi^f, 0)]$ is continuous and strictly increasing in η under Assumption 1. In stage 2, $\sum_{c'} q^{S2} \bar{V}$ is a weighted sum of continuous, strictly increasing functions of η with nonnegative weights, and $(1 - Q^{S2}) W_{t+1}^{S2}$ is continuous and weakly increasing by induction (where $W_T^{S2} = 0$), which together imply that $\mathcal{U}_t^{S2}$ is continuous and strictly increasing. The same reasoning shows that $\mathcal{U}_t^{S1}$ also has these properties. Then, to show that there exists a unique η for $\mathcal{U}_t^s = 0$ given s , it suffices to find values of η at which $\mathcal{U}_t^s$ is negative and positive, respectively. Note that when $\eta \rightarrow -\infty$, $\bar{V}_t \rightarrow 0$, and by induction, $W_t^{S2} \rightarrow 0$ for any $t < T$, which implies that there must exist a sufficiently negative η such that $\mathcal{U}_t^{S2} < 0$ for positive γ^{S2} . Similar reasoning shows that there must exist a sufficiently large η for $\mathcal{U}_t^{S2} > 0$. The same argument applies for $\mathcal{U}_t^{S1}$. Strict monotonicity implies the uniqueness of the zero. $\square$

The lemma thus shows that the optimal waiting decision is characterized by a cutoff

$$\eta^{s*}(x, e, t, c, z) = \inf\{\eta : \mathcal{U}_t^s(x, e, c, z, \eta) \geq 0\}, \quad (\text{A.3})$$

where a project continues if $\eta \geq \eta^{s*}$ and withdraws otherwise.

We further note that the project strategies can be simplified to a “bang-bang” solution, where projects either exit right after learning about new information from a study, or do not endogenously exit. This strategy is consistent with the empirical pattern we observe in the data. The following assumption implies such a strategy:

Assumption 2 (Waiting value increases over time). Given η and $s \in \{1, 2\}$, $\mathcal{U}_{t+1}^s(x, e, c, z, \eta) \geq \mathcal{U}_t^s(x, e, c, z, \eta)$.

We can prove this condition based on two more primitive assumptions: (1) study-arrival probabilities increase over time, which is consistent with improvement of queue positions as they wait longer in equilibrium, and (2) that the dispersions of ξ^{S2} and ξ^f increase with waiting time, raising the option value of remaining.

A.2 State Transitions

In the above, we have described the project strategies given their beliefs, and those beliefs are subject to the restriction in Assumption 2. This section describes a model of how PJM processes studies, which determines a project’s state transition probabilities. These probabilities must be consistent with project beliefs in an equilibrium.

A.2.1 Transitions of η

We first describe how the distribution of accumulated payoff information, η , changes over time. To keep the notation tractable, we explicitly use the parametric assumptions on ξ^{S1} and ξ^{S2} in Section 5. The innovation in payoff information is independent of queue statuses and of the observed study-state transition. The only transition in accumulated payoff information occurs when $\eta = \xi^{S1}$ in Stage 1 to $\eta' = \xi^{S1} + \xi^{S2}$ upon the second study’s arrival. We assume ξ^{S1} has standard deviation σ^{S1} , and ξ^{S2} has standard deviation $\sigma^{S2} \cdot (t - e)^{\kappa^{S2}}$, where $\kappa^{S2} > 0$.

After a non-terminal first study in period t , the improper density of $\eta \equiv \xi^{S1}$ among

projects who do not withdraw is

$$g_t^{S1}(\eta \mid x, e, c, z) = \frac{1}{\sigma^{S1}} \phi\left(\frac{\eta}{\sigma^{S1}}\right) \mathbf{1}\{\eta \geq \eta^{1*}(x, e, t, c, z)\}, \quad (\text{A.4})$$

where ϕ is the standard normal density. After a non-terminal second study in period t , the improper density of $\eta' = \eta + \xi^{S2}$ among projects who continue into Stage 2 in $t + 1$ and have waited $t - e$ periods is

$$g_t^{S2}(\eta' \mid \eta, x, e, c', z') = \frac{1}{\sigma^{S2} \cdot (t-e)^{\kappa^{S2}}} \phi\left(\frac{\eta' - \eta}{\sigma^{S2} \cdot (t-e)^{\kappa^{S2}}}\right) \mathbf{1}\{\eta' \geq \eta^{2*}(x, e, t + 1, c', z')\}. \quad (\text{A.5})$$

The normal density component in (A.5) does not depend on (c, z) , (c', z') , or $\mathcal{K}_t(e, x)$; these variables affect the mass of continuing projects only through the cutoff.

A.2.2 Transitions of c and z and Study Arrival Probabilities

We assume that the probabilities of (1) receiving a new terminal second study with cost c' in Stage 1, (2) receiving a new non-terminal second study with cost c' and updated study information z' in Stage 1, and (3) receiving the third study with cost information c' in Stage 2 are functions of $(x, e, t, c, z, \mathcal{K}_t(e, x))$, and are smooth in $\mathcal{K}_t(e, x)$. Thus Q^{S1} and Q^{S2} depend on the observed state and queue status, but not on η . In practice, we use an ordered-probit model for the transition of costs, and probit models for the transitions of other states (Appendix D).

A.3 Transitions of the Queue

We now describe, given how PJM processes studies and project strategies, how the composition of the queue changes over time. Let $m_t^s(x, e, c, z, \eta)$ denote the mass density of projects in stage $s \in \{1, 2\}$ at the start of quarter t , after the most recent withdrawal decision, where its integral with respect to η produces the mass of projects given (x, e, c, z) . Let $N_e(x)$ denote the exogenous mass of potential entrants in quarter e , and $q_0(c, z \mid x, e)$ the probability of a non-terminal first study with state (c, z) , so that $\sum_{c,z} q_0 = 1 - q^f$. The entry probability is

$$p^{\text{ent}}(x, e) = \Pr(EV(x, e) > C(x, e)), \quad (\text{A.6})$$

where C has a beta distribution specified in Section 6.3.

The Stage-1 mass of projects after entry is not exogenous: it equals the mass of potential projects multiplied by the entry probability, the probability of a non-terminal first-study state and the density of η for not withdrawing. A project arriving and choosing to enter in e joins the queue and receives its first study in $e + 1$. Thus,

$$m_{e+1}^{S1}(x, e, c, z, \eta) = N_e(x) p^{\text{ent}}(x, e) q_0(c, z \mid x, e) g_{e+1}^{S1}(\eta \mid x, e, c, z). \quad (\text{A.7})$$

To describe the changes of the queue mass over time, we require notation on the transition probabilities of (c, z) defined in Appendix A.2.2. For notational simplification, we slightly abuse notation and reuse q and Q from project beliefs but condition explicitly on queue statuses $\mathcal{K}_t(e, x)$.

The Stage-1 mass transition when $t \geq e + 1$ is:

$$m_{t+1}^{S1}(x, e, c, z, \eta) = \delta^{S1} [1 - Q^{S1}(x, e, t, c, z; \mathcal{K}_t(e, x))] m_t^{S1}(x, e, c, z, \eta). \quad (\text{A.8})$$

The projects leaving the queue include transitions to Stage 2 and exogenous exits at a per-period probability of $1 - \delta^{S1}$. For the projects receiving a terminal second study, the completed capacity in a period is the mass in Stage 1 receiving the terminal second study times the respective capacity and the probability of completion.

The Stage-2 mass consists of continuing Stage-2 projects and Stage-1 projects receiving a non-terminal second study:

$$\begin{aligned} m_{t+1}^{S2}(x, e, c', z', \eta') &= \delta^{S2} [1 - Q^{S2}(x, e, t, c', z'; \mathcal{K}_t(e, x))] \\ &\quad \times m_t^{S2}(x, e, c', z', \eta') \\ &\quad + \sum_{c, z} \int \delta^{S1} q^{S1}(c', z' \mid x, e, t, c, z; \mathcal{K}_t(e, x)) \\ &\quad \times g_t^{S2}(\eta' \mid \eta, x, e, c', z') m_t^{S1}(x, e, c, z, \eta) d\eta. \end{aligned} \quad (\text{A.9})$$

The leaving mass includes projects receiving the third study and exogenous exits at a per-period probability of $1 - \delta^{S2}$. For the projects receiving the third study, the completed capacity in a period is the mass in Stage 2 receiving the third study times the respective capacity and the probability of completion.

A.4 Queue Status and Operating Revenues

We now construct queue status variables based on the mass density functions defined above. The queue position in period t for projects that arrived in period e and joined the queue in $e + 1$ is the number of queued projects with higher-or-equal priority, i.e., those with arrival quarter $\tilde{e} \leq e$:

$$\sum_s \sum_{\tilde{x}} \sum_{\tilde{e} \leq e} \sum_{c,z} \int m_t^s(\tilde{x}, \tilde{e}, c, z, \eta) d\eta, \quad (\text{A.10})$$

Other queue statuses used in our empirical model defined in Section 5.3 are computed similarly. For example, the same-TO queue position restricts the sum above to projects in the same TO zone, same-TO queued capacity weights each project by capacity, and accumulated completed capacity sums completion flows by fuel type. We keep the notation focused on queue position because the other components are finite sums of the same mass and completion objects.

The wholesale price in t is determined by the queue status of accumulated capacity of each fuel type up to t . The accumulated completion adds to the supply curve in Appendix B, and its crossing with the exogenous demand yields the new wholesale price. A project's revenue, $\tilde{r}$, is computed as the discounted sum of wholesale price times the project capacity, adjusted for fuel-type efficiency, over a 25-year service life. We assume that the resulting $\tilde{r}$ is continuous in the queue status. In the actual implementation in Appendix B, we construct a step function based on the fuel types. For equilibrium existence purposes, the function can be smoothed at the break points to ensure existence. We use the step function directly in equilibrium calculation through an iterative algorithm. We do not find non-existence in practice.

A.5 Equilibrium Existence

A **non-stationary queuing equilibrium** is defined as the collection of cutoff strategies η^{s*} , entry probabilities p^{ent} , queue statuses $\{\mathcal{K}_t(e, x)\}_{t,e,x}$ and revenues $\tilde{r}$ that satisfy Equations (A.3), (A.6) and the conditions in Appendix A.4. Given the queue statuses and the resulting project revenue, backward induction yields the cutoffs and determine entry probabilities. Then, (A.7)–(A.10) produce new queue statuses. Call this mapping Γ . In the arguments below, we show that Γ maps a compact set into itself and is continuous. Existence thus follows from Brouwer's fixed point theorem.

First, consider the queue-position component. It lies in $[0, \bar{N}]$, where $\bar{N} = \sum_{e,x} N_e(x)$.

Other components of $\mathcal{K}$ are also bounded because they are finite sums of project counts, project capacity, or completed capacity over finite supports. Let $\bar{K}_j$ denote a finite upper bound for component j of $\mathcal{K}$. The domain $\mathcal{D} = \prod_j [0, \bar{K}_j]$ is compact and convex, and Γ maps $\mathcal{D}$ into itself.

Second, value functions are continuous in queue statuses by backward induction over continuous study-arrival probabilities and $\max\{\cdot, 0\}$. Entry probabilities are continuous in queue statuses via (A.6) and the continuity of value functions. The project revenue, $\tilde{r}$ in the payoff function, is based on wholesale prices, which continuously depend on $\mathcal{K}$ through accumulated completed capacity, and this function is continuous. As a result, cutoffs are continuous in queue statuses given the continuity in $\mathcal{U}_t^s$. The primitive transition of η is independent of queue statuses; the post-withdrawal densities in (A.4)–(A.5) are continuous in queue statuses only through the cutoffs. The transition probabilities of (c, z) in Appendix A.2.2 are continuous in queue statuses by assumption. Therefore, the mass transitions (A.7)–(A.9) are continuous in queue statuses. The mapping Γ is thus continuous. Given that Γ is a continuous mapping from a compact set into itself, Brouwer’s fixed point theorem implies the existence of the equilibrium.

B Construction of Revenue from Power Production

This section details how we construct the present value of wholesale electricity revenues for projects of different fuel types. For estimation, each project’s revenue is constructed using prices generated by the hourly supply curve and generation implied by generation fuel-specific capacity factors.⁴⁴ We construct the hourly supply curve for 2011 through 2021 and use it to price the corresponding 44 quarters. For quarters beyond this horizon, we repeat the last modeled year’s prices and capacity factors by calendar quarter. Present value is calculated as the sum of revenues over the 25 years beginning with the project’s in-service quarter, discounted at an annual rate of 5%.⁴⁵

For simulation, the mix of in-service projects would change, leading to different wholesale

⁴⁴In effect, we assume that the projects believe the power demand to be flat as in prior years, which implies that the surge in power demand driven by AI and data centers in later years would be a surprise to projects.

⁴⁵Revenue from power production excludes capacity market payments and tax credits. The fuel-type coefficients in the project payoff functions absorb capacity market payments in estimation, and Section 8.1 bounds the decline in capacity market prices that counterfactual capacity additions would induce. The production and investment tax credits enter through the construction cost ω_{it} (Section 6.2).

electricity prices. We assume that demand is perfectly inelastic, and we build each counterfactual supply curve from two pieces: the existing generators, taken as observed in the data with their entries and exits held fixed, and the new projects that complete in the queuing equilibrium model. Crossing the observed quantity (from the perfectly inelastic demand) with the equilibrium supply yields the equilibrium price. In counterfactuals, additional projects complete the queue, and their capacity adds to the supply curve, which changes the equilibrium price.

B.1 Supply Curve

We construct hourly supply curves for the PJM Interconnection from 2011 through 2021. The construction follows Cicala (2022) and Hausman (2025), with modifications for PJM-specific data availability.

The supply curve is based on a merit-order representation of competitive project offers, which we approximate by marginal costs. We rank dispatchable generators by estimated marginal cost and construct the hourly supply curve from their available capacity. We use this merit-order approach because it allows us to evaluate counterfactual capacity additions and produces normalized LMPs that closely track actual PJM prices. Building the curve for each hour requires two unit-level inputs, an estimated marginal cost and an available capacity; the next two subsections describe how we construct each, and the subsequent subsections cover demand, market clearing, and the use of the curve in counterfactuals.

The primary data sources are EIA Forms 860, and 923, which together provide generator-level capacity and sub-fuel characteristics, net generation, heat input, and fuel-price inputs; EIA operating-expense data on fuel-specific variable operation and maintenance costs (O&M costs); EIA Henry Hub and WTI spot prices; EPA Continuous Emissions Monitoring System (CEMS) data and Power Sector Programs Progress Reports on hourly unit-level emissions and heat input; PJM Data Miner data on hourly load and generation; NRC Daily Power Status reports on daily nuclear reactor operating status; S&P Global emissions allowance prices; NREL solar radiation data; PRISM county-level weather records; and U.S. Census geocode files.

We combine these sources by harmonizing generator, plant, and boiler identifiers across EIA and EPA records, including mapping fossil-plant boilers to generators using EIA linkages

and maintaining these mappings over time. We also map facilities to counties and time zones and convert hourly observations to a common time standard before merging them with PJM operations data.

Throughout this appendix, a “unit” refers to an individual generating unit represented separately in the merit-order supply curve and generally corresponds to a generator reported on EIA forms. A “plant” refer to a physical site that may contain multiple generating units.

B.1.1 Unit-level Marginal Cost

We estimate an hourly marginal cost for each dispatchable unit. Let j index dispatchable generating units and h index hours. Throughout this appendix, t indexes quarters. For fossil units, marginal cost is the sum of fuel costs, emissions allowance costs, and variable O&M costs:

$$MC_{jh} = \text{FuelCost}_{jh} + \text{EmissionsCost}_{jh} + \text{VOM}_{jh}, \quad (\text{B.11})$$

where FuelCost_{jh} is the fuel cost of producing one MWh, $\text{EmissionsCost}_{jh}$ is the emissions allowance cost per MWh, and VOM_{jh} is variable O&M costs.

Each unit’s fuel cost is determined by the unit’s heat rate, sub-fuel mix, and corresponding sub-fuel prices:

$$\text{FuelCost}_{jh} = \text{HR}_{jh} \sum_k \text{Share}_{jh}^k \text{FuelPrice}_{jh}^k. \quad (\text{B.12})$$

Here k indexes sub-fuel categories, Share_{jh}^k is the heat-input share of sub-fuel k , FuelPrice_{jh}^k is the corresponding sub-fuel price, and HR_{jh} is the unit’s heat rate. For natural gas and oil units, we use daily Henry Hub and WTI spot prices, adjusted by plant-year markups estimated from EIA fuel receipt data, to capture daily sub-fuel price variation not observed in the monthly EIA series. Missing plant-level fuel prices are filled using state-, division-, and national-level averages.

The heat rate HR_{jh} measures the amount of fuel energy required to produce one MWh of electricity. Multiplying the heat rate by the unit’s weighted-average sub-fuel price gives its fuel cost per MWh. For units monitored by the CEMS, we measure the monthly heat rate as reported heat input divided by net generation and winsorize it to a physically plausible range (roughly 6 to 100 MMBtu/MWh). Because a unit must carry a marginal cost even in months when it does not run, we fill missing and outlier months from unit-level regressions of the

heat rate on operating characteristics such as age and ramping, falling back to vintage- and size-specific averages where a unit has too little data of its own. Heat rates for units outside the CEMS are built from EIA-923 fuel consumption and generation by the same logic.

Emissions costs equal the unit’s emissions rate per MWh multiplied by the relevant allowance price. Units are assigned to SO₂ and NO_x programs based on location and program rules. For seasonal NO_x markets, emissions costs apply only during the ozone-season window for the unit’s county.

Variable O&M costs capture non-fuel and non-emissions expenses that vary with generation, including routine maintenance, other production-related expenses, and production-related materials. We assign each unit a fuel-type-year-specific variable O&M costs in dollars per MWh and add it to marginal cost as a flat per-MWh adder.

As a final step, fossil fuel-unit marginal costs are winsorized at the 1st to 99th percentiles within year to reduce the influence of extreme outliers.

For non-fossil units, we assign constant marginal costs within fuel type. Renewable resources, including wind, solar, hydro, and geothermal units, are assigned zero marginal cost. Nuclear units are assigned a constant marginal cost of \$17.13/MWh, based on the EIA average operating expense for nuclear generation.

B.1.2 Available Capacity for Generation

We next estimate the capacity available for generation from each unit in each hour. For fossil units, we first construct hourly net generation. For units monitored by CEMS, hourly net generation is calculated as reported gross load multiplied by the unit’s gross-to-net conversion factor, defined as the monthly ratio of net generation to reported gross load, with an estimated scale used when this ratio is missing or implausible. For units without CEMS coverage, monthly net generation from EIA-906 is allocated across hours of the month using the PJM-wide load. We then constrain each fossil unit’s hourly capacity by its observed daily maximum, preventing the merit order from dispatching capacity that the unit could not have produced on that day.

For non-fossil units, we use resource-specific availability measures. Wind capacity is taken hourly from PJM generation data. Solar generation is taken from PJM hourly data for 2019 onward. For earlier years, when PJM reports no hourly solar generation, we infer hourly PJM-

wide generation using plant-level monthly generation and irradiance data. For each plant, we allocate monthly generation across hours in proportion to irradiance at its location, then sum across plants to obtain PJM-wide hourly generation. Nuclear generation is constructed from daily NRC reactor capacity factors and spread uniformly across hours within each day. Hydroelectric, storage, biomass, and other minor fuel types are represented by a synthetic zero-cost unit with capacity equal to PJM’s reported hourly generation from these sources.

B.2 Demand

Hourly electricity demand is measured using PJM-wide metered load from 2011 through 2021. In the market-clearing exercise, demand is treated as perfectly inelastic within each hour.

B.3 Market-Clearing Price and Validation

For each hour, we sort available dispatchable units in ascending order of marginal cost and accumulate available capacity along the merit-order curve. The predicted market-clearing price, $\hat{P}_h$, is the marginal cost of the first unit for which cumulative available capacity equals or exceeds realized PJM-wide metered load.

Figure B.1 compares the quarterly average of $\hat{P}_h$ with the quarterly average of actual PJM real-time hourly LMPs from 2011 through 2021. The predicted series broadly tracks the actual LMP series.




Figure B.1: Predicted and Actual PJM Real-Time LMPs

B.4 Using the Supply Curve in the Structural Model

The structural model uses the constructed supply curve to predict how market prices respond to counterfactual capacity additions. To make these calculations computationally feasible, we approximate each quarter’s hourly demand distribution using 20 representative hours (Reguant, 2019). To define the representative hours, we cluster hours by net demand.⁴⁶ Within each quarter, we use k-means clustering to divide hours into 20 net-demand clusters and select the median-net-demand hour from each cluster as its representative. Each representative hour h receives a weight w_h , equal to the share of the quarter’s actual hours assigned to its cluster, so that the 20 weighted representative hours approximate the quarter’s full net-demand distribution. We write $\mathcal{C}_t$ for the set of representative hours in quarter t . Repeating this procedure for every quarter from 2011 through 2021 produces 880 representative hours. From here on, h indexes these representative hours, each of which identifies a unique quarter and clock hour; quarter-specific quantities therefore inherit their quarter from h and carry no separate t subscript.

The structural model determines the timing and amount of capacity completing the queueing process at the quarterly level. After a one-year construction lag, this capacity enters service and is added to the representative-hour supply curves, with renewable and non-renewable additions treated differently.

Wind and solar additions are assigned zero marginal cost, but their output depends on resource availability. Let g index fuel types. A nameplate capacity addition Cap_g has available generation of $\text{CF}_{h,g} \text{Cap}_g$ in representative hour h , where $\text{CF}_{h,g} \in [0, 1]$ denotes the fuel’s capacity factor in that representative hour. We estimate these wind and solar capacity factors $\text{CF}_{h,g}$ from observed PJM operations. For each quarter and each hour of the day, we average system-wide generation from g across all days in the quarter at that clock hour, and divide by that year’s installed nameplate capacity of g . This yields a diurnal capacity factor $\text{CF}_{h,g}$ for each hour and fuel-type g . As a result, wind and solar additions shift the supply curve more and depress the clearing price by more in representative hours when expected generating output is high, and less when output is low.

Non-renewable additions are treated as new natural-gas capacity. To allow new gas units

⁴⁶Net demand is defined as metered load minus available wind and solar generation, because variable renewable output shifts the residual demand that the non-renewable portion of the merit order must serve and that therefore pins down the clearing price.

to differ in efficiency, we divide the heat-rate distribution of natural-gas generators operating in PJM since 2011 into ten equal-width bins and weight each bin by its share of post-2011 new gas nameplate capacity. A counterfactual gas addition is then allocated across these bins according to those weights. This creates ten distinct additions to the supply curve, each with its own capacity and marginal cost. The marginal cost of each bin is computed using its mean heat rate, the prevailing annual natural-gas price, emissions cost, and variable operation and maintenance cost, following the same marginal-cost definition used for existing units. New gas capacity enters the supply curve at full nameplate capacity, while its dispatch is determined by the merit order. Its effective capacity factor is therefore endogenous.

This constructed supply curve may differ from that observed in the data because of retirements and other unmodeled capacity changes. To account for these differences, we compute a quarterly fuel-specific residual capacity wedge equal to the observed change in PJM nameplate capacity minus model-predicted completions. Adding this wedge to the merit order reconciles the change in model-implied capacity with the observed change in PJM capacity from one quarter to the next. In counterfactual exercises, we hold the wedge fixed, so the changes in predicted prices are driven only by changes in model-predicted capacity additions.

For each counterfactual capacity vector, we insert the resulting capacity additions and residuals into the existing merit-order supply curve, re-sort units by marginal cost, and determine the clearing price $\hat{P}_h$ in each representative hour by crossing the resulting supply curve with PJM demand. Each quarter therefore yields 20 representative-hour clearing prices, which Section B.5 combines with their weights w_h and project characteristics to construct revenue.

We also compute a single PJM-wide quarterly price, $\hat{P}_t = \sum_{h \in \mathcal{C}_t} w_h \hat{P}_h$, which weights representative hours by their frequency. This price does not enter project revenue. Instead, it serves as the convergence criterion when we solve for the queuing equilibrium: we iterate until the PJM-wide price path implied by projects' queue entry decisions reproduces the price path against which those decisions were made.

B.5 From Hourly Prices to Project-Level Revenue

We combine the representative-hour market clearing prices with each project's capacity and capacity factors, beginning in its in-service quarter, to construct project-level revenue. For each project, we calculate expected revenue from power production over a 25-year operating

horizon and discount future revenue at an annual rate of 5%. Projects differ in this revenue through their fuel type, in-service timing, and capacity.

Because the market clears at a single price in each representative hour, every dispatched unit earns $\hat{P}_h$ regardless of type. Fuel types nonetheless differ in their daily revenue per megawatt of nameplate capacity, because they differ in the hours they produce.

Consider first a single quarter, and let H_t denote the number of hours in quarter t . The revenue of a project of fuel g per megawatt of nameplate capacity in quarter t is the sum, over the quarter's hours, of its generation in each hour multiplied by that hour's clearing price. In terms of the representative hours, this revenue is

$$\text{Rev}_{g,t} = H_t \sum_{h \in C_t} w_h \text{CF}_{h,g} \hat{P}_h. \quad (\text{B.13})$$

In the above, the term $H_t w_h$ is how the number of type h hours in the quarter, and the weight is common to all fuels. The capacity factor $\text{CF}_{h,g}$ is how much energy a project of fuel g produces in this hour, and the capacity factor is specific to the fuel type. For wind and solar, $\text{CF}_{h,g}$ is the data-implied capacity factor; for new natural gas, it is the share of the hour's gas capacity that the merit order dispatches, which market clearing determines endogenously. Because wind, solar, and natural gas have different hourly capacity factors, they earn different revenue per megawatt in an hour.

A project of fuel g that goes into service in quarter t_0 earns this quarterly revenue over a 25-year life. Its discounted revenue per megawatt is

$$\sum_{n=0}^{99} \frac{\text{Rev}_{g,t_0+n}}{(1.05)^{n/4}}, \quad (\text{B.14})$$

rescaled to millions of dollars per megawatt. For quarters up to 2021, the first year of operation for projects completing in 2020, $\text{Rev}_{g,t}$ is based on that quarter's representative-hour prices and capacity factors in the data. For quarters after 2021, we carry forward the 2021 quarterly revenue for the corresponding calendar quarter, thereby preserving seasonal variation in prices and capacity factors. We map these discounted revenues to $\tilde{r}_{it}$ in Equation (1) of Section 5 based on completion time and fuel type. As described in Section 7, we compute the implied load (demand) in quarters after 2021, which, together with model predicted completions, generates equilibrium prices that match the observed prices in 2021.

C Identification of Payoffs and Waiting Costs

Fix project characteristics x and arrival quarter e , when the potential project makes its entry decision. We show below how to identify the payoff specific to each combination of (x, e) as well as distributions of unobserved heterogeneity and payoff shocks. We therefore suppress (x, e) from the arguments of all functions in this section, though e continues to appear where we write total elapsed time since arrival.

Notation. After a non-terminal first study with state (t, c, z) , define the probability that a project that waits receives a non-terminal second study with state (c', z') after having waited τ_1 quarters. Thus τ_1 is the elapsed time from this non-terminal first-study state to the next second study, whether the second study is terminal or non-terminal:

$$\Lambda^{S1}(c', z', \tau_1 \mid t, c, z) = \left[\prod_{h=0}^{\tau_1-1} \delta^{S1}(1 - Q^{S1}(t+h, c, z)) \right] \delta^{S1} q^{S1}(c', z' \mid t + \tau_1, c, z). \quad (C.15)$$

The analogous probability that the second study is terminal with interconnection cost c' is

$$\Lambda^{S1,f}(c', \tau_1 \mid t, c, z) = \left[\prod_{h=0}^{\tau_1-1} \delta^{S1}(1 - Q^{S1}(t+h, c, z)) \right] \delta^{S1} q^{S1,f}(c' \mid t + \tau_1, c, z). \quad (C.16)$$

After a non-terminal second study arrives at $\tilde{t} = t + \tau_1$, the project enters Stage 2 in $\tilde{t} + 1$ with state $(\tilde{t} + 1, c', z')$. Define the probability that a project that waits receives a terminal study with interconnection cost c'' after τ_2 additional no-arrival quarters, in period $\tilde{t} + 1 + \tau_2$:

$$\Lambda^{S2}(c'', \tau_2 \mid \tilde{t}, c', z') = \left[\prod_{h=0}^{\tau_2-1} \delta^{S2}(1 - Q^{S2}(\tilde{t} + 1 + h, c', z')) \right] \delta^{S2} q^{S2}(c'' \mid \tilde{t} + 1 + \tau_2, c', z'). \quad (C.17)$$

We write F^{S1} for the distribution of ξ^{S1} and condition the distributions F^{S2} and G on total elapsed time since arrival. Thus a second study at $\tilde{t} = t + \tau_1$ draws ξ^{S2} from $F^{S2}(\cdot \mid \tilde{t} - e)$ and, if terminal, ξ^f from $G(\cdot \mid \tilde{t} - e)$; a terminal third study in $\tilde{t} + 1 + \tau_2$ draws ξ^f from $G(\cdot \mid \tilde{t} + 1 + \tau_2 - e)$. We impose the normalizations $E[\xi^{S1}] = 0$, $E[\xi^{S2} \mid \text{total elapsed time}] = 0$, and $E[\xi^f \mid \text{total elapsed time}] = 0$. We use a_t to denote the present value of the profit from starting to operate the plant in t less the construction cost, which corresponds to $r_{it} - \omega_{it}$ in Equation (1).

Theorem 1. Under the conditions in Appendix A that yield cutoff strategies and ensure that a project continuing after a study does not withdraw before the next study, and the assumptions that (1) the probabilities defined in (C.15), (C.16), and (C.17), as well as the component transition objects Q^s , δ^s , and the discount factor β , are known and positive, (2) F^{S1} and the elapsed-time-indexed families F^{S2} and G are absolutely continuous with positive densities, have full support on $\mathbb{R}$, and have finite means, (3) conditional on the lagged observed state and total elapsed time since arrival, each new observed study-state realization is independent of the payoff shocks, the shocks ξ^{S1} , ξ^{S2} , and ξ^f are mutually independent, and each new study-cost realization has full support on $\mathbb{R}$,⁴⁷ (4) the newly realized study cost affects the waiting value only through the distribution of the next study's cost, and that distribution is continuously differentiable and strictly increasing in the realized cost in the sense of first-order stochastic dominance,⁴⁸ and the next study's cost diverges in probability as the realized cost tends to $\pm\infty$, with its expected negative part converging to zero as the realized cost tends to $+\infty$,⁴⁹ and (5) every total elapsed time reached with positive probability by a terminal second study after a non-terminal first study is also reached with positive probability by a terminal third study, the objects F^{S1} , F^{S2} , G , γ^{S1} , γ^{S2} , and a_t are identified at the study dates and total elapsed times reached with positive probability after a non-terminal first study. The payoff parameters θ are then identified whenever the matrix collecting $\chi_t(x_i)$ over these dates and project characteristics has full column rank.

Proof. All probabilities below condition on the lagged observed state and total elapsed time since arrival. By assumption, varying the newly realized study cost changes the continuation cutoff but not the distribution of the payoff shock at that decision node. Because a project that continues after a study does not withdraw before the next one, the Λ 's give the entire probability of surviving the wait and receiving the next study, and they do not depend on the

⁴⁷The empirical model in Section 5 assumes that c 's support is finite to simplify the notation for transition probabilities. Our characterization of strategies and equilibrium existence still goes through when c has full support and its conditional distributions are absolutely continuous.

⁴⁸In the estimated transition model, the most recent study cost enters only the ordered probit for the next study's cost and not the study-arrival probit, and its coefficients are increasing across prior-cost bins (Appendix Table D.1).

⁴⁹Precisely, let $\Pi(\cdot | c)$ denote the conditional distribution of the next study's cost given a realized cost c , with the other conditioning variables held fixed. For every fixed $k \in \mathbb{R}$, we require $\Pi(k | c) \rightarrow 0$ as $c \rightarrow +\infty$ and $\Pi(k | c) \rightarrow 1$ as $c \rightarrow -\infty$, together with $\int \max\{-c', 0\} d\Pi(c' | c) \rightarrow 0$ as $c \rightarrow +\infty$. The last condition rules out a rare but sufficiently heavy left tail that would keep the expected completion value away from zero. A location-shift specification in which the next cost equals a strictly increasing, unbounded function of c plus an independent error with a finite mean satisfies these conditions.

payoff shocks. Since the Λ 's are known and positive, we work with probabilities normalized by the relevant Λ ; the corresponding unnormalized probabilities are observed in the data.

Cutoff Regularity. Assumption (4) and backward induction from $W_T^s = 0$ imply that each waiting value $\mathcal{U}_t^s$ in Appendix A is continuous and strictly decreasing in the current study cost. For every fixed η , the tail conditions make $\mathcal{U}_t^s$ positive as that cost tends to $-\infty$ and negative as it tends to $+\infty$. Since $\mathcal{U}_t^s$ is continuous and strictly increasing in η , each cutoff is therefore continuous and strictly increasing in the current cost and ranges over $\mathbb{R}$. Steps 1–3 use these properties.

Step 1: Terminal Shock. Consider projects that have received a non-terminal second study in $\tilde{t}$ and enter Stage 2 with state $(\tilde{t}+1, c', z')$. Let $\eta = \xi^{S1} + \xi^{S2}$, and let H^{S2} denote the distribution of η at this decision node before the withdrawal decision, which by Assumption (3) does not vary with (c', z') . By Appendix A, the Stage-2 decision is characterized by a cutoff $\eta^{2*}(\tilde{t}+1, c', z')$, where projects continue if $\eta \geq \eta^{2*}(\tilde{t}+1, c', z')$.

The observed withdrawal probability upon entering Stage 2 after this non-terminal second study is

$$D^{S2}(\tilde{t}+1, c', z') = H^{S2}(\eta^{2*}(\tilde{t}+1, c', z')).$$

Next consider the observed joint probability, from the same non-terminal second-study state, of continuing into Stage 2, receiving a terminal third study in period $\tilde{t}+1+\tau_2$ with cost c'' , and withdrawing at the terminal study. This probability is

$$\Omega^{S2}(\tilde{t}, c', z', c'', \tau_2) = \Lambda^{S2}(c'', \tau_2 \mid \tilde{t}, c', z') M^{S2}(\tilde{t}, c', z', c'', \tau_2).$$

Since Λ^{S2} is known and positive, $M^{S2} = \Omega^{S2}/\Lambda^{S2}$ is known from data. It equals

$$M^{S2}(\tilde{t}, c', z', c'', \tau_2) = \int_{\eta^{2*}(\tilde{t}+1, c', z')}^{\infty} G(c'' - a_{\tilde{t}+1+\tau_2} - \tilde{\eta} \mid \tilde{t}+1+\tau_2 - e) dH^{S2}(\tilde{\eta}).$$

As c' varies, cutoff regularity makes D^{S2} continuous and strictly increasing from 0 to 1. Holding the remaining arguments fixed, we may therefore index M^{S2} by the observed probability

D^{S2} . Differentiating with respect to this probability gives

$$-\frac{dM^{S2}(\tilde{t}, c', z', c'', \tau_2)}{dD^{S2}(\tilde{t} + 1, c', z')} = G(c'' - a_{\tilde{t}+1+\tau_2} - \eta^{2*}(\tilde{t} + 1, c', z') \mid \tilde{t} + 1 + \tau_2 - e). \quad (C.18)$$

The left-hand side is known from data. As c'' varies over its support, the right-hand side traces $G(\cdot \mid \tilde{t} + 1 + \tau_2 - e)$ with a horizontal shift. Its mean-zero normalization identifies both this distribution and

$$b_{\tilde{t}+1+\tau_2}^{S2}(\tilde{t} + 1, c', z') \equiv a_{\tilde{t}+1+\tau_2} + \eta^{2*}(\tilde{t} + 1, c', z').$$

Repeating this argument across third-study arrival dates identifies the elapsed-time-indexed family G . By Assumption (5), this family also contains every distribution used when a second study is terminal.

Step 2: Second-Study Payoff Information. Now consider projects that have reached a non-terminal first study with state (t, c, z) . Let $\eta^{1*}(t, c, z)$ denote the Stage-1 cutoff. Entry precedes the realization of ξ^{S1} in (5), and ξ^{S1} is independent of the first-study state by Assumption (3), so its distribution at this node is F^{S1} and the immediate withdrawal probability is

$$D^{S1}(t, c, z) = F^{S1}(\eta^{1*}(t, c, z)).$$

Suppose a project that continues receives a non-terminal second study after τ_1 quarters, so that $\tilde{t} = t + \tau_1$ and it enters the new state $(c', z', \tilde{t} + 1)$ in the following period. The normalized joint probability of continuing after the first study and withdrawing upon entering Stage 2 is

$$M^{S1}(t, c, z, c', z', \tau_1) = \int_{\eta^{1*}(t, c, z)}^{\infty} F^{S2}(\eta^{2*}(\tilde{t} + 1, c', z') - \xi^{S1} \mid \tilde{t} - e) dF^{S1}(\xi^{S1}).$$

As in Step 1, cutoff regularity lets us index M^{S1} by the observed probability D^{S1} . Differentiating with respect to this probability gives

$$-\frac{dM^{S1}(t, c, z, c', z', \tau_1)}{dD^{S1}(t, c, z)} = F^{S2}(\eta^{2*}(\tilde{t} + 1, c', z') - \eta^{1*}(t, c, z) \mid \tilde{t} - e). \quad (C.19)$$

For any fixed τ_2 reached with positive probability, rewrite the argument on the right-hand side as

$$\eta^{2*}(\tilde{t} + 1, c', z') - \eta^{1*}(t, c, z) = b_{\tilde{t}+1+\tau_2}^{S2}(\tilde{t} + 1, c', z') - b_{\tilde{t}+1+\tau_2}^{S1}(t, c, z),$$

where

$$b_{\tilde{t}+1+\tau_2}^{S1}(t, c, z) \equiv a_{\tilde{t}+1+\tau_2} + \eta^{1*}(t, c, z).$$

Since $b_{\tilde{t}+1+\tau_2}^{S2}$ is identified from Step 1 and ranges over $\mathbb{R}$ through variation in c' , (C.19) identifies $F^{S2}(\cdot \mid \tilde{t} - e)$ and the shift $b_{\tilde{t}+1+\tau_2}^{S1}(t, c, z)$, using the mean-zero normalization of $F^{S2}(\cdot \mid \tilde{t} - e)$.

Step 3: Payoff Levels and Cutoffs. Using the identified object $b_{\tilde{t}+1+\tau_2}^{S1}$, the first-study withdrawal probability can be written as

$$D^{S1}(t, c, z) = F^{S1}\left(b_{\tilde{t}+1+\tau_2}^{S1}(t, c, z) - a_{\tilde{t}+1+\tau_2}\right). \quad (\text{C.20})$$

Variation in the first-study cost makes $b_{\tilde{t}+1+\tau_2}^{S1}$ range over $\mathbb{R}$ and thus traces out the CDF F^{S1} with a horizontal shift. Its mean-zero normalization identifies $a_{\tilde{t}+1+\tau_2}$ and F^{S1} . Repeating the argument over third-study dates, together with Assumption (5), identifies a_t at every terminal-study date covered by the theorem. The identified a_t then separately identifies the cutoffs $\eta^{1*}(t, c, z)$ and $\eta^{2*}(\tilde{t} + 1, c', z')$ from the identified shifts b^{S1} and b^{S2} . Finally, Equation (1) implies that $a_t - \tilde{r}_{it} + \omega_{it} = \chi_t(x_i)\theta$. Since the left-hand side is identified for each (x_i, t) , θ is identified when the resulting matrix of $\chi_t(x_i)$ has full column rank.

Step 4: Waiting Costs. At the Stage-2 cutoff in $\tilde{t} + 1$, the project is indifferent between waiting and withdrawing. By the waiting-value monotonicity conditions in Appendix A, a project that continues at this cutoff does not withdraw before the next study. Let

$$S^{S2}(\tau_2 \mid \tilde{t} + 1, c', z') = \prod_{h=0}^{\tau_2-1} \delta^{S2}(1 - Q^{S2}(\tilde{t} + 1 + h, c', z'))$$

denote the probability of still waiting through τ_2 Stage-2 periods before the next terminal-study arrival. The indifference condition implies

$$0 = A^{S2}(\tilde{t} + 1, c', z') - \gamma^{S2} B^{S2}(\tilde{t} + 1, c', z'),$$

where

$$A^{S^2}(\tilde{t} + 1, c', z') = \sum_{c'', \tau_2} \beta^{\tau_2+1} \Lambda^{S^2}(c'', \tau_2 \mid \tilde{t}, c', z') \\ \times \int \max\{a_{\tilde{t}+1+\tau_2} + \eta^{2*}(\tilde{t} + 1, c', z') \\ + v - c'', 0\} dG(v \mid \tilde{t} + 1 + \tau_2 - e)$$

and

$$B^{S^2}(\tilde{t} + 1, c', z') = \sum_{\tau_2} \beta^{\tau_2} S^{S^2}(\tau_2 \mid \tilde{t} + 1, c', z').$$

Both A^{S^2} and B^{S^2} contain only identified objects, so $\gamma^{S^2} = A^{S^2}/B^{S^2}$ is identified.

Similarly, define

$$S^{S^1}(\tau_1 \mid t, c, z) = \prod_{h=0}^{\tau_1-1} \delta^{S^1}(1 - Q^{S^1}(t + h, c, z)), \quad B^{S^1}(t, c, z) = \sum_{\tau_1} \beta^{\tau_1} S^{S^1}(\tau_1 \mid t, c, z).$$

The indifference condition implies

$$0 = A^{S^1}(t, c, z) - \gamma^{S^1} B^{S^1}(t, c, z),$$

where

$$A^{S^1}(t, c, z) = \sum_{c', \tau_1} \beta^{\tau_1+1} \Lambda^{S^1, f}(c', \tau_1 \mid t, c, z) \\ \iint \max\{a_{t+\tau_1} + \eta^{1*}(t, c, z) + u + v - c', 0\} dG(v \mid t + \tau_1 - e) dF^{S^2}(u \mid t + \tau_1 - e) \\ + \sum_{c', z', \tau_1} \beta^{\tau_1+1} \Lambda^{S^1}(c', z', \tau_1 \mid t, c, z) \\ \int W^{S^2}(t + \tau_1 + 1, c', z', \eta^{1*}(t, c, z) + u) dF^{S^2}(u \mid t + \tau_1 - e).$$

In the first term, Step 2 identifies $F^{S^2}(\cdot \mid t + \tau_1 - e)$ because a non-terminal second study at the same date has positive probability, while Step 1 and Assumption (5) identify $G(\cdot \mid t + \tau_1 - e)$. In the second term, W^{S^2} is identified from the Stage-2 problem. Thus $\gamma^{S^1} = A^{S^1}/B^{S^1}$ is identified. $\square$

D State Transitions, Study 1 State Distribution and Potential Entrants

D.1 Details on Project States in the Empirical Model

We first describe a project's post-entry state transitions. Interconnection costs are measured in millions of dollars per megawatt. To model cost transitions and simplify state spaces, we discretize the costs into four bins: $[0, 0.01]$, $(0.01, 0.05]$, $(0.05, 0.2]$, $(0.2, \infty)$. Let $\ell \in \{1, 2, 3, 4\}$ denote the latest study's cost level. When a project receives a terminal study of level ℓ , we assume its actual cost is drawn from the empirical distribution of final costs in the corresponding bin. We can express the states as $(s, \tau, \ell, z^{\text{network}}, z^{\text{test}})$, where s denotes which study has been received, $\tau = t - e$ denotes the waiting time since the potential project's arrival and entry decision, ℓ denotes the cost level, z^{network} indicates whether the project is assigned a network allocation, and z^{test} indicates whether a set of engineering tests have been performed. Because a project arriving in e joins the queue in $e + 1$, $\tau = 1$ in its first queue period.

The state transitions are built from five sets of probabilities that are functions of a project's current state and characteristics, including fuel type, size, and location. We use p^A for the probability that a study arrives during the current period t , which depends on the queue status of a project, $\mathcal{K}_t(e, x)$. Conditional on study arrival and the prior state, we use $p^c, p^{\text{test}}, p^{\text{network}}$, and p^f for the probabilities of a new cost level, engineering tests completion conditional on $z^{\text{test}} = 0$, being assigned a network allocation conditional on $z^{\text{network}} = 0$, and the second study being final. Both z^{test} and z^{network} are absorbing: once they become 1, they remain 1 in subsequent studies.

D.2 Transition Equations

We now express the transition probabilities in Section 5.1 using the primitive probabilities above. Let $z = (z^{\text{network}}, z^{\text{test}})$, and let $s = 1$ denote waiting for the second study and $s = 2$ denote waiting for the final study.

For a project in Stage 1, the probability that the second study arrives and is terminal with new cost level ℓ' is

$$q_i^{\text{S1},f}(\ell' \mid \ell, t, z; \mathcal{K}_t(e, x)) = p_i^A(s = 1, \tau, \ell, z; \mathcal{K}_t(e, x)) p_i^c(\ell' \mid s = 1, \tau, \ell, z) p_i^f(\tau, \ell, z).$$

The probability that a non-terminal second study arrives with updated state (ℓ', z') is

$$q_i^{S1}(\ell', z' \mid \ell, t, z; \mathcal{K}_t(e, x)) = p_i^A(s = 1, \tau, \ell, z; \mathcal{K}_t(e, x)) p_i^c(\ell' \mid s = 1, \tau, \ell, z) (1 - p_i^f(\tau, \ell, z)) \\ \times h^{\text{network}}(z^{\text{network}'} \mid z^{\text{network}}; \mathcal{K}_t(e, x)) h^{\text{test}}(z^{\text{test}'} \mid z^{\text{test}}; \mathcal{K}_t(e, x)).$$

where (suppressing dependence on $\mathcal{K}_t(e, x)$)

$$h^{\text{test}}(1 \mid 1) = 1, \quad h^{\text{test}}(1 \mid 0) = p^{\text{test}}, \quad h^{\text{test}}(0 \mid 1) = 0, \quad h^{\text{test}}(0 \mid 0) = 1 - p^{\text{test}}, \\ h^{\text{network}}(1 \mid 1) = 1, \quad h^{\text{network}}(1 \mid 0) = p^{\text{network}}.$$

Therefore

$$Q_i^{S1}(\ell, t, z; \mathcal{K}_t(e, x)) \\ = \sum_{\ell'} q_i^{S1,f}(\ell' \mid \ell, t, z; \mathcal{K}_t(e, x)) \\ + \sum_{\ell', z'} q_i^{S1}(\ell', z' \mid \ell, t, z; \mathcal{K}_t(e, x)) \\ = p_i^A(s = 1, \tau, \ell, z; \mathcal{K}_t(e, x)).$$

For a project in Stage 2, every arriving study is terminal. Thus

$$q_i^{S2}(\ell' \mid \ell, t, z; \mathcal{K}_t(e, x)) = p_i^A(s = 2, \tau, \ell, z; \mathcal{K}_t(e, x)) p_i^c(\ell' \mid s = 2, \tau, \ell, z), \\ Q_i^{S2}(\ell, t, z; \mathcal{K}_t(e, x)) = p_i^A(s = 2, \tau, \ell, z; \mathcal{K}_t(e, x)).$$

A project exits with 0 payoff when reaching the maximum duration (32 quarters after entry).

D.3 Transition Estimates

Table D.1 reports the parameter estimates used in these transition probabilities. The cost column reports parameters from the ordered-probit index. The other columns are from probit indices. We also report observed and fitted means. We find that receiving engineering tests, which indicate that PJM deems the project likely to trigger an overload, is correlated with slower study arrival. Being assigned network allocation, which may indicate that a project shares costs with others, is correlated with a lower next study cost. These estimates are robust to including more controls (Tables E.5 and E.11).

Table D.1: Transition Function Estimates

	Ordered-probit cost state	Study arrival	Network allocation	Receive enrg. tests	2nd study is final
Cutoff 2	0.807*** (0.037)				
Cutoff 3	1.177*** (0.103)				
Prior cost: (0.01, 0.05]	0.599*** (0.083)				
Prior cost: (0.05, 0.2]	1.211*** (0.079)				
Prior cost: > 0.2	2.117*** (0.094)				
Waiting for third study		-0.380* (0.209)			
Wind	0.594*** (0.111)	0.023 (0.052)	0.422*** (0.133)	0.909** (0.365)	-0.848*** (0.218)
Solar	0.526*** (0.073)	0.131*** (0.042)	0.671*** (0.090)	1.757*** (0.216)	-1.086*** (0.101)
Battery	0.508*** (0.094)	-0.054 (0.056)	1.169*** (0.132)	1.878*** (0.340)	-1.550*** (0.146)
Uprate	-0.259*** (0.064)	-0.053 (0.034)	0.218*** (0.082)	0.019 (0.213)	0.374*** (0.091)
10–20 MW		0.054 (0.049)	0.411*** (0.101)	0.746*** (0.234)	-0.413*** (0.096)
20–100 MW		0.022 (0.047)	0.785*** (0.105)	0.845*** (0.229)	-1.300*** (0.118)
> 100 MW		-0.141*** (0.051)	1.077*** (0.119)	1.209*** (0.314)	-1.999*** (0.148)
Receive enrg. tests	0.065 (0.070)	-0.172*** (0.039)			
Network allocation	-0.352*** (0.086)	-0.063 (0.046)			
ln(MW in TO Queue)			0.086*** (0.023)		
ln(# Higher priority)		-0.092*** (0.026)			
ln(# Higher priority) × Waiting for third study		-0.182*** (0.036)			
ln(# Higher priority in TO)		-0.125*** (0.015)			
ln(# Higher priority in TO) × Waiting for third study		-0.038 (0.026)			
Observed mean	0.324	0.112	0.697	0.744	0.192
Fitted mean	0.436	0.111	0.698	0.745	0.189
Observations	2648	23627	2002	394	2002

The ordered-probit cost-state column uses four cost categories: [0, 0.01], (0.01, 0.05], (0.05, 0.2], and > 0.2 million dollars per MW. The first cutoff is normalized to zero; Cutoff 2 and Cutoff 3 are the estimated higher cutoffs. Prior cost rows indicate the project's most recent prior study cost, in million dollars per MW. ln(MW in TO Queue) is the natural log total MW of projects in the same transmission-owner queue. For the cost-state column, observed and fitted means are shares with new cost above \$0.1 million/MW among study arrivals. Other means compare the average binary outcomes with the predicted probabilities. The study arrival probit model also controls for indicators for the number of periods in the queue and their interactions with whether they are waiting for the third study. Coefficient estimates are reported. Standard errors are in parentheses.

D.4 First Study State Distribution

The first-study state is constructed separately from the later transition probabilities. In the empirical implementation, the first-study quarter is the observed queue-entry quarter, corresponding to period $e + 1$ in the model. We group projects by observed queue-entry quarter, transmission owner, fuel type (wind, solar, battery, and other), whether it expands an existing plant, and project size bins with the cutoffs at 10, 20, 50, 100, 200, 500, 1000, 1250 MW. For each group, we compute the marginal distributions of each state. Under the assumption that these states are independently distributed, we construct the first study state of cost and study information by multiplying the respective marginals.

D.5 Potential Entrants

We construct the non-entering potential projects from the empirical distribution of observed queue entrants. We group observed entrants by entry year (first-study date), TO, fuel type indicators (wind, solar, storage, and other), whether the project expands an existing plant, and a coarse size bin: ≤ 20 MW, 20–100 MW, or > 100 MW. We then compute the maximum number of observed entrants of that group across cohort years. In each cohort year, if the observed number of entrants of a group is below this group-specific maximum (across years), we add potential non-entrants drawn from entrants in that group until the cohort-year count reaches the maximum. For each potential non-entrant, we draw its queue-entry quarter from the empirical distribution within the same entry year.

The full construction creates 10,275 potential non-entrants over 2012-2020. To keep the equilibrium simulation manageable, the baseline estimation uses a random 5 percent subsample of these non-entrants. Therefore the number of sampled potential non-entrants is 514. The robustness analysis in Section 6.3 uses a random 10 percent subsample.

D.6 Exogenous Long-Tail Withdrawal

The structural model implies that projects do not endogenously withdraw between studies. We therefore assume that the observed exits during these periods are exogenous, and we estimate the exogenous withdrawal probability with a probit model. The estimation sample covers (1) for projects receiving a non-terminal first study, project decisions between the first quarter after receiving the study and either the quarter before it receives second study or the

quarter it withdraws, (2) for projects receiving a non-terminal second study, the decisions between the first quarter after receiving the study and either the quarter before it receives the third study or the quarter it withdraws. Observations are at the project-quarter level. Table D.2 reports the estimates. The implied average probabilities of an exogenous exit are 0.029 in stages 1 and 2.


Table D.2: Exogenous Withdrawal Probability

	p^W
Waiting for Study 3	-0.058 (0.040)
Wind	-0.374*** (0.081)
Solar	-0.622*** (0.057)
Battery	-0.524*** (0.069)
Capacity increase at existing plant	-0.391*** (0.050)
ln(MW)	-0.189*** (0.015)
Constant	-0.546*** (0.089)
Observations	18,618
Fitted Mean	0.029

The dependent variable equals one if the project withdraws during the waiting sample. Waiting for Study 3 indicates projects that have received a non-terminal second study and are waiting for the third study. Coefficient estimates are reported. Standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E Additional Figures and Tables

Figure E.2: Costs by Location




Second study interconnection cost estimates by location. Costs in millions of 2020 dollars per MW.

Table E.3: Top 15 Developers and Selective Completion

Developer	No. of projects	Share of projects (%)	Share of capacity (%)	Cohorts with a project	Avg. proj. per cohort
Invenergy	83	4.3	8.1	20	4.2
Community Energy Solar	68	3.5	1.0	19	3.6
Dominion	63	3.3	5.1	21	3.0
PSEG	56	2.9	1.8	18	3.1
LS	47	2.4	3.8	22	2.1
NextEra	38	2.0	2.0	12	3.2
Apex	30	1.6	3.5	14	2.1
EDF	30	1.6	3.2	12	2.5
SunEnergy1	30	1.6	1.5	10	3.0
EffiSolar	29	1.5	0.2	5	5.8
IMG	29	1.5	0.3	10	2.9
EDP	28	1.5	1.8	13	2.2
Exelon	27	1.4	1.7	12	2.3
Urban Grid	27	1.4	1.2	6	4.5
AEP	25	1.3	0.3	12	2.1

Summary statistics for the top 15 developers in our sample by number of projects. Share of projects is the developer's project count as a percentage of the 1,922 queue projects for which a developer could be identified. Share of capacity is the developer's total nameplate capacity as a percentage of the 190.3 GW of capacity in those same projects. Cohorts with a project is the number of queue entry cohorts (out of 34) in which the developer has at least one project.

Figure E.3: Study Wait Time (months) by Date of Queue Entry



This figure shows the time to receive the terminal study, measured in months, by entry cohort. The horizontal bar represents the interquartile range (25th-75th percentiles), and the dot indicates the median. “% Terminal” is the share of projects that have received the terminal study among projects entering in a given cohort. “% Active” is the share of projects that are still active among projects in that cohort.

Table E.4: Summary Statistics by Transmission Owner

	Mean	Std. Dev.	25%	75%
<i>Panel A: Transmission owner-level totals and shares</i>				
No. of projects	190	191	33	292
Total capacity (GW)	18.5	22.7	3.6	21.1
Avg. capacity per project (MW)	101.9	54.9	60.4	139.1
No. of withdrawals	102	96	20	168
Total withdrawn capacity (GW)	10.20	10.63	2.12	14.56
No. of completions	33	34	4	55
Completed capacity (GW)	2.34	3.13	0.07	3.42
Share renewables	0.69	0.25	0.63	0.88
No. of unique developers	31	26	10	48
<i>Panel B: Within-transmission owner averages</i>				
New projects per year	14.8	14.5	2.5	22.5
Waiting time (mos.)	26.8	4.7	23.9	29.4
First study cost (\$M/MW)	0.14	0.15	0.07	0.16
Time to first study (mos.)	5.3	0.8	4.6	5.4
Second study cost (\$M/MW)	0.20	0.17	0.10	0.20
Time to second study (mos.)	14.7	4.4	12.4	15.7
Third study cost (\$M/MW)	0.11	0.09	0.05	0.15
Time to third study (mos.)	16.9	4.1	13.6	18.8
Transmission owners	22			

Panel A reports the distribution across transmission owners of aggregate statistics computed within each transmission owner. Unique developers counts distinct developers within each transmission owner; we are unable to match 2,365 projects (56.7%) in the full project-level sample to a developer that controls the project. Panel B first averages within each transmission owner across its projects, then reports the distribution of those transmission owner-level averages. New projects per year is the total number of projects in a transmission owner divided by the number of years spanned by its queue entries. Waiting time is months from queue entry to completion or withdrawal, excluding projects whose last known status is neither. Study costs (in millions of dollars per MW) and times (in months) are conditional on receiving the respective study. Renewables is solar, wind, and batteries.

Table E.5: New Study Arrival Probit Models With Additional Controls

	(1) Probit	(2) Probit	(3) Probit	(4) Probit	(5) Probit
ln(# Higher priority)	-0.031*** (0.006)	-0.032*** (0.006)	-0.020 (0.012)	-0.021* (0.012)	-0.028*** (0.007)
ln(# Higher priority) X Waiting for third study	-0.016** (0.008)	-0.016** (0.008)	-0.011 (0.007)	-0.011 (0.007)	-0.015** (0.008)
ln(# Higher priority in TO)	-0.016*** (0.003)	-0.014*** (0.004)	-0.020*** (0.007)	-0.020*** (0.007)	-0.023*** (0.006)
ln(# Higher priority in TO) X Waiting for third study	-0.012* (0.007)	-0.013* (0.007)	-0.011* (0.007)	-0.011* (0.007)	-0.011 (0.007)
Waiting for third study	-0.216*** (0.023)	-0.215*** (0.023)	-0.243*** (0.023)	-0.243*** (0.023)	-0.221*** (0.023)
Receive engr. tests	-0.027** (0.011)	-0.028** (0.012)	-0.030*** (0.010)	-0.030*** (0.010)	-0.028** (0.011)
Network allocation	-0.002 (0.012)	-0.003 (0.012)	0.000 (0.012)	-0.000 (0.012)	-0.003 (0.012)
ln(Size of withdrawals in queue within TO)					-0.000 (0.002)
ln(Size of projects in queue within TO)					0.008* (0.005)
Wholesale electricity price (100\$/MWh)		0.042** (0.020)		0.032 (0.024)	
Ordinance		-0.007 (0.005)		0.004 (0.005)	
Prior Study Controls	X	X	X	X	X
Transmission Owner FE			X	X	
Year FE			X	X	
Observations	28,917	28,917	28,917	28,917	28,917
Pseudo R ²	0.176	0.177	0.195	0.196	0.176

The sample includes projects queued from 2008-2020, observed over time at the quarterly level as they wait in the queue for study arrival. Marginal effects from probit regressions are reported. Standard errors in parentheses, clustered by transmission owner - entry cohort (416 clusters). The dependent variable is an indicator for study arrival (overall mean = 0.11; mean for second study is 0.16 and for third study is 0.06). All specifications control for project size (three bins), fuel type (wind, solar, storage), if the project is a capacity increase at an existing plant, time spent waiting in the queue (controlled for flexibly with a separate indicator for each quarter since queue entry), as well as their interactions with an indicator for whether the project is awaiting its third study, and prior study characteristics, including the project's own cost estimate (three bins) from the prior study (e.g., the first study costs for projects awaiting the second study), indicators for having ever received each of the three engineering tests and indicator for whether the project receives an explicit network upgrade cost. Specifications (2) and (4) control for wholesale electricity prices, defined as the average real-time locational marginal price at the project's nearest substation over peak hours (hours ending 07 through 23) in that quarter, in nominal \$/MWh, and ordinance, an indicator for a local ordinance restricting renewable energy development. Specification (5) further control for the log size of withdrawn projects waiting for the same or a later study within the same transmission owner over the past two quarters, and the log size of projects in the queue within the same transmission owner. *** p<0.01, ** p<0.05, * p<0.10.

Table E.6: New Study Arrival Probit Models With Different Time Controls

	(1) None	(2) Linear	(3) Quadratic	(4) YearxQtr FE	(5) TOxYear FE	(6) None
ln(# Higher priority)	-0.031*** (0.006)	-0.027*** (0.009)	-0.025** (0.011)	-0.033*** (0.010)	-0.005 (0.013)	-0.033*** (0.012)
ln(# Higher priority) X Waiting for third study	-0.016** (0.008)	-0.017** (0.008)	-0.017** (0.008)	-0.012 (0.008)	-0.012* (0.007)	0.014 (0.012)
ln(# Higher priority in TO)	-0.016*** (0.003)	-0.016*** (0.003)	-0.016*** (0.003)	-0.016*** (0.003)	-0.026** (0.010)	-0.025*** (0.005)
ln(# Higher priority in TO) X Waiting for third study	-0.012* (0.007)	-0.012* (0.007)	-0.012* (0.007)	-0.012* (0.007)	-0.010 (0.007)	-0.005 (0.011)
<i>Total queue-position effect (linear combinations)</i>						
At second study	-0.047*** (0.006)	-0.043*** (0.009)	-0.041*** (0.011)	-0.049*** (0.011)	-0.031*** (0.011)	-0.058*** (0.012)
At third study	-0.075*** (0.006)	-0.071*** (0.008)	-0.070*** (0.010)	-0.074*** (0.010)	-0.052*** (0.010)	-0.050*** (0.013)
Sample Period	2008-2020	2008-2020	2008-2020	2008-2020	2008-2020	2008-2018
Observations	28,917	28,917	28,917	28,295	28,105	14,109
R ²	0.176	0.176	0.178	0.210	0.233	0.077

The sample includes projects queued from 2008-2020, observed over time at the quarterly level as they wait in the queue for study arrival. Marginal effects from probit regressions are reported. Standard errors are in parentheses and clustered by transmission owner - entry cohort (416 clusters). The dependent variable is an indicator for study arrival (overall mean = 0.11; mean for second study is 0.16 and for third study is 0.06). All specifications control for project size (three bins), fuel type (wind, solar, storage), if the project is a capacity increase at an existing plant, time spent waiting in the queue (controlled for flexibly with a separate indicator for each quarter since queue entry), as well as their interactions with an indicator for whether the project is awaiting its third study, and prior study characteristics, including the project's own cost estimate (three bins) from the prior study (e.g., the first study costs for projects awaiting the second study), indicators for having ever received each of the three engineering tests and indicator for whether the project receive an explicit network upgrade cost. Columns (1)–(4) report estimates with different time controls; column (4) uses a full set of year-by-quarter fixed effects. Column (5) replaces the time controls with transmission-owner-by-year fixed effects, which absorb any shock common to a transmission owner in a given year. Column (6) uses the same time controls as column (1) but restricts the sample to projects with a queue year prior to 2019, as shown in the Sample Period row. The total-effect rows report linear combinations of the marginal queue-position effects. *At second study* sums the system-wide and transmission-owner queue effects, while *At third study* additionally includes their interactions with the indicator for awaiting the third study. *** p<0.01, ** p<0.05, * p<0.10.

Table E.7: Effect of Active Higher-Queued Capacity on Probability of High Interconnection Cost

	10km (1)	20km (2)	Transmission owner (3)	Hierarchical cluster (4)
Panel A: High Total Cost Estimate (mean = 0.37)				
Log active higher-priority MW	-0.034* [-0.072, 0.005]	-0.025 [-0.059, 0.009]	-0.010 [-0.036, 0.018]	0.003 [-0.023, 0.028]
Geographic Region FE	X	X	X	X
Queue Year FE	X	X	X	X
Observations	2,328	2,328	2,328	2,328
R^2	0.518	0.518	0.231	0.230
Panel B: High Network Upgrade Cost Estimate (mean = 0.13)				
Log active higher-priority MW	0.006 [-0.024, 0.036]	0.009 [-0.015, 0.034]	0.004 [-0.025, 0.028]	0.009 [-0.016, 0.033]
Geographic Region FE	X	X	X	X
Queue Year FE	X	X	X	X
Observations	2,328	2,328	2,328	2,328
R^2	0.477	0.477	0.178	0.204

Projects queuing from 2009–2020. Wild cluster bootstrap (9,999 replications) clustered at the substation level in the 10 km and 20 km columns, the transmission owner level in the transmission owner column, and the hierarchical cluster level in the hierarchical cluster column; 95% confidence intervals in brackets. In Panel A, the dependent variable is an indicator for whether the total cost estimate in the second or third study is high, defined as exceeding \$0.1 million per MW. In Panel B, the dependent variable is an indicator for whether the network upgrade cost estimate in the second or third study is high, defined as exceeding \$0.1 million per MW. The regressor of interest is the log of total MW from the focal project and other active projects in the same or earlier interconnection cohort in the local geographic region in the quarter when the focal study is issued. An earlier cohort code indicates higher queue priority. A project is active if it is still waiting in the queue, meaning it has not completed or withdrawn, in the focal study's issue quarter. Columns vary the geographic definition: 10 km, 20 km, transmission owner, and hierarchical cluster. All specifications control for project size (three bins), fuel type (wind, solar, storage), whether the project is a capacity increase at an existing plant; whether PJM Regional Transmission Expansion Plan investment was above the 75th percentile of the RTEP distribution within the same transmission owner, measured over a four-year window spanning three years before and one year after the study issue date. The geographic region fixed effects are substation fixed effects in the 10 km and 20 km columns, transmission owner fixed effects in the transmission owner column, and hierarchical cluster fixed effects in the hierarchical cluster column. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table E.8: Robustness for Effect of a Costly Prior Interconnection on the Probability of High Total Cost

	Definition of nearby prior interconnection							
	10km		20km		Transmission owner		Hierarchical cluster	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>A. Total cost > \$1m</i>								
Costly prior	0.066 [-0.041, 0.172]	0.056 [-0.045, 0.156]	-0.004 [-0.092, 0.084]	-0.005 [-0.091, 0.082]	-0.010 [-0.077, 0.060]	-0.012 [-0.078, 0.056]	-0.036** [-0.067, -0.007]	-0.031** [-0.063, -0.000]
<i>B. Total cost > 90th percentile</i>								
Costly prior	0.091 [-0.067, 0.247]	0.061 [-0.079, 0.203]	-0.064 [-0.177, 0.051]	-0.072 [-0.194, 0.061]	-0.038 [-0.170, 0.202]	-0.012 [-0.167, 0.236]	-0.049** [-0.094, -0.001]	-0.038* [-0.080, 0.006]
<i>C. Cost/MW > \$0.1m</i>								
Costly prior	0.010 [-0.118, 0.136]	0.018 [-0.116, 0.146]	-0.015 [-0.135, 0.104]	-0.014 [-0.140, 0.112]	-0.028 [-0.111, 0.052]	-0.016 [-0.081, 0.044]	0.043 [-0.012, 0.096]	0.045 [-0.011, 0.099]
<i>D. Cost/MW > 90th percentile</i>								
Costly prior	-0.061 [-0.312, 0.165]	-0.027 [-0.254, 0.176]	-0.026 [-0.170, 0.115]	-0.010 [-0.153, 0.132]	-0.026 [-0.091, 0.103]	-0.035 [-0.081, 0.134]	0.069* [-0.009, 0.144]	0.071* [-0.007, 0.148]
Geographic Region FE	X	X	X	X	X	X	X	X
Queue Year FE		X		X		X		X

Projects queuing from 2009–2020. Wild cluster bootstrap (9,999 replications) clustered at the substation level in the 10 km and 20 km columns, the transmission owner level in the transmission owner column, and the hierarchical cluster level in the hierarchical cluster column; 95% confidence intervals in brackets. The dependent variable is an indicator for whether the cost estimate in the second or third study is high, defined as exceeding \$0.1 million per MW (mean = 0.37). The 90th percentile for the cost of prior interconnection is 9.17, 9.37, 7.54, and 7.64 million dollars for total cost and 0.27, 0.19, 0.16, and 0.15 million dollars for cost per MW, for the 10km, 20km, transmission owner, and hierarchical cluster definitions, respectively. The mean of the key independent variable, Costly prior interconnection, is 0.10, 0.02, 0.06, and 0.03 for the 10km specification across Panels A–D; 0.16, 0.04, 0.09, and 0.04 for the 20km specification; 0.29, 0.08, 0.15, and 0.07 for the transmission owner level specification; and 0.22, 0.07, 0.11, and 0.07 for the hierarchical cluster level specification. Each estimate corresponds to a separate regression, with rows corresponding to different thresholds for costly prior interconnection and columns corresponding to different definitions of nearby. All specifications control for project size (three bins), fuel type (wind, solar, storage), whether the project is a capacity increase at an existing plant; whether PJM Regional Transmission Expansion Plan investment was above the 75th percentile of the RTEP distribution within the same transmission owner, measured over a four-year window spanning three years before and one year after the study issue date; the log of total entrant capacity (MW) within the previous three years in the local geographic region; and the log capacity (MW) of higher-queued projects in the local geographic region. The geographic region fixed effects are substation fixed effects in the 10 km and 20 km columns, transmission owner fixed effects in the transmission owner column, and hierarchical cluster fixed effects in the hierarchical cluster column. *** p<0.01, ** p<0.05, * p<0.10.

Table E.9: Robustness for Effect of a Costly Prior Interconnection on the Probability of High Network Upgrade Cost

	Definition of nearby prior interconnection							
	10km		20km		Transmission owner		Hierarchical cluster	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>A. Network upgrade cost > \$1m</i>								
Costly prior	0.146 [-0.225, 0.562]	0.111 [-0.217, 0.475]	0.036 [-0.156, 0.271]	0.060 [-0.129, 0.270]	-0.058 [-0.120, 0.046]	-0.044 [-0.106, 0.090]	0.012 [-0.059, 0.087]	0.037 [-0.032, 0.110]
<i>B. Network upgrade cost > 90th percentile</i>								
Costly prior	0.068 [-0.335, 0.573]	0.024 [-0.289, 0.439]	-0.019 [-0.223, 0.266]	0.004 [-0.211, 0.248]	-0.056 [-0.113, 0.071]	-0.043 [-0.090, 0.058]	0.011 [-0.060, 0.086]	0.034 [-0.037, 0.110]
<i>C. Network upgrade cost/MW > \$0.1m</i>								
Costly prior	0.167 [-2.823, 21.352]	0.163 [-4.064, 2.913]	-0.079 [-7.931, 1.709]	-0.008 [-3.421, 2.968]	-0.086 [-1.403, 2.081]	-0.096 [-0.732, 0.643]	-0.018 [-2.868, 2.530]	-0.066 [-2.085, 4.013]
<i>D. Network upgrade cost/MW > 90th percentile</i>								
Costly prior	0.207 [-0.366, 0.753]	0.178 [-0.219, 0.580]	0.096 [-0.135, 0.368]	0.113 [-0.089, 0.333]	-0.053 [-0.105, 0.071]	-0.041 [-0.084, 0.058]	-0.007 [-0.109, 0.066]	0.007 [-0.066, 0.080]
Geographic Region FE	X	X	X	X	X	X	X	X
Queue Year FE		X		X		X		X

Projects queuing from 2009–2020. Wild cluster bootstrap (9,999 replications) clustered at the substation level in the 10 km and 20 km columns, the transmission owner level in the transmission owner column, and the hierarchical cluster level in the hierarchical cluster column; 95% confidence intervals in brackets. The dependent variable is an indicator for whether the network upgrade cost estimate in the second or third study is high, defined as exceeding \$0.1 million per MW (mean = 0.13). In this table the costly-prior indicators are defined using the prior interconnection's network upgrade cost. In Panel C, a prior interconnection is classified as costly if its network upgrade cost per MW exceeds \$0.1 million. The 90th percentile for the network upgrade cost of prior interconnection is 2.03, 2.03, 0.07, and 0.61 million dollars for network upgrade cost and 0.04, 0.02, 0.00, and 0.01 million dollars for network upgrade cost per MW, for the 10km, 20km, transmission owner, and hierarchical cluster definitions, respectively. The mean of the key independent variable, Costly prior interconnection, is 0.01, 0.01, 0.00, and 0.01 for the 10km specification across Panels A–D; 0.02, 0.02, 0.00, and 0.02 for the 20km specification; 0.02, 0.05, 0.01, and 0.05 for the transmission owner level specification; and 0.03, 0.03, 0.00, and 0.03 for the hierarchical cluster level specification. Each estimate corresponds to a separate regression, with rows corresponding to different thresholds for costly prior interconnection and columns corresponding to different definitions of nearby. All specifications control for project size (three bins), fuel type (wind, solar, storage), whether the project is a capacity increase at an existing plant; whether PJM Regional Transmission Expansion Plan investment was above the 75th percentile of the RTEP distribution within the same transmission owner, measured over a four-year window spanning three years before and one year after the study issue date; the log of total entrant capacity (MW) within the previous three years in the local geographic region; and the log capacity (MW) of higher-queued projects in the local geographic region. The geographic region fixed effects are substation fixed effects in the 10 km and 20 km columns, transmission owner fixed effects in the transmission owner column, and hierarchical cluster fixed effects in the hierarchical cluster column. *** p<0.01, ** p<0.05, * p<0.10.

Table E.10: Predictors of Cost-Share Group Assignment

	<i>Dependent variable: Cost-share group assignment (mean = 0.68)</i>							
	Panel A: Higher-queued capacity				Panel B: Total active capacity			
	10km (1)	20km (2)	Trans. owner (3)	Hierarchical (4)	10km (5)	20km (6)	Trans. owner (7)	Hierarchical (8)
Log active MW (local region)	0.011 (0.019)	0.010 (0.018)	0.039* (0.021)	0.016 (0.014)	0.025 (0.018)	0.031* (0.018)	0.046*** (0.012)	0.032** (0.012)
Geographic Region FE	X	X	X	X	X	X	X	X
Queue Year FE	X	X	X	X	X	X	X	X
Observations	1,398	1,398	1,398	1,398	1,398	1,398	1,398	1,398
R^2	0.653	0.653	0.458	0.468	0.654	0.655	0.461	0.471

Projects queuing from 2009–2020 at the second study. The dependent variable is an indicator for whether the project is assigned to a cost-sharing group (sample mean = 0.68); estimates are from linear probability models. Standard errors are in parentheses and clustered by hierarchical cluster. Each column is a separate regression. Panel A (columns (1)–(4)) measures local capacity as the log MW of active higher-queued projects; Panel B (columns (5)–(8)) uses the log total active MW (all cohorts). Active projects are those still waiting in the queue in the study issue quarter. Within each panel, columns vary the geographic definition: 10 km, 20 km, transmission owner, and hierarchical cluster. All specifications control for project size (three bins), fuel type (wind, solar, storage), whether the project is a capacity increase at an existing plant; and whether PJM Regional Transmission Expansion Plan investment was above the 75th percentile of the RTEP distribution within the same transmission owner, measured over a four-year window spanning three years before and one year after the study issue date. The geographic region fixed effects are substation fixed effects in the 10 km and 20 km columns, transmission owner fixed effects in the transmission owner column, and hierarchical cluster fixed effects in the hierarchical cluster column.

*** p<0.01, ** p<0.05, * p<0.10.

Table E.11: Local Queue Scale and Network Allocation: Effect on High Interconnection Cost at Increasing Thresholds

	Total cost				Network upgrade cost			
	>0.10 (1)	>0.15 (2)	>0.20 (3)	>0.25 (4)	>0.10 (5)	>0.15 (6)	>0.20 (7)	>0.25 (8)
Log active higher-priority MW (local)	0.022 (0.014)	0.014 (0.013)	-0.002 (0.011)	-0.006 (0.010)	0.013 (0.008)	0.006 (0.008)	0.004 (0.008)	0.004 (0.007)
Network allocation	-0.054* (0.032)	-0.110*** (0.028)	-0.090*** (0.024)	-0.072*** (0.021)	-0.055*** (0.015)	-0.034*** (0.011)	-0.029*** (0.011)	-0.026*** (0.009)
Transmission Owner FE	X	X	X	X	X	X	X	X
Queue Year FE	X	X	X	X	X	X	X	X
Mean of dep. var.	0.34	0.23	0.19	0.15	0.11	0.08	0.08	0.07
Observations	3,121	3,121	3,121	3,121	3,121	3,121	3,121	3,121
R ²	0.298	0.295	0.307	0.293	0.150	0.145	0.146	0.146

Cost-transition estimates. Projects queuing from 2009–2020, observed at the first study or the second study as they transition to the next study. The dependent variable is an indicator for whether the *next* study’s cost per MW exceeds the column threshold, in millions of dollars per MW. Columns (1)–(4) use total interconnection cost; columns (5)–(8) use network upgrade cost. Standard errors are in parentheses and clustered by substation. Log active higher-priority MW (local) is the log of total MW of the focal project and other active projects in the same or earlier interconnection cohort within the focal project’s transmission owner. Network allocation is an indicator of whether a project receives an explicit allocation of network upgrade costs. All specifications control for the current study’s cost-category dummies (cutoffs at \$0.01, \$0.05, and \$0.20 per MW, in the same cost type as the outcome), fuel type (wind, solar, storage), whether the project is a capacity expansion at an existing plant, an engineering-test indicator, and a Study 2 indicator. *** p<0.01, ** p<0.05, * p<0.10.


Table E.12: Construction Cost Inputs by Fuel Type

	2012	2013	2014	2015	2016	2017	2018	2019	2020
Gas (combined cycle)	1.017	0.917	0.917	0.917	0.978	0.978	0.978	0.958	0.958
Wind (onshore)	–	2.105	1.918	1.814	1.758	1.739	1.424	1.408	1.498
Solar PV	–	2.408	2.270	1.899	1.391	1.523	1.201	1.167	1.076
Storage (4-hour battery)	1.750	1.489	1.331	1.026	0.849	0.701	0.621	0.572	0.521

The costs are in \$ million/MW. Gas uses EIA natural-gas combined-cycle installation costs. Wind uses EIA onshore-wind construction costs. Solar uses EIA Form 860/923 utility-scale PV installed costs in \$/kWAC, converted to \$ million/MW. Storage uses BloombergNEF battery pack prices plus NREL ATB-based balance-of-system costs for a 4-hour utility-scale battery. Both solar and battery are multiplied by 0.7 to reflect the investment tax credit. “–” denotes no observation for that year.

Figure E.4: Effect of Interconnections with Network Upgrade Costs on Local Wholesale Prices

(a) Adding 20 MW New Capacity that Pays Network Upgrade Costs (b) Adding 100 MW New Capacity that Pays Network Upgrade Costs



Panel (a) plots event-study estimates of the effect of a new generator with capacity of at least 20 MW beginning operation at a substation on local wholesale electricity prices, focusing on substations where the entering generator is responsible for transmission network upgrade costs as part of the interconnection process. Panel (b) uses a 100 MW threshold. Treated observations are substation-months in which a generator above the corresponding capacity threshold begins operation and is responsible for network upgrade costs; control observations are substation-months without such entry and include substation-months with entry of smaller or non-network-upgrade-paying generators. The dependent variable is the monthly average peak-hour locational marginal price at the substation level (mean = \$34.89/MWh), constructed from PJM locational marginal price data over 2008-2020. Estimates are obtained using the Callaway and Sant'Anna (2021) difference-in-differences estimator with doubly robust inverse probability weighting. Standard errors are clustered at the substation level.


F Additional Counterfactual Tables and Figures

Table F.13: By-Size Entry Surplus Changes: \$20,000/MW and \$400,000 Flat

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	As-is	Per-MW fee			Flat fee		
Renewable	160.6	-15.7	-13.3	-10.2	-3.3	-3.1	-1.7
≤20 MW	118.7	-16.0	-13.8	-10.6	-21.5	-18.4	-14.4
20–100 MW	167.8	-16.3	-13.6	-11.2	-4.7	-4.0	-3.0
>100 MW	163.9	-15.4	-13.1	-9.7	-0.1	-0.4	+0.7
Non-renewable	191.3	-14.0	-11.2	-8.9	-0.3	-0.2	+0.3
≤20 MW	253.9	-16.2	-11.9	-9.1	-27.6	-19.1	-14.6
20–100 MW	268.8	-16.0	-12.2	-10.0	-5.3	-3.9	-3.3
>100 MW	184.6	-13.8	-11.1	-8.9	+0.9	+0.6	+1.0
<i>Design: share of fee forfeited</i>							
If withdraw before requesting the 2nd study		100%	100%	50%	100%	100%	50%
If withdraw later		100%	100%	100%	100%	100%	100%
If complete		100%	0%	0%	100%	0%	0%

Column (1) reports the as-is equilibrium in levels; columns (2)–(4) and (5)–(7) report changes relative to as-is under the fees of \$20,000 per MW of nameplate capacity (per-MW fee) and \$400,000 per project (flat fee), respectively, at queue entry. The design rows report the share of that payment forfeited at each outcome. In column (2), the entry fee is not refundable. In column (3), the deposit is refundable when the project completes but forfeited if the project withdraws. In column (4), the deposit is 50% refundable upon withdrawal at any point before the project chooses to receive the second study; thereafter, the deposit is refunded if the project completes but forfeited if it withdraws. Columns (5)–(7) report the same outcomes under the flat fee. Entry surplus is $E \max\{0, EV_i - C_i\}$.

Figure F.5: Change in Waiting Time by Generator Group under a \$20,000/MW Nonrefundable Fee



The figure plots the change in waiting time (quarters) after imposing a nonrefundable entry fee of \$20,000/MW, relative to the as-is equilibrium. Each group is defined as a fuel type, TO, arrival quarter, and size-bin combination. We group generators into eight size bins with cutoffs at 10, 20, 50, 100, 200, 500, and 1000 MW. There are 1,845 generator groups. For each group, we plot the change in time in queue. The time decreases for 1,761 groups, representing about 2,575 projects and 243 GW, and the time increases for 84 groups, representing about 131 projects and 19 GW.




Figure F.6: Completed GW vs. Per-MW Entry Fee Level

Each point is an equilibrium simulation. Five series: nonrefundable fee (solid), fee with budget-neutral rebate (dotted, triangle), staged deposit (dotted, diamond), staged deposit with budget neutral rebates (dashed, inverted triangle), and deposit (dashed, square). The as-is equilibrium is when fees are 0.

Table F.14: By-Size Entry Count Changes: \$20,000/MW and \$400,000 Flat

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	As-is	Per-MW fee		Flat fee			
Renewable	2061	-63	-58	-27	-196	-163	-102
≤ 20 MW	1030	-48	-46	-17	-193	-160	-101
20–100 MW	613	-8	-7	-6	-3	-2	-2
> 100 MW	418	-7	-6	-4	+0	+0	+0
Non-renewable	644	-8	-6	-5	-36	-29	-23
≤ 20 MW	322	-4	-3	-2	-36	-29	-23
20–100 MW	133	-1	-1	-1	-1	+0	+0
> 100 MW	189	-3	-3	-2	+0	+0	+0
<i>Design: share of fee forfeited</i>							
If withdraw before requesting the 2nd study		100%	100%	50%	100%	100%	50%
If withdraw later		100%	100%	100%	100%	100%	100%
If complete		100%	0%	0%	100%	0%	0%

Column (1) reports the as-is equilibrium in levels; columns (2)–(4) and (5)–(7) report changes relative to as-is under the fees of \$20,000 per MW of nameplate capacity (per-MW fee) and \$400,000 per project (flat fee), respectively, at queue entry. The design rows report the share of that payment forfeited at each outcome. In column (2), the entry fee is not refundable. In column (3), the deposit is refundable when the project completes but forfeited if the project withdraws. In column (4), the deposit is 50% refundable upon withdrawal at any point before the project requests the second study; thereafter, the deposit is refunded if the project completes but forfeited if it withdraws. Columns (5)–(7) report the same outcomes under the flat fee.




Figure F.7: Completed GW vs. Flat Fee Level

Each point is an equilibrium simulation. Three series: nonrefundable fee (solid), deposit (dashed), and staged deposit (dotted) where projects receive 50% of the deposit if they withdraw before requesting the second study, no deposit if they withdraw after, and full deposit if they complete.