State Capacity and Clean Energy Investment

Greg C. Wright

Department of Economics, University of California, Merced
gwright4@ucmerced.edu

May 10, 2026

Abstract

Infrastructure investment depends not only on regulatory standards, but also on agencies' capacity to administer them. To study this margin in U.S. clean-energy development, I compare solar projects near the same Army Corps of Engineers district boundary but assigned to different regulatory offices. A one-standard-deviation increase in non-energy workload is associated with delays of three to five months in final queue resolution. The shock is not associated with project build-out or size. Within the border sample, the delay implies 3.7–5.8 gigawatt-years of delayed capacity worth $411–$630 million. Limited public capacity reshapes when investment arrives, even without blocking projects.

Keywords: state capacity; administrative capacity; permitting; renewable energy; project delay.

JEL Codes: H54, Q42, Q48, Q58.

This draft studies administrative workload pressure and clean-energy project timing. All errors are my own.

1 Introduction

Large infrastructure projects cannot break ground until public agencies review, permit, and approve them. When these agencies are capacity constrained, public investment goals can be delayed even when projects are ultimately approved. This margin is central to current debates over state capacity and abundance: governments cannot build at scale if the institutions that review projects cannot process them at scale (Besley and Persson, 2010; Bagley, 2019; Klein and Thompson, 2025). Permitting is a useful setting because administrative throughput is both visible and consequential. The policy debate often asks whether permitting prevents projects from being built. A growing set of reforms instead targets the time it takes agencies to process permits, without necessarily changing the underlying legal standards. This paper studies that timing margin.[1]

The workload is a natural empirical measure of this form of state capacity. Regulatory offices have staff, procedures, and review obligations that cannot adjust one-for-one with incoming caseloads. When unrelated work rises, the same office must process more matters with limited review capacity. If that pressure delays private projects, state capacity becomes an economic constraint through processing speed rather than through formal rejection.

The workload shocks occur in a setting of tightening administrative resources. The Corps' public budget justifications report the amount of Regulatory Program funding sent to districts and divisions. In nominal terms, this field-funding line rose from $195.6 million in FY2016 to $201.8 million in FY2023 and $212.0 million in FY2024. In real 2018 dollars, however, it fell from $204.6 million in FY2016 to $166.3 million in FY2023 and $169.7 million in FY2024. Observed non-energy records in the public data were about 38 percent higher in FY2023 than in FY2016. Appendix Table 13 reports a simple pressure index: observed non-energy records per real field-funding dollar, normalized to FY2018. This descriptive fact does not identify the causal effect. Identification comes from local comparisons near district boundaries. It does, however, give the workload measure a direct state-capacity interpretation: offices process year-to-year caseload shocks against a background in which real resources have not kept pace with work.

This paper asks which margin matters in clean-energy development. Does limited administrative capacity stop observed projects from being built, reduce the scale of projects that do build, or mostly delay when capacity arrives? The distinction matters for both measurement and policy. If the main effect is cancellation, the relevant cost is missing investment. If the main effect is delay, the relevant cost is capacity-time: productive capacity and its benefits arrive later than they otherwise would. The research design separates margins that are often conflated. It cannot observe projects that developers never propose. Conditional on observed queue entry, it tests whether office pressure changes eventual build-out, matched built capacity, and the timing of the final queue outcome.

I study this question in renewable-energy development. I ask whether solar projects move more slowly when the regulatory office assigned to them is unusually busy, whether that pressure changes eventual build-out, and how large the resulting delay costs are. I focus on the Corps, whose

[1] One example is Maryland's use of artificial-intelligence tools to modernize environmental site-assessment requests (McKinsey & Company, 2023; Washington Post, 2026).

regulatory offices process records related to wetlands, waterways, and jurisdictional determinations (formal findings about whether a particular waterbody falls under federal regulation). I construct a district-year measure of non-energy workload in the Corps' public regulatory records and link it to solar projects in the Lawrence Berkeley National Laboratory (LBNL) interconnection queue.

The main design compares projects near Corps district boundaries. These projects are geographically close and often share state-year conditions, but they are assigned to different regulatory offices. The empirical challenge is that slow projects differ from fast projects. Projects with difficult permits may be larger, closer to regulated resources, or located in places where development is more contested. A simple comparison of fast and slow projects therefore mixes administrative delay with project selection. The border design is meant to isolate variation in office pressure from these project-level and local-market differences.

The empirical pattern has three parts. First, solar projects entering Corps districts during high-workload years take longer to reach a final queue outcome: either operation or withdrawal. The result survives county controls, boundary-pair controls, and comparisons among projects near the same boundary in the same entry year. In that most demanding local comparison, a one-standard-deviation increase in non-energy office workload is associated with a 0.172 increase in log days to final queue resolution, corresponding to about 144 additional days for the median project with a final outcome. A more conservative boundary-pair comparison implies 96 additional days.

Second, the same shock has little association with whether the project ultimately builds. This result is conditional on observed queue entry. Developers may decide not to propose projects when they expect permitting to be difficult, and that margin is not observed here. Within the queue, however, the evidence does not look like a build/no-build effect (a change in whether the project is ultimately constructed). A supplementary match to Energy Information Administration (EIA) plant records also does not show robust evidence that high-workload projects are smaller when they do build. The pattern looks like a timing effect.

Third, the timing effect has economically meaningful scale. The main economic object is capacity-time, defined as the megawatts (MW) that eventually arrive but arrive later than they otherwise would. In the built-project border sample, moving above-normal workload years back to normal workload implies 3.7 to 5.8 gigawatt-years (GW-years) of delayed built solar capacity. Under central assumptions, this corresponds to about $411 million to $630 million in delayed generation value. These dollar figures are the paper's main quantitative scale. Applying the same per-project effect proportionally across the full LBNL solar interconnection queue with terminal outcomes would imply a back-of-envelope national scale on the order of $2 to $3 billion, but that queue-wide number is illustrative rather than separately identified. The boundary-pair design isolates within-pair, within-year office-pressure variation among projects within 50 km of a true neighboring district boundary, and projects in the interior of large districts are not part of the identifying variation. Even within that scope, the capacity-time magnitudes are comparable to recent estimates of the per-program welfare cost of regulatory delay in interstate infrastructure (Liscow, 2025). The EIA match provides suggestive support for the timing interpretation outside

the queue data. Among projects that eventually operate, high-workload offices are also associated with later recorded operation in federal plant records.

The paper makes three contributions. First, it provides quasi-experimental evidence on the implementation margin of state capacity using a within-boundary comparison rather than an external instrument. State-capacity and abundance work emphasizes this distinction (Besley and Persson, 2009; Klein and Thompson, 2025; Finan, Olken, and Pande, 2017) but rarely measures it directly in the U.S. infrastructure setting. Second, conditional on observed queue entry, the binding margin in the queue is timing, not completion: limited administrative capacity changes when capacity arrives, without detectably changing whether projects build or how large they are when built. The design does not identify whether developers propose fewer projects when they expect a slow office, or what happens to projects in the interior of large districts that lack a neighboring Corps boundary. Within those boundaries, however, staffing and processing speed — not legal standards or rejection-margin reforms — target the constraint that shows up in the data. Third, the paper develops capacity-time as a measurement framework that translates timing-margin frictions into a unit comparable to the value of the delayed investment. The framework lets a delay coefficient be expressed as gigawatt-years of delayed capacity and as the dollar value of delayed generation, allowing administrative-capacity costs to be compared directly with the marginal value of additional agency resources. It connects to work on regulatory delay in infrastructure (Liscow, 2025), energy-market frictions (Davis, Hausman, and Rose, 2023; Johnston, Liu, and Yang, 2023; Gorman et al., 2024), and environmental cleanup (Sigman, 2001).

The closest energy-market evidence studies interconnection queues rather than environmental permitting. Johnston, Liu, and Yang (2023) show that longer interconnection queues increase waiting times and that interconnection costs help explain project withdrawal. Gorman et al. (2024) document the scale of the grid-connection bottleneck using project-level interconnection data, including rising connection times and high withdrawal rates. I use the same broad project-development setting, but shift the source of variation from grid-connection rules and costs to administrative workload pressure in permitting offices. This lets the paper ask whether a public-sector capacity constraint affects observed projects by stopping them, delaying them, or changing when their output arrives.

The design turns on a specific comparison. It does not compare busy offices to quiet offices in levels. It compares nearby projects, often on opposite sides of the same Corps boundary and entering in the same year, when one assigned office has unusually high non-energy workload relative to its own recent history. The main identifying concern is correspondingly specific. A contemporaneous shock on one side of a boundary could both raise non-energy regulatory workload and independently slow solar development. The empirical checks target that concern. Future workload is not associated with current delay in the timing audit. Workload is not systematically associated with project size, local energy entry, or local employment. Record-level checks reconstruct the workload measure after removing nearby non-energy Corps records, and the result survives. Linked Corps records point to jurisdictional-determination review during queue spells. Taken together, the evidence

supports an administrative-capacity interpretation. The result is not merely that busy offices are slower. It shows that the estimates are consistent with limited state capacity imposing economically meaningful time costs even when observed projects eventually build.

2 Motivating Facts

Multi-year waiting is not unique to the Corps solar setting. Figure 1 summarizes two descriptive benchmarks. In Council on Environmental Quality (CEQ) data on 1,903 final Environmental Impact Statements (EISs) completed from 2010 through 2024, the median time from Notice of Intent to final EIS is 3.1 years, and the 90th percentile is 8.3 years (Council on Environmental Quality, 2025). The 175 Corps-led EIS reviews in the same CEQ data have a median of 3.9 years and a 90th percentile of 10.2 years. In the LBNL solar queue, the median observed time from queue entry to operation or withdrawal is 1.6 years, with a 90th percentile of 5.1 years.

These series are not estimates of the same regulatory process. EIS timelines measure formal federal environmental review for projects that require an EIS, while the solar queue measures a broader project-development spell. The purpose is descriptive: public review and project-development pipelines often involve multi-year waiting. The empirical design below asks whether one measurable source of public-sector capacity pressure contributes to that broader delay problem.

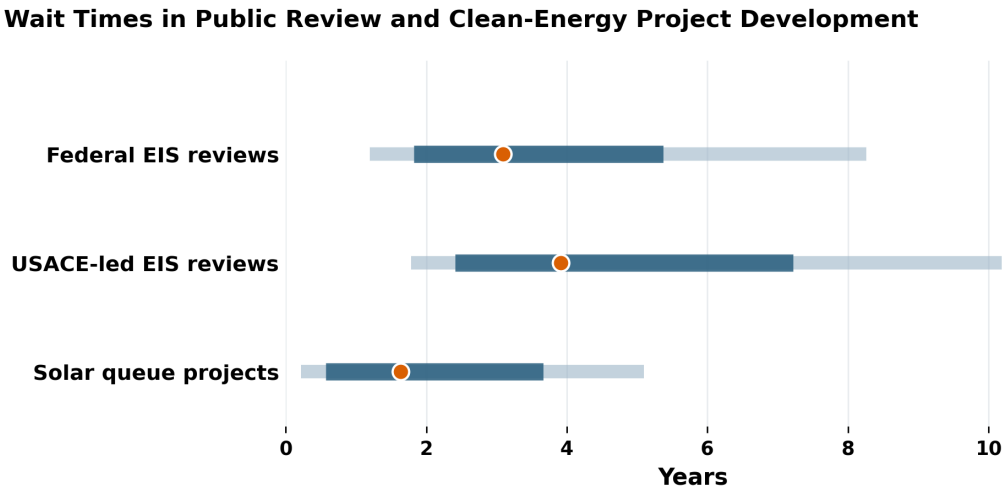

Figure 1: Motivating facts on permitting and project-development delay

Notes: Thin bars show the 10th to 90th percentile range. Thick bars show the 25th to 75th percentile range. Dots show medians. Federal and Corps-led Environmental Impact Statement reviews use the Council on Environmental Quality's 2025 timeline workbook for final statements completed from 2010 through 2024. Solar queue projects use Lawrence Berkeley National Laboratory solar projects with observed queue-entry and operation or withdrawal dates.

3 Data

The project data come from the LBNL interconnection queue. I use solar projects with improved coordinates, meaning geocodes that are more precise than the county centroid available in the raw queue data. For each project I observe queue-entry dates, final queue outcomes, project size, technology, and whether the project becomes operational or is withdrawn. The main outcome is the number of days from queue entry to final resolution. I also study an indicator for resolution within 365 days and an indicator for whether the project is ultimately built.

The Corps is useful for this question because its permitting authority is organized through local regulatory districts. These offices process records related to wetlands, waterways, jurisdictional determinations, and permit decisions. The relevant district is therefore an administrative assignment, not simply a state or county attribute. A project near a district boundary can be close to similar projects on the other side of the boundary while being served by a different regulatory office.

I assign each project to a Corps regulatory district using its coordinates and construct a border sample around Corps district boundaries. The preferred sample includes solar projects within 50 kilometers of a true neighboring Corps district boundary. This restriction removes cases where the nearest district boundary is only an exterior or coastal boundary. The resulting sample contains 2,118 solar projects, including 1,809 projects with observed final queue outcomes.

The main explanatory variable is a district-year non-energy regulatory-workload shock. For each district-year, I count Corps public regulatory records with non-energy keywords, compare that workload to the district's trailing three-year history, and standardize the deviation within year. The trailing-history comparison removes persistent differences in office scale, while the within-year standardization removes national changes in regulatory-record activity. The non-energy filter removes records with solar, wind, transmission, battery, substation, and other directly energy-related terms, so the variation reflects broader office pressure rather than demand from the same solar queue.

The annual measure should be interpreted as office pressure over a sustained district-year period rather than as a precisely timed contemporaneous shock. Districts move through multi-year high-pressure and low-pressure phases, and the measure captures the state of administrative pressure facing projects that enter each district-year. Sub-annual diagnostics support this interpretation: monthly and entry-quarter variants of the measure yield similar headline coefficients, but a 12-month-after-entry window does not survive a live-spell audit (its apparent association with delay disappears once the sample is restricted to projects whose post-entry window falls within their actual review spell). The persistent-phase structure of district workload is the right frequency at which to read the measure.[2]

[2]To make concrete what such a phase looks like: in 2012–2013 the Jacksonville district recorded 1,785 non-energy records, of which 1,138 were Letter of Permission records. The timing aligns with a Corps announcement on May 9, 2012 that two heavily used Florida regional general permits for docks (SAJ-20 and SAJ-33) had expired while endangered-species consultation continued, and that the district would process work through nationwide permits or Letters of Permission instead (USACE Jacksonville, May 9 2012). A related 2013 episode is the reissuance of programmatic permit SAJ-91 for Cape Coral docks and seawalls following smalltooth-sawfish critical-habitat consultation; the announcement explicitly framed programmatic permits as a way to streamline processing by reducing

The Corps regulatory records are used to measure office pressure, not to define the project sample. This distinction is important. The LBNL queue provides a broad set of proposed solar projects and their final outcomes. The regulatory records provide a district-year measure of administrative workload. The empirical question is whether queue projects assigned to offices with unusually high non-energy workload take longer to resolve. The design does not require every queue project to have a matched Corps record.

4 Empirical Design

The estimating equation is:

$$y_{idst} = \beta\, Shock_{dt} + X_i'\gamma + \alpha_d + \lambda_{st} + \pi_b + \varepsilon_{idst}, \tag{1}$$

where i indexes projects, d indexes Corps districts, s indexes states, t indexes queue-entry years, and b indexes nearest Corps boundary pairs. The baseline includes district indicators, state-by-year indicators, project size, and county-level controls. The preferred local designs add either county indicators, boundary-pair indicators, or boundary-pair-by-year indicators. Standard errors are clustered by Corps district.

Table 1: Identifying variation in the USACE border design

Design object	Support	Role in identification
Preferred border universe	2,118 projects; 36 districts; 70 border pairs	Restricts comparison to solar projects near true neighboring USACE district boundaries.
Terminal-delay estimating sample	1,344 projects; 31 districts; 54 border pairs	Compares terminal projects near the same border after district, state-year, and border-pair controls.
Border-pair-year comparisons	97 varying pair-years; 1,082 projects	Compares projects near the same district boundary in the same entry year but assigned to different offices.
Build-out estimating sample	1,613 projects; 36 districts; 69 border pairs	Tests whether the same office-pressure shock predicts ultimate completion, not only delay.
Observed workload exposure	132.7 GW with nonmissing shock; 46.9% of MW above normal	Shows that the shock is economically relevant in the proposed-capacity sample.

Notes: The preferred border universe is solar projects within 50 km of a true neighboring USACE district boundary. Estimating-sample rows use the complete-case samples from the paper's local-control regressions. A varying pair-year contains projects assigned to at least two USACE offices and at least two values of the workload shock.

The identifying comparison is local and administrative. Projects near the same Corps district boundary face similar geography and regional conditions, but their assigned regulatory office can experience different workload pressure. A boundary-pair-by-year cell is the set of projects that entered the queue in the same calendar year and that fall within 50 km of the same Corps district

separate approvals (USACE Jacksonville, Feb. 28 2013). These are illustrative of the kind of administrative event that pushes work back into Corps offices and shows up as a non-energy workload spike in the public records. They are reported here as institutional context for what spikes are, not as a stand-alone causal case.

boundary, but that are administratively assigned to different districts (one on each side of the boundary). A boundary-pair-by-year fixed effect therefore holds constant any condition shared by all projects in that cell. Table 1 summarizes the support for this comparison. The preferred border universe contains 2,118 solar projects near 70 Corps boundary pairs. The final-resolution estimating sample includes 1,344 projects, and the most demanding boundary-pair-by-year design uses 97 pair-years in which projects are assigned to at least two offices and experience different workload shocks.

The border design disciplines two common sources of confounding. First, renewable projects are not randomly located. Solar projects cluster where land, interconnection opportunities, and local development conditions are favorable. Comparing projects near the same administrative boundary reduces the role of broad spatial differences in these fundamentals. Second, permitting workloads can move with local economic conditions. State-year indicators absorb statewide policy and market shocks, while boundary-pair-by-year indicators compare projects near the same boundary in the same entry year.

The identifying assumption is not that office workload is random in levels. It is that, after these controls, remaining deviations in non-energy workload do not coincide with same-side, same-year shocks that would slow solar projects for reasons other than administrative pressure. District indicators absorb permanent differences across Corps offices. State-year indicators absorb statewide policy, market, and queue conditions in a given entry year. Boundary-pair indicators compare projects near the same administrative boundary. Boundary-pair-by-year indicators go further by comparing projects near the same boundary in the same entry year. To explain the result, residual confounding would therefore have to be local to one side of a Corps boundary, contemporaneous with queue entry, correlated with an office's deviation from its own recent non-energy workload trend, and directly related to solar project delay. The timing, falsification, and mechanism checks below evaluate this threat directly.

This is why the shock is constructed from non-energy regulatory records. If the workload measure simply counted solar-related permit activity, the coefficient could reflect solar demand itself: offices with more solar projects would mechanically have more solar permitting records and longer queues. The non-energy measure captures office pressure coming from other work. The resulting coefficient captures the association between overall administrative pressure and project delay, rather than a mechanical comparison of offices with more or fewer solar filings.

The identification is therefore not an external instrument for office workload but the boundary-pair-by-year comparison itself. Appendix Table 13 reports the field-funding and caseload series behind the state-capacity interpretation. The series is descriptive context: real Corps Regulatory Program funding fell over a period in which observed non-energy workload did not fall commensurately. It is not part of the identifying variation. The contribution of this paper is to measure what limited administrative capacity costs along the timing margin, conditional on observed queue entry, using a within-boundary-pair comparison that holds local geography, market conditions, and entry-year shocks fixed.

5 Results

Table 2 reports the main results. The office-pressure shock is associated with longer resolution times in every local-control specification. The coefficient on log days to final resolution is 0.172 in the boundary-pair-by-year specification, the most demanding comparison because it compares projects near the same administrative boundary in the same entry year. The more conservative boundary-pair estimate is 0.117. These estimates imply 144 and 96 additional days at the median resolution time.

Table 2: U.S. Army Corps office-pressure shock and solar project outcomes

Outcome	Base	County	Border	Pair-year	Pair-year + county
Log days to final resolution	0.109**	0.162**	0.117**	0.172***	0.156**
	(0.045)	(0.060)	(0.045)	(0.047)	(0.057)
Final resolution within 365d	-0.034**	-0.043	-0.038**	-0.051***	-0.038
	(0.016)	(0.027)	(0.015)	(0.010)	(0.032)
Built	0.007	0.000	0.004	-0.001	0.002
	(0.015)	(0.022)	(0.015)	(0.019)	(0.029)
Observations	1,344	1,351	1,344	1,344	1,351
Extra days at median spell	88	135	96	144	130

Notes: The sample is solar projects within 50 km of a true neighboring U.S. Army Corps district boundary. The shock is the district-year non-energy regulatory-workload deviation from the district's trailing three-year history, standardized within year. All specifications include district indicators and project size. Base, county, and border columns include state-year indicators. Standard errors are clustered by U.S. Army Corps district. *, **, and *** denote significance at the 10, 5, and 1 percent levels.

The timing pattern is concentrated in the delay estimates. The same shock is associated with a lower probability that a project reaches a final queue outcome within 365 days in the baseline, boundary-pair, and boundary-pair-by-year specifications. In contrast, the built coefficients are close to zero and statistically insignificant across all columns. This pattern is consistent with administrative pressure slowing project resolution rather than changing ultimate completion. Queue resolution is broader than direct permit-processing time, but it is the economically relevant project-development spell. Until a built project reaches operation, its capacity is not available to generate electricity. I therefore interpret this pattern as evidence that office pressure is associated with longer project-resolution times, not as a direct estimate of Corps permit duration.

The magnitude is economically meaningful because delay changes capacity-time. The boundary-pair-by-year estimate implies an 18.7 percent longer resolution time, or about 144 additional days at the median resolution time of 768 days. The conservative boundary-pair estimate implies a 12.4 percent longer resolution time, or about 96 additional days. For a median 48 megawatt (MW) solar project, the conservative estimate is roughly 13 MW-years of delayed capacity. This capacity-time translation is the main economic scale of the timing result.

Table 3 translates the timing estimates into capacity-time. Under the table's central assumptions, a 25 percent solar capacity factor and \$50/MWh generation value, one delayed gigawatt-year of solar capacity corresponds to 2.19 million megawatt-hours (MWh) and \$109.5 million in

Table 3: Scale of delayed solar capacity under office-pressure shocks

Counterfactual	Estimate	Active projects	GW-years delayed	Gen. value ($m)	Gen. + CO2 ($m)
One-SD shock	Conservative estimate	233	8.9	979	1,860
One-SD shock	Preferred estimate	233	13.5	1,475	2,802
Above-normal to normal	Conservative estimate	120	3.7	411	780
Above-normal to normal	Preferred estimate	120	5.8	630	1,198
High workload to normal	Conservative estimate	35	2.8	309	587
High workload to normal	Preferred estimate	35	4.4	479	911

Notes: Central assumptions use a 25 percent solar capacity factor (EIA Electric Power Monthly utility-scale PV average, 2022-2024), $50/MWh generation value (EIA Annual Energy Outlook 2024 reference-case wholesale electricity price, 2025-2030 window, 2023 dollars), 0.45 tons CO2/MWh (US average grid intensity per EPA eGRID 2022, rounded), and $100/ton carbon value (intermediate between the IWG 2021 interim SCC at 3 percent discount rate and the EPA 2023 SCC at 2 percent discount rate). Low and high scenarios bracket each parameter. The policy counterfactual moves positive workload deviations to zero in the built-terminal border sample. These are capacity-time translations, not lifetime lost-output or welfare estimates. Full active-MW support is reported in the underlying CSV.

delayed generation value. Applying the conservative estimate to built projects with final queue outcomes implies 8.9 GW-years delayed under a one-standard-deviation workload shock. The preferred boundary-pair-by-year estimate implies 13.5 GW-years. A more policy-like counterfactual moves observed above-normal workload years back to normal workload. In that exercise, the conservative estimate implies 3.7 GW-years of delayed built capacity and the preferred estimate implies 5.8 GW-years. These correspond to about $411 million to $630 million in delayed generation value under the same assumptions. Adding a carbon-value calculation raises the scale to about $780 million to $1.2 billion. These are capacity-time translations of delayed generation, not lifetime lost-output estimates. The dollar figures are roughly proportional to the assumed capacity factor: a 20 percent factor (representative of less-sunny regions) reduces them by about 20 percent, while a 30 percent factor (representative of utility-scale Southwest installations) raises them by about 20 percent.

The headline estimate is also stable across bandwidths from 25 to 100 km of the nearest neighboring district boundary; the 75 and 100 km samples add few projects beyond the 50 km headline cut, so the design is effectively saturated at the headline bandwidth (Appendix Table 10).

I also match the LBNL projects to EIA plant records as an exploratory check outside the interconnection-queue data. Appendix Table 15 reports this supplementary exercise. In the matched built-project sample, the office-workload shock is associated with later recorded operation, providing suggestive support for the timing interpretation outside the queue data. A one-standard-deviation increase is associated with a 0.18-year increase in queue-to-EIA operation time, with $p = 0.079$. The matched capacity ratio shows no relationship with the shock. That null is mechanically conservative because the matching algorithm uses queued-versus-observed capacity similarity as a match criterion. I therefore treat the plant-record exercise as suggestive operation-data evidence rather than as validation of the main timing estimate. Post-operation capacity-factor

estimates are reported in the appendix only and are not part of the headline claim. Capacity factors reflect solar resource, curtailment, interconnection constraints, plant quality, and market conditions, and the post-operation coefficients are sensitive to dropping the Jacksonville and Savannah districts (Appendix Table 16).

6 Credibility and Mechanism

Table 4 brings the main credibility checks into the short paper. The central issue is whether the remaining variation in non-energy office workload reflects administrative pressure rather than local development conditions, project composition, or a few influential offices. The record-level leave-local check, described later in this section, is the most direct test of this concern. It reconstructs the workload measure after removing all Corps records within 50 km of the focal solar site so that the shock cannot mechanically reflect permitting activity at the solar project itself.

Table 4: Credibility and mechanism checks

Check	Evidence	Interpretation
Influence and random reassignment	Leave-one-district estimates remain positive in 31/31 cases; permutation positive-tail p-value is 0.016.	The main delay estimate is not driven by one district or by arbitrary entry-year reassignment.
Timing	One-at-a-time t-1 coefficient 0.175 ($p = 0.011$); joint future positive-significant count = 0.	Future office pressure does not explain current project delay.
Project composition	Project-name and size checks have minimum local-control p-value of 0.108.	High-workload offices are not simply receiving observably larger or more documented projects.
Local development pressure	Solar/wind entry and capacity checks have no 5 percent rejections; county employment and establishment checks have no 5 percent rejections (minimum $p = 0.131$).	Observable local entry and county economic conditions do not appear to drive the shock.
Record-level leave-local workload	Using individual Corps records, the 50 km leave-local shock predicts log resolution time (0.132, $p = 0.011$) but not build-out (-0.000, $p = 0.991$). Mean 50 km local share is 0.064.	The result is not driven by non-energy permit records located near the solar project.
Corps record linkage	1,587 projects queried; 10 strong/medium solar links; all 10 are jurisdictional determinations. All 10 links occur after queue entry; 6 of 7 with observed final dates fall inside the queue spell.	Sparse linked records align with Corps review during the project-development spell.

Notes: The table collects paper-facing checks for the office-workload design. Permutation p-values are based on within-entry-year reassignment of district-year workload shocks. The leave-local row rebuilds workload from individual Corps records and removes records within 50 km of each solar project from both current workload and the lag baseline.

Sub-component reduced form on independent caseloads. A direct check on the channel uses three components of the non-energy caseload that are generated by parties unrelated to solar development: modifications to projects originally authorized under the Water Resources Development Act (WRDA), which are initiated by infrastructure owners (ports, cities, utilities) modifying existing Corps locks, dams, and levees; cases flagged by the Corps as requiring substantive review; and a residual "other" category. None of the three is generated by solar or wind project sponsors. WRDA modifications in particular have no plausible link to solar siting: the projects are pre-existing federal water-resource infrastructure whose modification is driven by their operators, not by interconnection-queue activity. This sub-component reduced form is corroborating evidence on the channel; it is not used as an instrument. Internal probes of WRDA-wave Bartik instruments did not produce both a strong first stage and a clean second stage in the boundary-pair design. Appendix Table 8 reports the reduced-form association between each case-type caseload and solar delay at the same district. With district and boundary-pair-by-year fixed effects, a one-log-unit increase in WRDA caseload is associated with a 21 percent longer time to final queue resolution (coefficient 0.210, cluster-robust standard error 0.054, $p < 0.001$). The substantive-review and other-case-type coefficients are 0.157 and 0.166 (both $p < 0.001$). Under district plus state-year fixed effects, the three coefficients are 0.166, 0.113, and 0.111 (all $p \leq 0.05$). All six estimates are positive and individually significant. Districts with more independent non-energy caseload process solar projects more slowly, even when the caseload comes from cases that solar developers play no role in generating. The pattern supports the same channel as the headline border design.

District trends and two-way clustering. The headline shock is constructed as a deviation from the trailing three-year mean within district. District fixed effects absorb the level mean across years but do not absorb a district-specific linear trend. To address this, Appendix Table 9 re-estimates the headline boundary-pair-by-year specification with district-specific linear time trends as additional controls. The headline coefficient remains 0.136 (cluster-robust standard error 0.049, $p = 0.009$), about 21 percent smaller than the 0.172 baseline but still significant at the 1 percent level. The same table reports the headline coefficient with two-way clustering (district + year, Cameron-Gelbach-Miller decomposition), which raises the standard error from 0.047 to 0.059 ($p = 0.013$) but does not change the point estimate. The result is robust to both critiques.

Pre-trend test. The proper falsification check asks whether *future* shocks are associated with current project delay, not whether past shocks are. I lead with the pooled solar-and-wind specification because the placebo concern is not technology-specific (whatever local condition might independently raise both later workload and current delay would do so for projects of any technology). The pooled joint F-test on the two leads ($t+1$ and $t+2$) is 0.75 ($p = 0.483$, Appendix Table 6). On the smaller solar-only sample, the same test gives $F = 2.97$ ($p = 0.069$), close to but above conventional cutoffs (Appendix Table 12). The individual lead coefficients are small and within sampling noise of zero. The $t-2$ lag, by contrast, carries a positive and significant coefficient, consistent with persistent workload pressure. To address the concern that the contemporaneous estimate might partly proxy for serially correlated backlog, I orthogonalize the contemporaneous

shock against its own two-year lag at the district-year level and re-estimate the headline regression using the residual (Appendix Table 7). The R^2 of the contemporaneous shock on the two-year lag is 0.016, indicating that the contemporaneous shock is already largely orthogonal to its lag. On the 1,192 boundary-pair projects with a non-missing two-year lag, the headline coefficient is 0.239 (cluster-robust standard error 0.037, $p < 0.001$). Replacing the contemporaneous shock with its orthogonalized residual yields 0.265 (cluster-robust standard error 0.040, $p < 0.001$).

Project composition. If high-workload offices simply receive larger, more documented, or otherwise harder projects, the shock should be associated with observable project characteristics. The shock is not associated with project size or with whether the LBNL row has a project name in the same local-control designs. The minimum local-control p-value across those project-level falsification checks is 0.108.

A stress test asks whether the delay result is driven by a single Corps district or by the particular assignment of district-year shocks. The leave-one-district estimates remain positive throughout the local-control ladder. Within-entry-year permutation tests put the preferred estimate in the upper tail of the permutation distribution. This exercise supports the influence and specification stability of the main delay result.

The tightest workload-construction check uses individual Corps records. I harvest 429 Corps district-year cells covering the solar border-sample project years and the prior years needed for the lag baseline. This produces 2,118 project-local rows. Of those, 1,669 have at least two harvested lag years, and 1,367 are complete for the 50 km leave-local delay regression. For each project, I remove non-energy Corps records within 50 km from both current workload and the trailing baseline. The resulting leave-local shock is still associated with log resolution time, with a coefficient of 0.132 and $p = 0.011$, and it has no relationship with build-out. This check addresses a concrete threat. The main result is not driven by other permit records located near the same solar site.

A broader local-pressure audit reaches a similar conclusion. The shock is not associated with local solar or wind entry counts or MW in the border universe, and it is not associated with county-year total or construction employment and establishments in the Quarterly Census of Employment and Wages for the solar border sample. Across these core checks, there are no 5 percent rejections and the smallest p-value is 0.131. The one imbalance is that the shock is associated with the district-year share of entries with project names. To check whether this compositional pattern explains the null build-out result, I re-estimate the boundary-pair-by-year built regression on the subsample of solar projects with a project name in the LBNL row. The panel contains 582 named projects; 369 have non-missing values of the workload shock and enter the regression. The 213 named projects without shock data are concentrated in the pre-2012 cohort and have a markedly lower built rate than the included sample, so the conditional analysis is implicitly post-2012. On that subsample the conditional built coefficient is 0.003 (standard error 0.076), with a 95 percent confidence interval roughly spanning -0.15 to $+0.15$. The estimate is not distinguishable from zero, but the precision is too coarse to rule out a moderate negative effect. The log-delay coefficient on the same named subsample remains positive and significant at 0.118 ($p = 0.038$). The compositional

imbalance therefore reinforces the interpretation of workload as overall administrative pressure, and is consistent with, rather than overturning, the null build-out finding.

Taken together, the checks clarify what the estimates measure. Conditional on observed queue entry, workload is associated with the speed at which projects move through the development process. It is not associated with whether those projects ultimately build (the boundary-pair-by-year built coefficient is -0.001 with a 95 percent confidence interval of roughly -0.038 to $+0.036$, against a sample mean build rate of about 30 percent, so the design rules out reductions in the build probability larger than about 4 percentage points). The plant-record match does not show robust evidence that delayed built projects are smaller. This timing margin is directly connected to staffing, triage, and process-improvement reforms.

7 Conclusion

Public agencies shape private investment not only by choosing rules, but also by determining how quickly projects move through those rules. This paper studies that implementation margin in clean-energy development. In local comparisons near Corps district boundaries, solar projects are associated with longer resolution times when their assigned regulatory office is unusually busy. A one-standard-deviation increase in non-energy office workload is associated with a three-to-five-month increase in resolution time, with no detectable change in whether observed projects ultimately build.

These estimates identify an implementation margin that is narrower than the total welfare effect of permitting requirements, but central to policy: the speed at which projects move through existing rules. Limited state capacity is associated with delayed timing of clean-energy investment while leaving ultimate build-out and matched capacity largely unchanged. For policy, this points toward a cost that is often missed in build/no-build analyses: administrative frictions can be costly because they delay capacity, not only because they stop projects. Reforms that target this margin — additional reviewer staffing, Section 214 funded positions, triage and process-improvement programs, and modernization of district workflows — are evaluable using the capacity-time framework introduced here, because their benefits accrue through the same channel the design identifies.

The capacity-time magnitudes are large relative to the resources of the program that processes these reviews. Real Corps Regulatory Program field funding fell from an FY2016 index of 106 to an FY2024 index of 88, a decline of about 17 percent in real terms; by FY2023, before the low 2024 observed-record count, the decline was about 19 percent. Applying the border-sample timing estimates proportionally to the full LBNL solar interconnection queue implies annual delayed-generation value on the order of \$200 million to \$300 million per year of the analysis window. This queue-wide figure is illustrative rather than separately identified: the boundary-pair design only identifies within-pair variation among projects within 50 km of a true neighboring district boundary, and projects in the interior of large districts are not part of the identifying variation. The annual real shortfall in Corps Regulatory Program field funding relative to the FY2016 baseline is on the

order of $35 million to $40 million. This is an illustrative scale comparison, not an estimate of the marginal return to additional agency funding. Even with that scope, the comparison shows why administrative capacity belongs in the economics of infrastructure investment. Building more clean infrastructure requires not only choosing the right legal standard, but also having administrative systems that can process projects quickly enough for those standards to matter.

A Appendix Tables

The current draft is built from the following reproducible exhibit files:

- main local-control table,
- timing table and compact credibility table,
- project-composition and entry-falsification table,
- local-development pressure falsification table,
- record-level leave-local workload check,
- regulatory-record linkage and jurisdictional-determination mechanism table,
- regulatory-record timing validation,
- identifying-variation table,
- motivating-facts table on broad review and queue delay,
- delay-value and capacity-time translation table,
- capacity-online exploratory screen,
- EIA-923 post-operation output check and robustness table,
- EIA-860 match-quality audit for eventual matched capacity,
- referee stress tests for the main delay coefficient.

Table 5 and Figure 2 report leave-one-district and permutation stress tests for the main delay coefficient.

Table 5: Referee stress tests for the office-pressure delay coefficient

Specification	Estimate	Std. error	Leave-one range	Positive drops	Placebo p_+	Placebo p_{2s}
Baseline	0.109**	(0.045)	[0.043, 0.150]	31/31	0.067	0.148
County FE	0.162**	(0.060)	[0.139, 0.191]	31/31	0.033	0.061
Border-pair FE	0.117**	(0.045)	[0.059, 0.166]	29/29	0.088	0.152
Border-pair-year FE	0.172***	(0.047)	[0.130, 0.231]	31/31	0.016	0.035
BP-year + county FE	0.156**	(0.057)	[0.130, 0.184]	31/31	0.058	0.105

Notes: The outcome is log days from queue entry to terminal resolution in the solar neighbor-border sample. Leave-one range reports the minimum and maximum coefficient after dropping one USACE district at a time. Positive drops counts the number of valid leave-one estimates that remain positive. Placebo p_+ is the positive-tail probability and placebo p_{2s} is the two-sided coefficient-tail probability from within-entry-year permutations of district-year workload shocks. All specifications use the same controls as the main local-control table. Standard errors are clustered by USACE district. *, **, and *** denote significance at the 10, 5, and 1 percent levels.

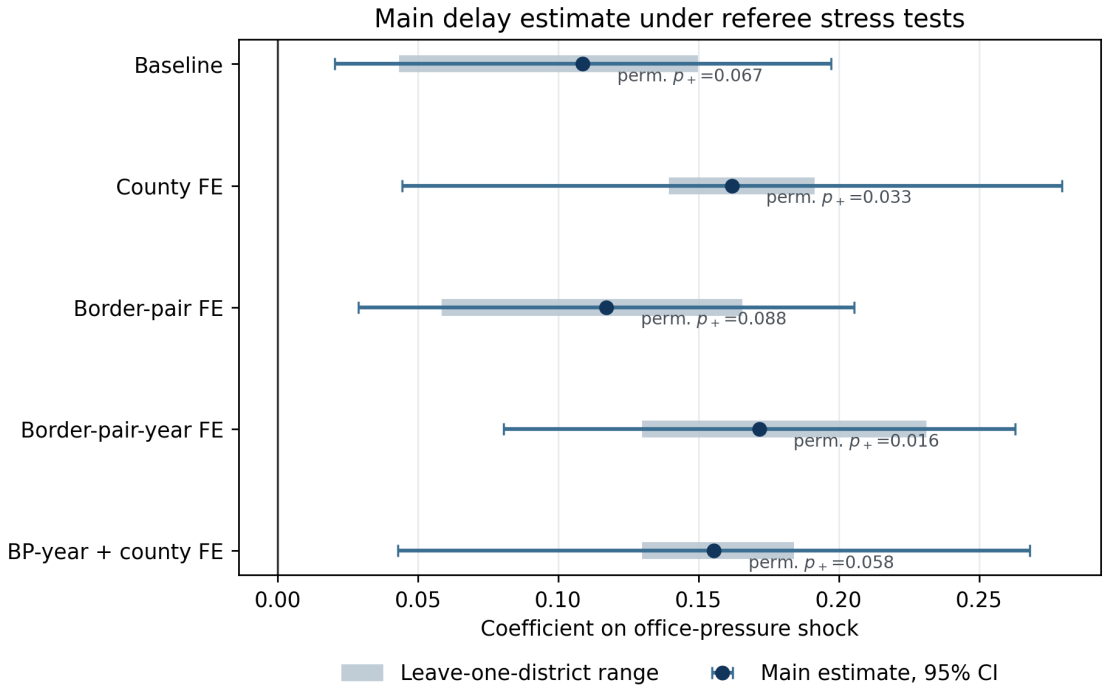

Figure 2: Referee stress-test ladder for the main delay coefficient

Notes: Points are the main local-control estimates for log days to final queue resolution. Horizontal whiskers show 95 percent confidence intervals using district-clustered standard errors. Thick shaded bars show the range of estimates obtained by dropping one USACE district at a time. Permutation labels report positive-tail probabilities from within-entry-year permutations of district-year workload shocks.

Appendix Table 15 reports the EIA-based output check, and Table 16 reports the capacity-factor robustness checks discussed in the text. Appendix Table 13 documents the Corps Regulatory Program field-funding and observed caseload trends discussed in the empirical-design section. Appendix Tables 10, 11, 12, 6, 7, 8, and 9 (with companion Figure 3) report seven additional defensive checks: bandwidth sensitivity at 25, 50, 75, and 100 km; wild cluster bootstrap p-values for the headline coefficients given the small number of Corps districts ($G \approx 31$); a distributed-lag specification that tests whether shock leads at $t+1$ and $t+2$ are associated with current delay (the formal pre-trend test); a pooled solar-and-wind version of the same leads-only test that improves statistical power on the placebo question; an orthogonalized-shock specification that residualizes the contemporaneous shock against its two-year lag to address the concern that the headline may be partly driven by autocorrelated workload pressure; a reduced-form workload-sub-component specification that examines whether three case-type caseloads (WRDA modifications, substantive-review cases, and an "other" residual) are independently associated with solar delay at the same district; and a district-trend-and-two-way-clustering robustness check that re-estimates the headline with district-specific linear time trends and with two-way clustering at the district and year levels.

Table 6: Pooled leads-only joint F-test with statistical power improved by pooling solar and wind

	Solar only (headline)	Wind only	Pooled solar + wind
Coef. on $t-2$ shock	+0.212	−0.769	+0.179
Coef. on $t-1$ shock	+0.091	+2.551	+0.063
Coef. on t shock (current)	+0.309	−0.964	+0.107
Coef. on $t+1$ shock (lead)	−0.077	+0.251	−0.048
Coef. on $t+2$ shock (lead)	+0.092	+0.449	+0.053
Joint F, leads = 0	2.97	0.53	0.75
Joint p, leads = 0	0.069	0.597	0.483
Observations	1,001	177	1,178
USACE districts	26	20	27
Boundary-pair-by-year FE	Yes	Yes	Yes
District FE	Yes	Yes	Yes
log MW + county controls	Yes	Yes	Yes

Notes: Each column reports the same five-term distributed-lag specification used for the solar-only headline (Appendix Table 12). The pooled column adds wind projects to the boundary-pair sample with a technology fixed effect. The contemporaneous coefficient for wind is negative (wind does not typically go through Corps wetlands review at the same intensity as solar, so workload pressure is not the binding constraint), which mutes the pooled $t = 0$ coefficient. For the leads-only joint F-test, however, the technology of the focal project does not matter: the placebo concern is whether *future* workload predicts current delay, which would only happen if local conditions independently raise both later workload and current delay — a channel that is not technology-specific. The pooled joint $F = 0.75$ ($p = 0.483$) shows that the leads are jointly insignificant in the larger sample. Cluster-robust standard errors at the Corps district level.

The replication archive records the script-level provenance for all exhibits. The paper-facing build includes generators for the main local-control results, mechanism checks, delay-value trans-

Table 7: Orthogonalized-shock robustness: the headline survives residualization against the $t-2$ lag

	Headline (full sample)	Headline (common sample)	Orthogonalized shock
Workload (z) coefficient	$+0.171^{***}$	$+0.239^{***}$	$+0.265^{***}$
Cluster-robust SE	(0.050)	(0.037)	(0.040)
p (asymptotic)	0.002	0.000	0.000
Observations	1,312	1,192	1,192
USACE districts	30	29	29
Boundary-pair-by-year FE	Yes	Yes	Yes
District FE	Yes	Yes	Yes
log MW + county controls	Yes	Yes	Yes

Notes: The headline column reports the boundary-pair-by-year specification on solar projects within 50 km of a true neighboring Corps district boundary. The orthogonalized column repeats the same regression, replacing the contemporaneous workload shock with the residual from a district-year-level OLS regression of the contemporaneous shock on its own two-year lag. The orthogonalization removes the autocorrelated component, leaving only the within-year innovation. The R^2 of the contemporaneous shock on its two-year lag is 0.016, so the orthogonalized residual retains nearly all the contemporaneous-shock variation. The orthogonalized coefficient is larger than the headline rather than smaller, indicating that the result is not driven by the autocorrelated component of workload. Cluster-robust standard errors at the Corps district level. Stars: $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

Table 8: Reduced-form association between non-energy case-type caseloads and solar delay

	District + BPY FE			District + state-year FE		
	WRDA	Review	Other	WRDA	Review	Other
log caseload	$+0.210^{***}$	$+0.157^{***}$	$+0.166^{***}$	$+0.166^{***}$	$+0.113^{***}$	$+0.111^{**}$
	(0.054)	(0.035)	(0.036)	(0.046)	(0.032)	(0.042)
Observations		1,397			1,397	
USACE districts		31			31	
District FE		Yes			Yes	
log MW + county controls		Yes			Yes	

Notes: Reduced-form OLS regressions of log days to final queue resolution on the log of one plus the count of district-year non-energy cases in three sub-components: modifications to projects authorized under the Water Resources Development Act ("WRDA"), cases flagged as requiring substantive review ("Review"), and the residual "Other" category. None of the three case types is generated by solar or wind project sponsors, so each is plausibly exogenous to interconnection-queue developer behavior. The first three columns use district + boundary-pair-by-year fixed effects; the last three use district + state-year fixed effects. Each column is a separate regression with the column header as the right-hand-side variable; all specifications include log project size and four county controls (population, median household income, poverty rate, terrain ruggedness). Cluster-robust standard errors at the Corps district level. Stars: $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

Table 9: District-trend and two-way clustering robustness for the headline

Specification	Coef. on workload shock	SE	N	Districts
Headline	$+0.172^{***}$	(0.047)	1,344	31
+ district linear trend	$+0.136^{***}$	(0.049)	1,344	31
+ district trend + state-year FE	$+0.153^{**}$	(0.072)	1,344	31
Two-way clustering (district + year)	$+0.172^{**}$	(0.059)	1,344	31

Notes: Each row reports the coefficient on the non-energy workload shock from the headline boundary-pair-by-year specification. The Headline row matches the Pair-year column of Table 2. The "+ district linear trend" row adds a district-specific linear time trend (district indicator interacted with q_year_int) as an additional control, addressing the concern that the trailing-mean shock construction could embed district-level trending. The "+ district trend + state-year FE" row adds state-year fixed effects on top. The "Two-way clustering" row reports the headline coefficient with two-way clustering at the district and year levels, computed via the Cameron-Gelbach-Miller (2011) decomposition $V_{2way} = V_d + V_y - V_{d\times y}$ with t reference distribution at $\min(G_d, G_y) - 1$ degrees of freedom. All specifications include district fixed effects, project size, and county controls. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 10: Bandwidth sensitivity for the headline boundary-pair-by-year specification

Outcome	25 km	50 km	75 km	100 km
Log days to terminal	$+0.202^{*}$	$+0.172^{***}$	$+0.171^{***}$	$+0.171^{***}$
	(0.102)	(0.047)	(0.046)	(0.046)
Terminal within 365d	-0.071^{***}	-0.051^{***}	-0.051^{***}	-0.052^{***}
	(0.024)	(0.010)	(0.010)	(0.010)
Built	+0.033	−0.001	−0.001	−0.002
	(0.025)	(0.019)	(0.019)	(0.019)
Observations	803	1,344	1,354	1,361

Notes: Each column reports the headline boundary-pair-by-year fixed-effects regression on the subsample of solar projects within the indicated distance of a true neighboring USACE district boundary. Coefficients are for a one-standard-deviation increase in the non-energy regulatory-workload shock. Standard errors clustered by USACE district. The 50 km column is the headline reported in the main results; estimates at 75 km and 100 km are essentially identical because few additional projects are added beyond 50 km. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 11: Wild cluster bootstrap p-values, headline boundary-pair-by-year specification

Outcome	Coef.	SE (cluster)	p (asymptotic)	p (wild boot.)
Log days to terminal	+0.172	(0.042)	0.000	0.021
Terminal within 365d	-0.051	(0.009)	0.000	0.002
Built	-0.001	(0.017)	0.973	0.971
N	1,344; 31 USACE districts; 999 bootstrap reps			

Notes: Wild cluster bootstrap implements the Cameron, Gelbach, and Miller (2008) procedure with Rademacher weights at the USACE district level and 999 bootstrap replications. The wild bootstrap p-value is the recommended inference for a small number of clusters ($G = 31$ here). Asymptotic p-values are based on cluster-robust standard errors with a Student-t reference distribution. The headline log-delay coefficient remains significant at the 5 percent level under wild cluster bootstrap. The built coefficient remains statistically indistinguishable from zero under both procedures.

Table 12: Distributed-lag pre-trend specification (log days to terminal)

Event-time shock	Coefficient	SE
Shock at $t-2$	$+0.212^{**}$	(0.094)
Shock at $t-1$	$+0.091$	(0.093)
Shock at t (current)	$+0.309^{***}$	(0.075)
Shock at $t+1$ (lead)	-0.077	(0.068)
Shock at $t+2$ (lead)	$+0.092$	(0.060)
Joint test, leads $= 0$	$F = 2.97$	$p = 0.069$
Observations	1,001; 26 USACE districts	

Notes: Single regression that includes the non-energy workload shock at five event times relative to project queue entry ($t-2$, $t-1$, t, $t+1$, $t+2$). All other controls and fixed effects match the headline boundary-pair-by-year specification. The joint F-test on the two leads ($t+1$ and $t+2$) is the formal pre-trend / falsification check: workload deviations after queue entry should not predict project delay if the contemporaneous estimate is causal. The leads are individually statistically indistinguishable from zero and jointly insignificant at the 5 percent level. The current shock effect ($t = 0$) is large and highly significant. The $t-2$ shock is also positive and significant, consistent with persistent workload pressure at the district level. The body addresses this by orthogonalizing the contemporaneous shock against its two-year lag and re-estimating the headline regression (Appendix Table 7); the R^2 of the contemporaneous shock on the two-year lag is only 0.016, so the contemporaneous variation is already largely orthogonal to its lag. Standard errors clustered by USACE district. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 13: USACE Regulatory Program field funding and observed caseload, FY2016–FY2024

FY	Field funding ($M nom.)	Field funding ($M 2018$)	Funding index	Caseload index	Pressure index
2016	195.6	204.6	106	77	72
2017	193.0	197.7	102	113	110
2018	193.0	193.0	100	100	100
2019	192.0	188.6	98	113	115
2020	190.7	185.0	96	86	89
2021	193.0	178.9	93	221	239
2022	196.9	168.9	88	105	120
2023	201.8	166.3	86	106	123
2024	212.0	169.7	88	51	58

Notes: Nominal field funding is the Corps Regulatory Program line labeled "Funding to Districts/Divisions" in the U.S. Army Corps of Engineers Civil Works Budget Justification, Other Business Programs, Regulatory Program pages. Real (2018$) values are computed using annual averages of the Consumer Price Index from the U.S. Bureau of Labor Statistics. Caseload is the count of non-energy regulatory records observed in the Corps' Online Records Module extract used to construct the leave-local workload shock. The caseload and real-field-funding indices are normalized to FY2018 = 100. The pressure index is the caseload index divided by the real-field-funding index, again normalized to FY2018 = 100. The FY2016–FY2024 trend in real field funding is -4.9 million per year ($p < 0.001$). Observed public-record caseload is noisier year to year, but was about 38 percent higher in FY2023 than in FY2016. The implied workload-per-real-dollar pressure tightened over the period when caseload remained elevated.

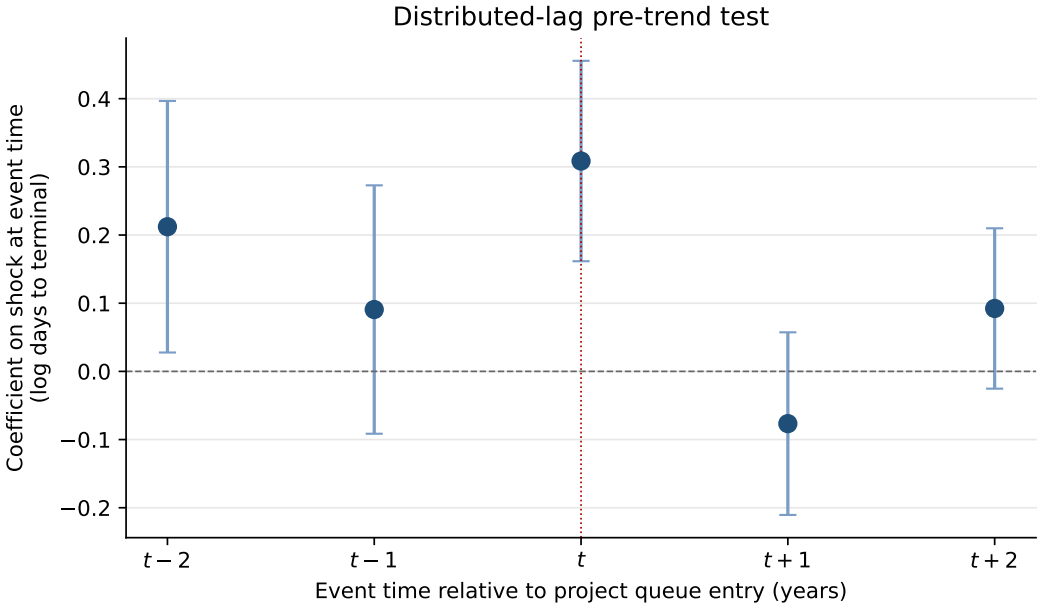

Figure 3: Distributed-lag pre-trend test: log days to final queue resolution

Notes: Coefficients on the non-energy workload shock at event time $t-2$, $t-1$, t (current, dashed vertical line), $t+1$, and $t+2$ from a single distributed-lag regression with boundary-pair-by-year fixed effects. Whiskers are 95 percent confidence intervals using district-clustered standard errors. Sample is solar projects within 50 km of a true neighboring USACE district boundary.

Table 14: Motivating facts on permitting and project-development delay

Series	N	Mean	P25	Median	P90
Federal EIS reviews	1,903	4.1	1.8	3.1	8.3
USACE-led EIS reviews	175	5.2	2.4	3.9	10.2
Solar queue projects	5,229	2.5	0.6	1.6	5.1

Notes: Entries are years. Federal and USACE-led EIS reviews use CEQ's 2025 EIS timeline workbook for final EISs completed from 2010 through 2024. Solar queue projects use LBNL solar projects with observed queue-entry and terminal operation or withdrawal dates. The series are descriptive and measure different stages of public review and project development.

Table 15: EIA post-operation output pilot

Shock	Outcome	Estimate	High-conf. CF	Projects	Obs.
Office workload	Queue-to-EIA operation years	0.18*		276	276
		(0.10)			
Office workload	log(EIA MW / queued MW)	0.036		276	276
		(0.052)			
Office workload	Capacity factor, post-op year 1	-0.007**	-0.007*	221	221
		(0.004)	(0.004)		
Office workload	Capacity factor, post-op year 2	-0.002	-0.004***	175	175
		(0.003)	(0.001)		
Office workload	Capacity factor, post-op year 3	-0.002	-0.003***	158	158
		(0.003)	(0.001)		

Notes: This appendix table reports an exploratory cross-check, not a separate causal design. The sample is solar projects near USACE district borders that match to EIA-860 plant records and have EIA operation years no earlier than queue entry. Capacity factors use EIA-923 annual solar net generation divided by matched EIA capacity times 8,760 hours. Post-operation year 1 is the first full calendar year after EIA's operating year. Main estimates use USACE district and border-pair-year fixed effects, project size, and county controls. High-conf. CF restricts to EIA matches based on tight county-fuel-capacity-year or fuzzy-name matches. The log(EIA MW / queued MW) outcome is mechanically conservative because the EIA-860 match algorithm uses queued-versus-observed capacity similarity as a match criterion. A parallel set of estimates using the Section 214 sponsor-pressure measure was computed but is not shown: those rows yielded opposite-signed estimates on queue-to-EIA operation years, which is consistent with Section 214 activity being demand-driven rather than capacity-driven. Standard errors are clustered by USACE district. *, **, and *** denote significance at the 10, 5, and 1 percent levels.

lation, EIA validation, referee stress tests, motivating facts, the compact credibility table, the record-level leave-local workload check, and the numerical-claim audit.

References

McKinsey & Company. 2023. Unlocking US federal permitting: A sustainable growth imperative.

Bagley, Nicholas. 2019. The procedure fetish. *Michigan Law Review* 118(3): 345–401.

Besley, Timothy, and Torsten Persson. 2009. The origins of state capacity: Property rights, taxation, and politics. *American Economic Review* 99(4): 1218–1244.

Besley, Timothy, and Torsten Persson. 2010. State capacity, conflict, and development. *Econometrica* 78(1): 1–34.

Council on Environmental Quality. 2025. Environmental Impact Statement timelines, 2010–2024.

Finan, Frederico, Benjamin A. Olken, and Rohini Pande. 2017. The personnel economics of the developing state. In *Handbook of Field Experiments*, Vol. 2, edited by Abhijit Banerjee and Esther Duflo, pp. 467–514. Amsterdam: North-Holland.

Davis, Lucas W., Catherine Hausman, and Nancy L. Rose. 2023. Transmission impossible? Prospects for decarbonizing the US grid. *Journal of Economic Perspectives* 37(4): 155–180.

Table 16: EIA post-operation capacity-factor robustness

Sample/specification	Year 1	Year 2	Year 3
Main all matches	-0.007**	-0.002	-0.002
	(0.004)	(0.003)	(0.003)
High-confidence matches	-0.007*	-0.004***	-0.003***
	(0.004)	(0.001)	(0.001)
Add output-year FE	-0.008**	-0.004	-0.002
	(0.003)	(0.002)	(0.002)
Add region-by-output-year FE	-0.009*	-0.005***	-0.003
	(0.005)	(0.001)	(0.002)
Drop Jacksonville/Savannah	-0.028	0.021***	0.042
	(0.025)	(0.003)	(0.036)
Drop Jacksonville/Savannah, high-conf.	-0.038	0.001	0.019
	(0.024)	(0.026)	(0.037)
Size bin: ¡20 MW	0.012	-0.002	0.021
	(0.047)	(0.010)	(0.016)
Size bin: 20-100 MW	-0.005***	-0.004*	-0.002
	(0.001)	(0.002)	(0.002)
Size bin: ¿=100 MW			

Notes: The table reports coefficients on the office-workload shock for EIA-923 solar capacity factors in the first three full calendar years after EIA operation. Capacity factors are not a headline outcome and reflect solar resource, curtailment, interconnection constraints, plant quality, and market conditions. The main row uses district and border-pair-year fixed effects, project size, and county controls. High-confidence matches are tight county-fuel-capacity-year or fuzzy-name EIA matches. Jacksonville/Savannah drops omit SAJ and SAS; the post-operation coefficients are sensitive to this drop (sign flips from negative to positive at year 2 and year 3 in the high-confidence sample), so the capacity-factor pattern should not be over-interpreted. Size-bin rows use queued MW. Blank cells indicate that the split has no valid clustered inference. Standard errors are clustered by USACE district. *, **, and *** denote significance at the 10, 5, and 1 percent levels.

Table 17: Appendix audit and validation checks

Check	Evidence	Interpretation
Numerical audit	145 of 145 paper-facing checks pass after regenerating exhibits.	Draft numbers match generated artifacts.
Main timing estimate	BP-year log-delay coefficient 0.172 ($p = 0.001$); border-pair FE coefficient 0.117 ($p = 0.014$).	The most demanding local comparison predicts slower terminal resolution.
Timing windows	One-at-a-time t-1 coefficient 0.175 ($p = 0.011$); joint future positive-significant count = 0.	Pressure appears over a short administrative cycle; future shocks do not drive the result.
Project composition	Project-name and log-MW placebos have minimum local-control $p = 0.108$.	Main result is not explained by observed project size or name availability.
Entry composition	Entry counts and mean log MW are quiet ($p = 0.528$ and 0.796); project-name share moves ($p < 0.001$).	Use reduced-form framing; the shock is not orthogonal to every entry-year feature.
Local pressure audit	Solar/wind entry and MW checks have 0 5% rejections; QCEW employment and establishment checks have 0 5% rejections (minimum $p = 0.131$).	Observable local entry and county economic pressure do not appear to drive the shock.
Record linkage	1587 projects queried; 10 strong/medium solar links; 10 are JDS.	Sparse linkage, but matched records point to administrative review.
ORM timing validation	10 of 10 strong/medium links occur after queue entry; 6 of 7 with observed terminal dates fall inside the queue spell.	Matched records align with project development timing, but sparse links keep this as validation evidence.
JDS workload proxy	JDS shock coefficients are positive: 0.403 ($p = 0.118$) and 0.705 ($p = 0.052$).	Mechanism proxy points the same way but is too sparse to replace the main shock.
Stakes translation	One-SD shock delays 9.1 GW-years and 20.0 million MWh among eventual built projects.	Scale is economically meaningful but not a welfare estimate.

Notes: This table summarizes generated artifacts used to audit the short-paper draft. The claim audit is produced after regenerating the paper-facing exhibits. JDS denotes jurisdictional determination records in the ORM public data.

Gorman, Will, Julie Mulvaney Kemp, Joseph Rand, Joachim Seel, Ryan H. Wiser, Nick Manderlink, Fredrich Kahrl, Kevin Porter, and Will Cotton. 2024. Grid connection barriers to renewable energy deployment in the United States. *Joule* 9: 101791.

Johnston, Sarah, Yifei Liu, and Chenyu Yang. 2023. An empirical analysis of the interconnection queue. NBER Working Paper No. 31946.

Liscow, Zachary. 2025. Getting infrastructure built: The law and economics of permitting. *Journal of Economic Perspectives* 39(1): 151–180.

Klein, Ezra, and Derek Thompson. 2025. *Abundance*. New York: Avid Reader Press.

Sigman, Hilary. 2001. The pace of progress at Superfund sites: Policy goals and interest group influence. *Journal of Law and Economics* 44(1): 315–344.

Washington Post. Maryland uses AI to fix permitting, create jobs. April 21, 2026.