Permitting Delays, Regulatory Regimes, and Investment: Evidence from the Clean Water Act

William Swanson*

May 1, 2026

Abstract

This paper develops a framework for measuring how changes in federal environmental permitting requirements affect private investment. We use longitudinal variation in Clean Water Act (CWA) jurisdictions—generated by a sequence of Supreme Court decisions and administrative rules between 2015 and 2026—to estimate the effect of regulatory stringency on the time-to-build distribution of planned electric generation projects. Project phases are observed in the Energy Information Administration's monthly 860M database, allowing us to estimate a Discrete Phase-Type (DPH) hazard model in which the CWA regulatory regime shifts each project's exit hazard and approval rate. We find that stricter CWA regulation slows the permitting phase and reduces the probability of a successful exit, while the Sackett v. EPA (2023) decision, which narrowed CWA jurisdiction, reverses these effects. We embed the estimated time-to-build distributions in a structural model of capital investment in which permitting delays raise the user cost of capital and discourage entry. Simulating the Sackett decision for solar energy projects, we estimate a long-run elasticity of investment with respect to a reduction in permitting duration of approximately -0.36.

1 Introduction

Efforts to regulate the use of the nation's waterways are among the more consequential and contested interventions in federal environmental law. The Clean Water Act's contribution to

*Congressional Budget Office, Washington, D.C. (William.Swanson@cbo.gov). This manuscript benefited greatly from the suggestions and guidance from Dorian Carloni, Nicholas Chase, Ron Gecan, Evan Herrnstadt, Joseph Kile, Jeffrey Kling, and Heidi Williams.

The views expressed in this article are solely those of the author and do not necessarily reflect the views, positions, or policies of the Congressional Budget Office or the United States government.

water quality, aquatic habitat, and public health comes at the price of imposing regulatory requirements on developers, such as the obligation to obtain permits before beginning construction near protected waters. This paper provides new evidence on a cost channel that has received relatively little attention: the effect of permitting delays on investment. When a project developer waits for regulatory approval before committing capital to production, resources are tied up during the waiting period without generating a return. Longer or more uncertain permitting process therefore raise the effective cost of capital, discouraging investment at the margin. How large this effect is in practice depends on the length of the delay, the probability of eventual approval, and the rate at which investors discount future returns—quantities that vary across projects, sectors, and regulatory regimes, and that have been difficult to measure convincingly

We estimate these quantities using the variation in Clean Water Act jurisdictional stringency generated by a sequence of Supreme Court decisions and administrative rules between 2006 and 2023 provided in data from Greenhill et al. (2026). The most significant of these are Rapanos v. United States (2006) and Sackett v. EPA (2023), together with the Clean Water Rule (2015), the Navigable Waters Protection Rule (2020), and their subsequent vacaturs and replacements. These changes altered the scope of CWA jurisdiction across locations and over time in ways that were largely driven by legal and constitutional considerations outside the energy sector. This variation provides us with something close to exogenous shifts in regulatory exposure for the electric generation projects in our sample, which we use to identify the causal effect of CWA stringency on project timelines and outcomes.

Our empirical strategy combines two datasets. The first is the Energy Information Administration's monthly EIA-860M database, which tracks the status of all planned and existing grid-connected electric generators with nameplate capacity above one megawatt. Because the EIA reports project stage (permitting or construction) monthly, we directly observe the phase path of each project—a feature that substantially simplifies the likelihood

function relative to settings where project phases are latent. The second dataset is the simulated CWA jurisdictional stringency measures of Greenhill et al. (2026), who apply deep-learning methods to predict the probability of CWA jurisdiction at approximately four million geographic coordinates under each of four regulatory regimes. Matching these measures to the locations of electric generators creates a project-level panel of regulatory exposure.

We use this panel to estimate a two-phase Discrete Phase-Type (DPH) hazard model in which each project's transition matrix $\mathbf{A}^{(l)}$, exit hazard $h_j^{(l)}$, and success probability $s_j^{(l)}$ depend on the prevailing CWA regime l. The treatment effect parameters for the regime, ζ^h, ζ^S, and ζ^s measure the CWA's impact on exit speed, stage-transition routing, and approval probability, respectively. Likelihood ratio tests reject the null of no treatment effect through the hazard and routing channels at conventional significance levels. The results imply that broader CWA jurisdiction lengthens the permitting phase and reduces the probability of successful completion. We estimate that the recent Sackett supreme court decision that narrowed CWA jurisdictions relative to the prior "Rapanos" regime, reduced the average project durations by around 1.75 months and raised the probability of project completion.

We embed these estimated distributions in a structural model of capital investment in which project sponsors queue for permits before purchasing productive capital. The model nests the DPH time-to-build distribution within a neoclassical investment framework: permitting delays enter through the user cost of capital (Hall and Jorgenson, 1967), and the elasticity of investment with respect to the UCC equals the price elasticity of demand for capital (Chodorow-Reich, 2025). Simulating the Rapanos-to-Sackett transition for solar energy projects—the largest technology group in our sample—we estimate a long-run elasticity of investment with respect to a reduction in permitting duration of approximately -0.36. These estimates should be understood as measuring the investment costs of the delay channel only. A full assessment of the CWA's net social impact would require weighing these costs against the environmental benefits of the regulations, which are beyond the scope of this

paper.

Related literature. Our paper relates to several strands of the investment and regulatory literature. On the measurement of regulatory costs, Al-Ubaydli and McLaughlin (2017) use text analysis of the Federal Register to construct indices of regulatory constraints, while Trebbi and Zhang (2022b) infer regulatory costs from compliance wage data. Our approach instead uses quasi-experimental variation in a specific regulatory regime to recover causal effects on a single, well-defined margin: the duration of the permitting queue.

The most closely related paper on the investment-delay channel is Gabriel and Kung (2024), who estimate the effect of approval-time uncertainty on new housing supply in Los Angeles, finding an elasticity of at least -0.54. Our estimate of -0.36 is smaller in magnitude, consistent with electric power generation facing somewhat less elastic demand than housing in the short run. Glancy et al. (2024) and Bennon et al. (2024) also model the option value of investment under regulatory uncertainty, while Johnston et al. (2023) and Rand et al. (2023) examine queuing delays specifically in electricity markets. Our contribution relative to these papers is the use of a phase-type model that separately identifies how regulations affect the permitting and construction stages, the probability of successful completion, and aggregate investment through a calibrated structural model.

Other work on specific environmental statutes includes Greenstone (2002) on the Clean Air Act, Ryan (2012) and Lewis (2019) on the CWA, and Fraas et al. (2021) and Nilson et al. (2026) on the barriers to electric power generation at the federal and state levels.

A version of the structural model described here was used by the Congressional Budget Office to estimate the private investment response to the "Opt-in Fee" provision included in the 2025 Reconciliation Recommendations of the House Natural Resources Committee.[1] Additional details are available in CBO's published cost estimate[2] and CBO's dynamic

[1]Congressional Budget Office (2026), GitHub Permitting-Model, `www.github.com/US-CBO/permitting-model`.

[2]Congressional Budget Office, cost estimate for Reconciliation Recommendations of the House Committee on Natural Resources (May 19, 2025), `www.cbo.gov/publication/61415`.

estimate of the reconciliation act of 2025.[3]

The remainder of the paper proceeds as follows. Section 2 describes the data and the CWA regulatory history. Section 3 presents the structural model of investment. Section 4 develops the DPH likelihood and estimation strategy and Section 5 embeds the DPH distribution in a sector-level structural model. Section 6 calibrates the model and Section 7 presents the structural results. Section 8 concludes and an appendix details the estimation algorithm.

2 Data and Identification

2.1 Electric Generation Projects: EIA 860M

Our primary source of project-level data is the Energy Information Administration's monthly 860M survey, which tracks the development status of all planned and existing electric generators with nameplate capacity above one megawatt that are or will be connected to a local or regional electric grid. The survey is legally required and covers both the planning and construction stages of a project's life, reporting status monthly from initial announcement through commercial operation. We use the 860M records to construct a panel in which each project-month observation records whether the project was in the permitting phase or the construction phase. The sample runs from August 2015, the first month of available data, through February 2026.

Table 1 summarises project characteristics by technology type. Solar photovoltaic projects dominate the sample (5,202 observations), followed by battery storage (1,283) and fossil fuel generation (1,109). Permitting durations vary substantially across technologies: hydro and nuclear projects face median permitting times exceeding 50 months, while solar and battery projects clear permitting in under 10 months on average. These differences reflect both the underlying complexity of each technology's environmental footprint and the scale of capital

[3]Congressional Budget Office, cost estimate for H.R. 1, One Big Beautiful Bill Act (Dynamic Estimate), (June 17, 2025), `www.cbo.gov/publication/61486`.

involved.

Two limitations of the 860M data bear noting. First, there is ambiguity about the precise moment at which projects begin appearing in the database: some developers may delay reporting until a project is relatively mature in its planning stages, which would left-censor early phases. To minimise this concern, we restrict the sample to projects first observed after July 2015 and exclude Alaska and Hawaii (which lack CWA stringency data from Greenhill et al. 2026). Second, the database excludes "behind-the-meter" generation below one megawatt or not connected to the grid, so our sample skews toward utility-scale projects.

Table 1: Electric Generation Summary Statistics by Technology

	Capacity	Duration (Months)		Total
Technology	MW	Permitting	Construction	Observations
Solar	55.2	7.8	6.6	5202
Batteries	93.3	9.3	7.7	1283
Fossil Fuels	114.1	10.5	7.8	1109
Onshore Wind	179.2	9.1	8.9	655
Hydro	23.2	52.2	6.0	124
Other Energy	8.2	10.4	11.9	91
Geothermal	13.3	12.4	3.2	22
Nuclear	70.2	66.1	0.0	13
Offshore Wind	731.0	23.8	7.5	13

Notes: Summary statistics are calculated from EIA 860M for projects in the pre-construction phases of development between August 2015 and February 2026, excluding Hawaii and Alaska. Capacity is average nameplate capacity in megawatts. Duration figures are averages across completed spells in each phase.

2.2 CWA Jurisdictional Stringency: Greenhill et al. (2026)

Our second dataset provides predicted Clean Water Act jurisdictional stringency at approximately four million geographic coordinates across the contiguous United States. Greenhill et al. (2026) train deep-learning models on historical Approved Jurisdictional Determinations (AJDs) issued by the Army Corps of Engineers under past CWA rules to predict, for each coordinate, the probability of being subject to CWA jurisdiction under each of four

regulatory regimes (Rapanos, Clean Water Rule, NWPR, and Sackett). AJDs are legally binding determinations of whether a given site falls under CWA protection; a developer who needs to know whether CWA permits are required may request an AJD from the Corps before beginning a project. The Greenhill et al. predictions therefore reflect the geophysical characteristics—wetness, hydrological connectivity, proximity to navigable waters—that have been legally determinative under each regime.

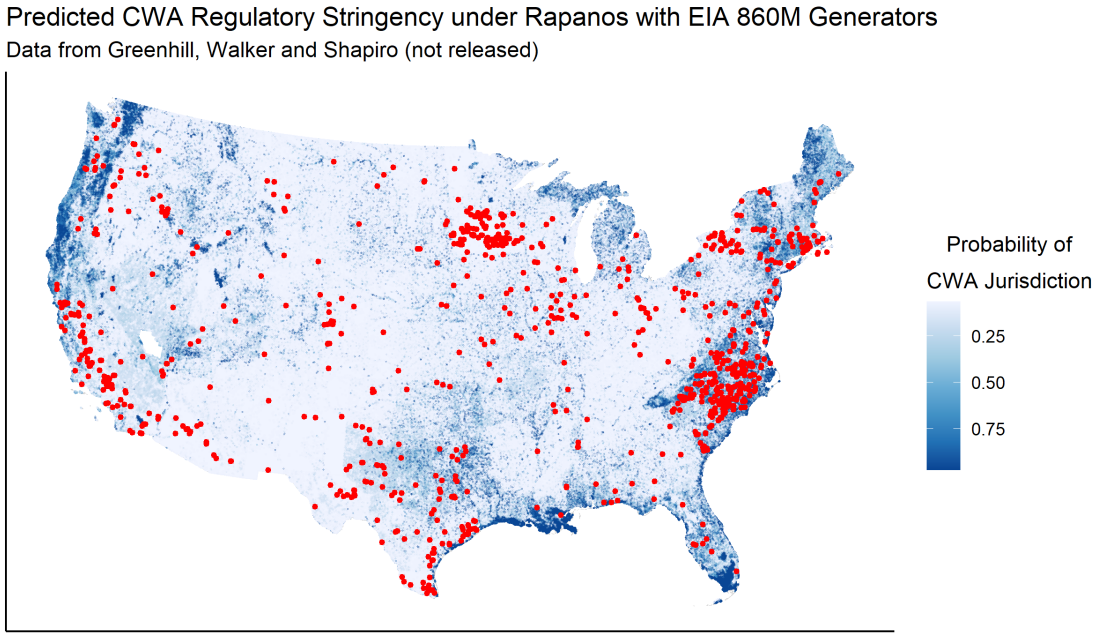
2.3 Constructing Project-Level CWA Exposure

We assign each EIA 860M project a CWA stringency measure using a nearest-neighbor approach. Because both the EIA and Greenhill et al. provide geographic coordinates, we identify the three prediction points closest to each generator and average the predicted jurisdictional probability across those three coordinates. This produces, for each project, a stringency measure $z^{(l)} \in [0, 1]$ under each of the four regimes, interpretable as the predicted probability of being subject to CWA jurisdiction at that location.

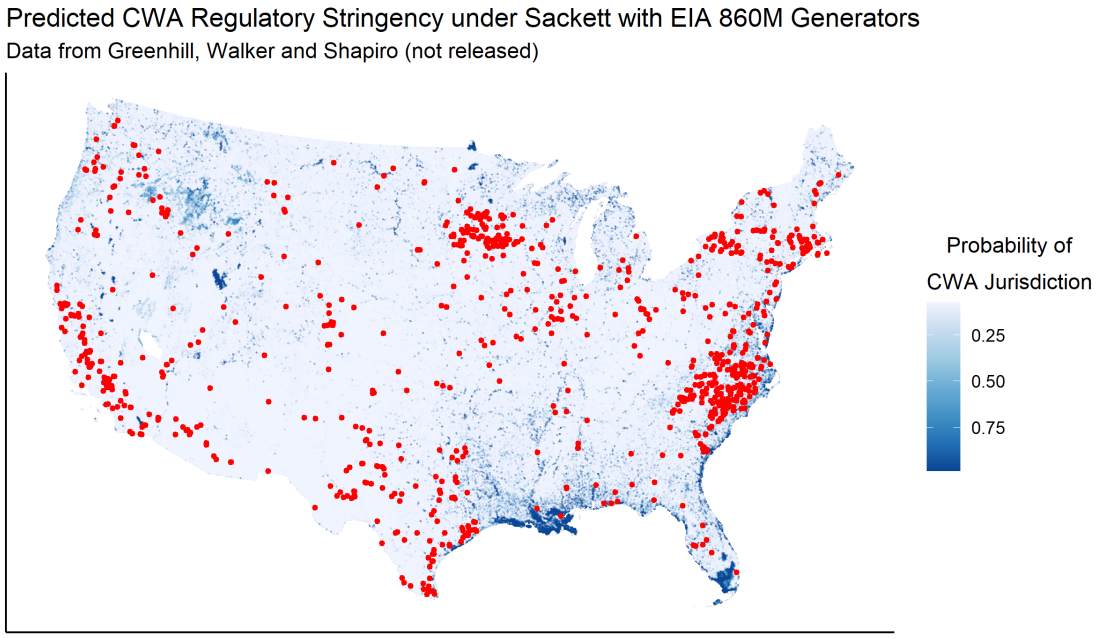
We classify project stages into two groups: permitting and construction. A project is in the permitting phase if its reported status is "Planned for installation," "Regulatory approvals pending, not under construction," or "Regulatory approvals received, not under construction." A project is in the construction phase if its status is "Under construction, less than or equal to 50 percent complete," "Under construction, more than 50 percent complete," or "Construction complete, but not yet in commission." This two-stage classification corresponds to the two phases $j \in \{1, 2\}$ in the DPH model.

2.4 History of CWA Regulatory Regimes

The boundary of CWA jurisdiction—which waters are "Waters of the United States" (WOTUS) and therefore subject to the Act's permitting requirements—has shifted repeatedly since 2006 through a combination of Supreme Court decisions, EPA rules, and state-level litigation. These shifts generate the cross-sectional and temporal variation in regulatory stringency that



(a) The CWA jurisdiction under the Rapanos decision was applied periodically through our sample period from 2015 until February of 2026.



(b) The Sackett decision in 2023 curtailed the CWA stringency for many locations in the United States, particularly the West Coast and the mid-atlantic regions. See Greenhill et al. (2026) for details about the rendering of these data, and alternative maps under NWPR and CWR.

we use for identification.

Rapanos (2006–2015, 2019–2020, 2021–2023). The Supreme Court's 2006 *Rapanos* decision established a "significant nexus" standard: to be subject to CWA, a water body must have a biological, physical, or chemical connection to navigable waters. This became the baseline regulatory framework for most of our sample period.

Clean Water Rule (2015–2019). The 2015 Clean Water Rule (CWR) broadened jurisdiction by pre-classifying certain tributary types—including intermittent and ephemeral streams—as protected without requiring case-by-case nexus determinations. The CWR was immediately contested; thirteen states obtained injunctions and remained under Rapanos guidance throughout the period. By 2019, when the rule was formally repealed, twenty-two states had adopted it while twenty-eight remained under Rapanos.

Navigable Waters Protection Rule (2020–2021). The 2020 Navigable Waters Protection Rule (NWPR) narrowed jurisdiction relative to both the CWR and Rapanos by explicitly excluding ephemeral streams. Colorado remained under Rapanos due to separate state litigation until April 2021; all other states adopted the NWPR in June 2020. The NWPR was vacated by federal courts in August 2021, returning all states to the Rapanos framework.

Sackett (2023–present). The Supreme Court's May 2023 decision in *Sackett v. EPA* substantially narrowed WOTUS by eliminating the significant-nexus test and requiring a continuous surface-water connection to navigable waters. A conforming administrative rule took effect September 8, 2023, though litigation has delayed its full application in twenty-six states.

Table 5 in the appendix summarizes the effective dates and geographic scope of each regime change. Table 2 shows the distribution of projects across regime transitions: the large diagonal entries reflect projects that remained entirely within a single regime, while off-diagonal entries capture the within-project variation we use to identify treatment effects.

Table 2: Project-level Regulation Regime Transition Matrix

From \ To	NWPR	CWR	Rapanos	Sackett
NWPR	2042	0	1198	0
CWR	0	2032	762	0
Rapanos	1184	0	4043	1494
Sackett	0	0	0	3804

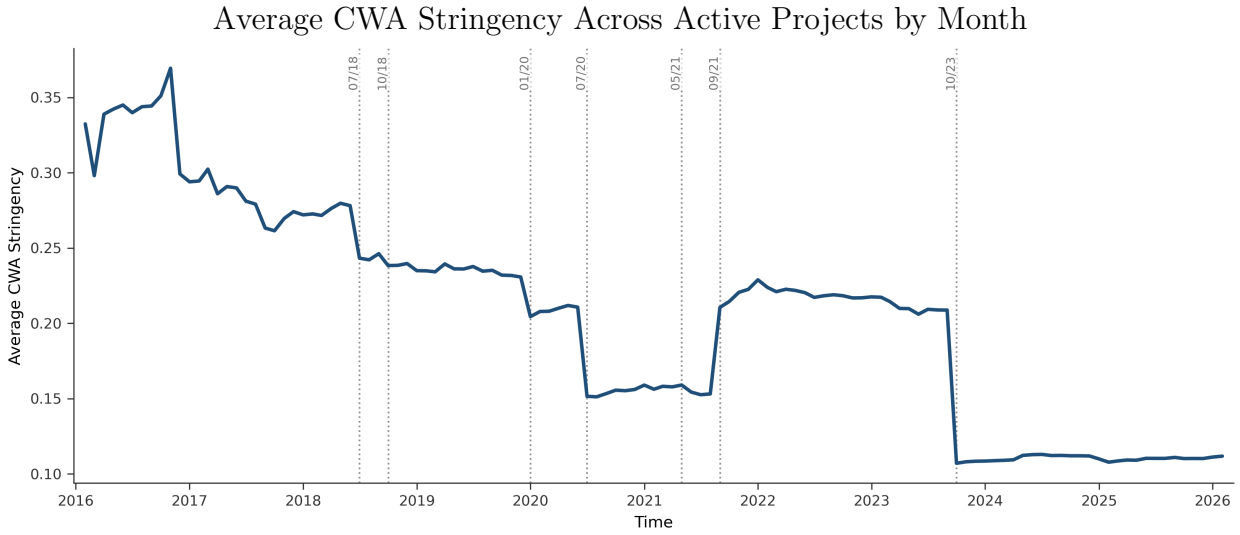
Notes: Each cell reports the number of projects that experienced at least one transition from the row regime to the column regime. Diagonal entries indicate projects with at least one within-regime one-step stay. The abbreviated regimes are the Navigable Waters Protection Rule (NWPR) and the Clean Water Rule (CWR).

2.5 Identification Strategy

We use variation in CWA jurisdiction across locations and over time to identify the effect of regulatory stringency on the duration and outcome of electric generation projects. Identification rests on three features of the setting. First, regime changes were determined by judicial decisions and administrative rulemaking processes largely independent of the energy industry, providing an arguably exogenous source of variation in regulatory exposure. Second, because regulations were applied unevenly across states—due to state-level litigation and injunctions—the same national-level rule change generated differential treatment across project locations, enabling cross-sectional comparisons. Third, individual projects often span multiple regulatory regimes, as shown in Table 2, allowing within-project variation to contribute to identification.

Two measurement challenges merit discussion. First, the DPH model must account for the fact that many projects begin and end under different CWA regimes, so the regime to which a project is "subject" is not a fixed characteristic but a time-varying exposure. Our likelihood function handles this directly: it marginalises over the possible assignment of regime-months to the project's phase path, subject to calendar constraints on the feasible regime sequence (Section 4).

Second, we estimate the direct effect of CWA stringency on time-to-build, holding project size fixed. If stricter regulation also induced projects to scale back their planned capacity—



and smaller projects take less time to permit—then the direct effect understates the total effect. This indirect channel is left for future work; our structural model abstracts from scale effects by assuming constant returns to scale.

3 A Structural Model of Permitting and Investment

This section develops a structural model of investment in which project sponsors must queue for regulatory approval before purchasing productive capital. The model links the time-to-build distribution—estimated empirically in Section 4—to aggregate investment through the user cost of capital.

The sector consists of *ex-ante* identical projects that produce a homogeneous commodity using capital as the only input. Before earning any revenue, each project's sponsor must obtain a permit. The permitting process lasts an uncertain number of periods, after which the project either exits successfully (receiving approval to build and operate) or exits unsuccessfully (being abandoned or denied). Duration and outcome are determined entirely by idiosyncratic project-level shocks, which arrive unexpectedly each period. Project sponsors have rational expectations over the distribution of these shocks but cannot predict any individual project's exact waiting time or outcome in advance.

Upon entering the queue, sponsors incur a sunk upfront cost U covering planning, design, and initial permitting submissions. They also pay a per-period carrying cost ω for each month they remain in the queue, and accumulate interest and dividend obligations on the debt and equity raised to finance these expenditures. Productive capital is purchased at price p^k only upon a successful exit. Project sponsors discount future payoffs at rate β and optimally balance the expected costs of waiting against the expected revenues from eventual production.

This framework shares features with Cacciatore and Fiori (2016) (regulatory entry costs as sunk costs), Kalouptsidi (2014) and Glancy et al. (2024) (option value of investment under uncertainty), and Bennon et al. (2024) (environmental review as a survival process). The model's results should be interpreted as percentage changes in investment for the portion of activity directly affected by a policy change. Translating these estimates to aggregate effects requires scaling by the share of total sector investment subject to the regulation; details on CBO's scaling procedures are available in related CBO documentation.[4]

3.1 Project Value and the Time-to-Build Distribution

This section formalises the project value function and derives how the time-to-build distribution $f(\tau)$ enters through each cost component. The key distribution to estimate is $f(\tau)$; the remaining structural parameters are calibrated from external data sources. Time is discrete and indexed by $t \geq 0$ and projects are indexed by i. Before a project enters the permitting queue, the project commits to the size of the capital investment, k_{it}. Also, upon entering the queue, the project incurs a one-time cost, U, for planning-related expenses, and they pay ω every period while they wait in the queue. Capital can be purchased at price p^k only after the project exits the queue successfully. Share ϵ of their obligations are financed by equity, and share $1 - \epsilon$ are financed with debt. While waiting in the queue, the project realizes three possible random outcomes at the start of each period. Projects may:

[4]Congressional Budget Office, Estimating the Effects of Permitting Requirements on Private Investment (February 2026), `www.cbo.gov/publication/62100`.

- **Exit Successfully**: If a project successfully completes the permitting and building process, it earns a net present value of future profits, V_{it}^{Π}, minus the cost of the productive capital investment, $p^k k_{it}$, minus all of their accumulated obligations, D_{it}.
- **Exit Unsuccessfully**: If a project is abandoned or forced to stop the permitting process, the project earns no profits, but is still compelled to pay the accumulated debt, $D_{it}(1-\epsilon)$. They are not compelled to pay back equity obligations, $D_{it}\epsilon$.
- **Stay in the Queue**: The present value of staying in the queue as of period t is V_{it}.

Writing this in terms of the model's variables,

$$\begin{aligned} \text{Exit Successfully:} \quad & V_{it}^{\Pi} - p^k k_{it} - D_{it} \\ \text{Exit Unsuccessfully:} \quad & -D_{it}(1-\epsilon) \\ \text{Stay in the Queue:} \quad & V_{i,t+1} \end{aligned}$$

Two random variables, θ_{it} and ψ_{it}, govern which of these outcomes happen for a project in the current period. They are idiosyncratic to the project and can take on values of 0 or 1.

- **Exit Shock** (θ_{it}): This random variable is 1 if the project exits the queue, and 0 otherwise.
- **Approve Shock** (ψ_{it}): This random variable is 1 if the project is approved, and 0 otherwise.

Combining the three possible states with the two shocks gives the total value of starting the project in time t:

$$V_{it} = E\Bigg[\underbrace{\theta_{it}\psi_{it}\big(V_{it}^{\Pi} - p^k k_{it} - D_{it}\big) - \theta_{it}(1-\psi_{it})D_{it}(1-\epsilon)}_{\text{Exit from Queue}} + \underbrace{(1-\theta_{it})\beta V_{i,t+1}}_{\text{Stay in Queue}}\Bigg] \tag{1}$$

Each project's owner has a constant discount rate of β that is between 0 and 1. Iterating Equation (1) forward and applying the expectations operator:

$$V_{it} = \underbrace{\sum_{\tau=0}^{\infty} \beta^{\tau} V_{i,t+\tau}^{\Pi}\, g(\tau)}_{\text{PV of producing}=V^{\text{Total Revenue}}} - \underbrace{\sum_{\tau=0}^{\infty} \beta^{\tau} k_{it} p^{k}\, g(\tau)}_{\text{PV of investment}=V^{\text{Capital Cost}}} - \underbrace{\sum_{\tau=0}^{\infty} \beta^{\tau} D_{i,t+\tau}\big[\epsilon\, g(\tau) + f(\tau)(1-\epsilon)\big]}_{\text{PV of Debt and Equity}=V^{\text{Upfront}}} \tag{2}$$

The function $f(\tau)$ describes the probability of any type of exit after exactly τ periods in the queue:

$$f(\tau) = E\left[\theta_{i,t+\tau} \prod_{s=0}^{\tau} (1 - \theta_{i,t+s})\right]$$

The function $g(\tau)$ represents the probability of only a successful exit after exactly τ periods in the queue, so $g(\tau) = E[\psi_{i,t+\tau}\, \theta_{i,t+\tau} \prod_{s=0}^{\tau}(1-\theta_{i,t+s})]$. The next section describes the parameterisation of $f(\tau)$ and $g(\tau)$, and the following sections discuss the three components on the right side of Equation (2)—total revenue, capital cost, and upfront costs—re-written here:

$$V_{it} = V_i^{\text{Total Revenue}} - V_i^{\text{Capital Cost}} - V_i^{\text{Upfront}}$$

3.2 Distribution of a Project's Time to Build

The function $f(\tau)$ represents the probability of any type of exit from the permitting queue after τ periods, and $g(\tau)$ is the probability of only a successful exit. To model the survival time of projects in the queue, this framework uses a generalized version of the Geometric distribution, called the Discrete Phase-type Distribution (DPH), where the time to absorption is a Markov chain sequence of stages.

More stages in the Markov chain of $f(\tau)$ allow the distribution to match more moments in the data—for example, with more than one stage, the model can match not just the mean, but both the mean and median of an observed distribution. In our empirical application, we have

observed projects phases, so we know when and for how long each project is in construction or permitting. Also, multiple stages are a stylized way of representing distinct permitting requirements, or distinct phases of the project lifecycle, such as planning, permitting, and construction. For the case when there is only one stage, the time-to-build distribution is a simple Geometric distribution, and the Markov chain reduces to a set of independent Bernoulli trials.

The distribution for $f(\tau)$ is defined below for $\tau \geq 0$. Letting p be the number of stages, $\mathbf{A}^{(l)}$ is a $p \times p$ sub-stochastic transition matrix whose entries depend on the regulatory regime l; it represents the probabilities of transitioning between one of p stages of the permitting process:

$$f(\tau) = \boldsymbol{\pi}' \mathbf{A}^{\tau} (\mathbf{e}^{G} + \mathbf{e}^{B}) \tag{3}$$

The vectors $\mathbf{e}^{G}$ and $\mathbf{e}^{B}$ are $p \times 1$ vectors that represent the probability of exiting successfully ($\mathbf{e}^{G}$) or unsuccessfully ($\mathbf{e}^{B}$) given that the project started that period in one of the p states. Finally, $\boldsymbol{a}_0$ is a $p \times 1$ vector that describes how projects are initially distributed across the p stages. The probability of a successful exit, $g(\tau)$, is the probability of exiting successfully, so $g(\tau) = \boldsymbol{\pi}' \mathbf{A}^{\tau} \mathbf{e}^{G}$.

In this model, the number of stages is set to two, construction and permitting ($p = 2$), so the matrices are:

$$\boldsymbol{\pi} = \begin{pmatrix} \pi_1 \\ 1 - \pi_1 \end{pmatrix}, \quad \mathbf{A}^{(l)} = \begin{pmatrix} \alpha_{11} & \alpha_{12} \\ \alpha_{21} & \alpha_{22} \end{pmatrix}$$

$$\mathbf{e}^{G} = \begin{pmatrix} (1 - \alpha_{11} - \alpha_{12})\, s_1 \\ (1 - \alpha_{22} - \alpha_{21})\, s_2 \end{pmatrix}, \quad \mathbf{e}^{B} = \begin{pmatrix} (1 - \alpha_{11} - \alpha_{12})(1 - s_1) \\ (1 - \alpha_{22} - \alpha_{21})(1 - s_2) \end{pmatrix}$$

The parameters have the following interpretations. The entry α_{ij} is the joint probability of surviving period t and transitioning to stage j, given that the project started in stage i under regime l. The parameter $s_i^{(l)}$ is the probability of a successful exit conditional on exiting from stage i, and a_{0i} is the initial probability of starting in stage i. The row deficit

$(1 - \sum_j \alpha_{ij}) = h_i^{(l)}$ is the per-period exit probability (the hazard) from stage i. The entries of $\mathbf{e}^G$ and $\mathbf{e}^B$ are then $e_i^{(l,G)} = h_i^{(l)} s_i^{(l)}$ and $e_i^{(l,B)} = h_i^{(l)}(1 - s_i^{(l)})$, consistent with the absorption vector $e_j^{(l,\omega)}$ defined in Section 4. The structural and empirical DPH distributions are therefore the same object, calibrated at regime l.

4 Estimating the Likelihood

We observe n infrastructure projects in the EIA860M dataset, indexed by i. Each project passes through p discrete phases — in our application permitting and construction, indexed $j \in \{1, \ldots, p\}$ — before exiting. The project's phase at period t is $x_t \in \{1, \ldots, p\}$, and crucially the full phase path $(x_1^i, x_2^i, \ldots, x_{d_i}^i)$ is observed in the data: we know whether each project was in permitting or construction in every month of our sample. The project's total observed duration is d_i, a discrete count of months.

The regulatory environment is characterised by m discrete treatment states $l \in \{1, \ldots, m\}$, each corresponding to a distinct CWA regulatory regime (NWPR, Rapanos, CWR, or Sackett) with an associated numerical stringency value $z^{(l)}$. The treatment state prevailing in calendar time at period t is l_t, but the project-level treatment exposure is *latent*: we know the calendar sequence of regimes but not exactly which regime governed each period of a given project's life.

The calendar history for each project is encoded as a sequence of blocks $(l_1^i, b_1^i), (l_2^i, b_2^i), \ldots$, where l_r^i is the prevailing CWA regime during the r-th block and b_r^i is its length in months, with $\sum_r b_r^i = d_i$. The first block (l_1, b_1) is particularly important: the project is locked in regime l_1 with certainty for all b_1 periods. When considering whether the project may have switched to a target regime $l_{\text{tgt}} \neq l_1$, the quantity $B_{\text{tgt}} = \sum_{r:\, l_r = l_{\text{tgt}}} b_r$ is the total calendar time available in that target regime over the project's life.

Each project has a vector of covariates $\mathbf{x}_i \in \mathbb{R}^q$, including characteristics such as generating capacity (MW), location, and technology type. The three core model components — all

depending on the current phase j and regime l — are the per-period exit hazard $h_j^{(l)}$, the routing probabilities $S_{jk}^{(l)}$ governing transitions between phases, and the success probability $s_j^{(l)}$ conditional on exit. The initial phase distribution $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_p)$, with $\pi_j = P(x_1 = j)$, gives the probability of starting in each phase.

Projects exit with type $\omega_i \in \{G, B\}$, where G (good) denotes a successful approved exit and B (bad) denotes an unsuccessful termination. The exit type is observed but the treatment state at exit is not. The switch time κ counts the number of additional l_1-steps taken after the first calendar block before the project switches permanently to the target regime; $\kappa = s_{\text{after}} \equiv d_i - b_1 - 1$ corresponds to the no-switch case.

The phase path is fully observed and the treatment path $\{l_t\}$ is marginalised under three structural constraints described in Section 4.2. The resulting marginalised likelihood for project i is a sum over the switch time κ, each term being the probability of the observed data under the treatment path that switches at step κ:

$$\mathcal{L}_i(\theta) \;=\; \left(\underbrace{\pi_{x_1} \prod_{t=1}^{b_1} A_{x_t,x_{t+1}}^{(l_1)}}_{\text{first block (certain)}} \right) \cdot \left(\sum_{\kappa=\kappa_{\min}}^{s_{\text{after}}} \underbrace{\prod_{t=b_1+1}^{b_1+\kappa} A_{x_t,x_{t+1}}^{(l_1)} \cdot \prod_{t=b_1+\kappa+1}^{d_i-1} A_{x_t,x_{t+1}}^{(l_{\text{tgt}})} \cdot e_{x_{d_i}}^{(l^*(\kappa),\,\omega_i)}}_{P(\text{data}|\text{switch at extra-step } \kappa)} \right), \tag{4}$$

where $l^*(\kappa) = l_1$ if $\kappa = s_{\text{after}}$ (no switch) and $l^*(\kappa) = l_{\text{tgt}}$ otherwise.

The κ-th term inside the sum is the probability of the observed data $(x_1, \ldots, x_{d_i}, \omega_i)$ under the specific treatment path that stays in l_1 for $b_1 + \kappa$ periods then switches to l_{tgt} for the remainder. Summing over κ averages over treatment paths consistent with the calendar constraints. Since the constraints rule out impossible switch times but do not assign differential probabilities to valid ones, the sum is an unweighted sum over valid paths.

The transition matrix $\mathbf{A}^{(l)}$ has a natural factorisation that separates the decision to exit from the decision of which stage to move to next. Specifically,

$$\mathbf{A}^{(l)} = \operatorname{diag}(\mathbf{1} - \mathbf{h}^{(l)})\, \mathbf{S}^{(l)}, \tag{5}$$

where $\mathbf{h}^{(l)}$ is a vector of per-period exit hazards and $\mathbf{S}^{(l)}$ is the routing matrix. The row sums satisfy $\sum_k A_{jk}^{(l)} = 1 - h_j^{(l)} < 1$; the row deficit is the probability of exiting the queue in that period. The routing matrix $\mathbf{S}^{(l)}$ is fully populated: all off-diagonal entries $S_{jk}^{(l)} > 0$ for $k \neq j$, permitting backward transitions from construction back to permitting as well as forward ones.

The absorption vector $e_j^{(l,\omega)}$ collects the probability of exiting from phase j under regime l with outcome ω:

$$e_j^{(l,\omega)} = h_j^{(l)} \times \begin{cases} s_j^{(l)} & \omega = G, \\ 1 - s_j^{(l)} & \omega = B. \end{cases} \tag{6}$$

This expression shows that exit events have two sequential components: first the project exits (probability $h_j^{(l)}$), then the outcome is determined by the success probability $s_j^{(l)}$.

4.1 Parameter Function Forms

Each of the three DPH model components—the per-period exit hazard $h_j^{(l)}$, the routing probabilities $S_{jk}^{(l)}$, and the success probability $s_j^{(l)}$—is parameterized in two steps. First, a baseline value is estimated as a function of project covariates $\mathbf{x}_i$. Second, the regulatory regime l shifts the baseline through a power transformation, so that higher or lower CWA stringency scales the probability multiplicatively. This two-step structure ensures that all probabilities remain in $(0, 1)$ for any parameter value, avoiding constrained optimization.

Baseline hazard. The per-period probability that a project exits the queue from stage j, absent any treatment effect, is specified using the complementary log-log link. This link arises naturally from the grouped-data proportional hazards model of Prentice and Gloeckler (1978) and preserves a proportional hazards interpretation for continuous-time processes observed at discrete intervals (Jenkins, 1995):

$$h_j^{\text{base}}(\mathbf{x}_i) = 1 - \exp\big(-\exp(\gamma_j + \mathbf{c}_j^\top \mathbf{x}_i)\big). \tag{7}$$

The parameter γ_j is a phase-specific intercept and $\mathbf{c}_j$ is a vector of covariate coefficients. Larger values of $\gamma_j + \mathbf{c}_j^\top \mathbf{x}_i$ produce higher exit probabilities; when the linear index is very negative the project is nearly certain to remain in the queue each period.

Baseline routing matrix. The routing probabilities $S_{jk}^{\text{base}}(\mathbf{x}_i)$ govern which stage k a project moves to, conditional on not exiting in the current period. The routing matrix is fully populated—all off-diagonal entries are positive, allowing transitions in either direction between stages, including backward transitions from construction back to permitting. The self-loop probability ($k = j$, staying in the same stage) serves as the reference category in a softmax parameterisation:

$$S_{jk}^{\text{base}}(\mathbf{x}_i) = \frac{\exp(\mu_{jk} + \boldsymbol{\nu}_{jk}^\top \mathbf{x}_i)}{1 + \sum_{r \neq j} \exp(\mu_{jr} + \boldsymbol{\nu}_{jr}^\top \mathbf{x}_i)}, \quad k \neq j, \tag{8}$$

$$S_{jj}^{\text{base}}(\mathbf{x}_i) = \frac{1}{1 + \sum_{r \neq j} \exp(\mu_{jr} + \boldsymbol{\nu}_{jr}^\top \mathbf{x}_i)}. \tag{9}$$

Here μ_{jk} is an intercept and $\boldsymbol{\nu}_{jk}$ is a covariate coefficient vector for the $j \to k$ transition. A large positive μ_{12} means projects in the permitting stage advance to construction with high probability relative to remaining in permitting.

Baseline success probability. The probability of a good outcome conditional on exiting from stage j is modelled with a logistic link:

$$s_j^{\text{base}}(\mathbf{x}_i) = \sigma\big(\alpha_j + \mathbf{b}_j^\top \mathbf{x}_i\big), \quad \sigma(u) = \tfrac{1}{1+e^{-u}}. \tag{10}$$

Treatment scaling. The regulatory regime l enters by raising the complement of each baseline probability to a power $\exp(\zeta \cdot z^{(l)})$, where $z^{(l)}$ is the numerical CWA stringency value for regime l and ζ is an estimated treatment effect parameter. This power-transformation form keeps all probabilities in $(0, 1)$ for any real ζ and is analogous to proportional hazards

scaling (Lancaster, 1990). The three components scale as follows:

$$1 - h_j^{(l)} = \left(1 - h_j^{\text{base}}\right)^{\exp(\zeta_j^h z^{(l)})}, \tag{11}$$

$$S_{jj}^{(l)} = \left(S_{jj}^{\text{base}}\right)^{\exp(\zeta_j^S z^{(l)})}, \tag{12}$$

$$1 - s_j^{(l)} = \left(1 - s_j^{\text{base}}\right)^{\exp(\zeta_j^s z^{(l)})}. \tag{13}$$

Off-diagonal routing entries rescale proportionally to preserve row sums: $S_{jk}^{(l)} = S_{jk}^{\text{base}} \cdot (1 - S_{jj}^{(l)})/(1 - S_{jj}^{\text{base}})$ for $k \neq j$. A positive ζ_j^h means higher stringency (larger $z^{(l)}$) raises the exit hazard; a negative ζ_j^h means higher stringency slows exit. The full transition matrix entries combine the hazard and routing scalings:

$$\alpha_{jj}^{(l)} = \left(1 - h_j^{\text{base}}\right)^{\exp(\zeta_j^h z^{(l)})} \cdot \left(S_{jj}^{\text{base}}\right)^{\exp(\zeta_j^S z^{(l)})}, \tag{14}$$

$$\alpha_{jk}^{(l)} = \left(1 - h_j^{\text{base}}\right)^{\exp(\zeta_j^h z^{(l)})} \cdot S_{jk}^{\text{base}} \cdot \frac{1 - \left(S_{jj}^{\text{base}}\right)^{\exp(\zeta_j^S z^{(l)})}}{1 - S_{jj}^{\text{base}}}, \quad k \neq j. \tag{15}$$

The treatment effect parameters ζ_j^h, ζ_j^S, and ζ_j^s are the objects of primary interest: they measure the CWA's impact on exit speed, stage-transition routing, and approval probability respectively. Their identification comes from cross-sectional and over-time variation in $z^{(l)}$ across project locations and the sequence of CWA regime changes.

4.2 Effect of Observed Phases and Treatment Constraints

Detailed timelines on project phases greatly reduces the computational burden of calculating the data likelihood. Because the phase path $(x_1, \ldots, x_{d_i})$ is observed, each project's phase transition $x_t \to x_{t+1}$ contributes a simple product of scalars to the likelihood: the product of phase transitions $A_{x_t, x_{t+1}}^{(l_t)}$ and the probability of terminal period result $e_{x_{d_i}}^{(l_{d_i}, \omega_i)}$. The remaining burden is to integrate over the *treatment* path $\{l_t\}$, which is restricted by three structural constraints that we impose to align the framework the real-world observations and restrictions:

1. **First-block lock-in.** The project is in regime l_1 with certainty for all periods $t =$

$1, \ldots, b_1$ of the first calendar block. For example, if a project starts under the Clean Water Rule, they are surely subject to that rule for as long as the rule is in place.

2. **At most one transition.** After the first block, the treatment path takes the form $l_1^{b_1+\kappa}, \ldots, l_{\text{tgt}}^{d_i-b_1-\kappa}$ for some $\kappa \geq 0$. This means that a project can be given at most one AJD, and it may not necessarily be granted under the initial CWA regime. Setting $\kappa = s_{\text{after}} \equiv d_i - b_1 - 1$ corresponds to no switch.

3. **Unobserved terminal treatment.** The exit type ω_i is observed but the treatment state at exit is not:

   - $\kappa = s_{\text{after}}$: absorbed in l_1, using $e_{x_{d_i}}^{(l_1,\omega_i)}$.
   - $\kappa < s_{\text{after}}$: absorbed in l_{tgt}, using $e_{x_{d_i}}^{(l_{\text{tgt}},\omega_i)}$.

A calendar capacity constraint further bounds κ from below: the project cannot spend more periods in l_{tgt} than the calendar offered. With $B_{\text{tgt}} = \sum_{r:\, l_r = l_{\text{tgt}}} b_r$ the total calendar time in that regime,

$$\kappa \;\geq\; \kappa_{\min} = \max(0,\; s_{\text{after}} + 1 - B_{\text{tgt}}). \tag{16}$$

4.3 Estimation Results

Table 3 reports likelihood ratio tests for four restricted models, each of which sets one vector of treatment effect parameters to zero while retaining the others. The three treatment effect vectors are ζ^h (CWA's effect on the per-period exit hazard), ζ^S (CWA's effect on the routing matrix governing transitions between permitting and construction), and ζ^s (CWA's effect on the probability of a successful exit conditional on exiting).

The LRT results indicate that CWA stringency has statistically significant effects through the hazard and routing channels. Restricting $\zeta^h = 0$ yields an LR statistic of 8.12 with 2 degrees of freedom ($p = 0.017$), and restricting $\zeta^S = 0$ yields a statistic of 40.61 ($p < 0.001$), suggesting that the routing channel—how stringency affects the transition from permitting

to construction—is the most precisely identified effect. The success-probability channel (ζ^s) is not individually significant ($p = 0.49$), which is unsurprising: approximately one-third of the projects in our sample are right-censored, meaning we do not observe their exit outcome, which limits identification of this parameter. The joint null that all $\zeta = 0$ is rejected with an LR statistic of 51.07 ($p < 0.001$).

The estimated DPH matrices under Rapanos and Sackett, displayed in Section 6 on calibration, are qualitatively consistent with the expected effects of a more permissive regulatory regime. Under Sackett, the self-loop probability in the permitting stage (α_{11}) falls from 0.941 to 0.938, meaning projects exit the permitting phase faster. The transition probability to construction (α_{12}) rises from 0.049 to 0.051. The construction self-loop (α_{22}) also falls modestly, and the successful-exit probability rises. These changes are consistent with a less burdensome permitting environment: projects move through both stages more quickly and are more likely to complete successfully. The simulated distribution moments in Table 3 show that expected project duration rises from 21.96 months under the baseline ($\zeta = 0$) to 23.71 months under the estimated Rapanos treatment effects, a difference of 1.75 months attributable to CWA regulatory exposure.

5 A Structural Model of Investment Subject to the Clean Water Act

We now detail the structural components of the model and connect them to the estimated DPH distributions. The simulation focuses on solar electric generation projects, which account for 5,202 of the 8,512 projects in our sample (Table 1) and for which our estimates are most precisely identified. We use the DPH distributions estimated under Rapanos and Sackett (see Section 6), evaluated at the average megawatt capacity for solar projects, to simulate the investment response to the Sackett decision. We treat the Rapanos-to-Sackett transition as a one-time, permanent discrete shock to the time-to-build distribution affecting

Table 3: Key Model Results

Likelihood Ratio Tests				
Restriction	LR stat	DF	p-value	
$\zeta_h = 0$	8.12	2	0.0173	*
$\zeta_S = 0$	40.61	2	0.0000	***
$\zeta_s = 0$	1.42	2	0.4926	
All $\zeta = 0$	51.07	6	0.0000	***
Simulated Distribution Moments				
Moment	Baseline: $\zeta = 0$	Treated: ζ unrestricted per Greenhill et al. (2026)		
E[Duration]	21.960	23.707		
SD[Duration]	24.519	26.315		
Skewness	2.553	2.415		
Estimated Treatment Effects				
Parameter	Estimate	Reg. SE	p-value	
$\zeta_{h,1}$	-0.1442	0.0512	0.0049	**
$\zeta_{h,2}$	-0.2014	0.0547	0.0002	***
$\zeta_{S,1}$	-0.4763	0.0493	0.0000	***
$\zeta_{S,2}$	1.3193	0.0534	0.0000	***
$\zeta_{s,1}$	-0.2661	0.0766	0.0005	***
$\zeta_{s,2}$	-0.1593	0.0468	0.0007	***
Log-Likelihood	-33168.11			
Sample size	8512			

all projects simultaneously—both those already in the queue and those considering entry. This thought experiment isolates the investment-channel effect of the policy change and provides an upper-bound estimate of the speed of adjustment.

5.1 Present Value of Upfront Costs

Obligations accumulate while the project waits in the queue. In the initial period of entry, the project purchases a sunk initial investment, U. In later periods, the project incurs additional costs, ω, and rolls over obligations from previous periods. Each period the project sponsors pay interest on debt and issue dividends on equity shares at rate r_d and r_e, respectively. This paper defines $R \equiv 1 + r_e\epsilon + r_d(1-\epsilon)$ and r_e and r_d are set larger than $1/\beta - 1$. Again, the parameter ϵ represents the share of financing that comes from equity.

Obligations accumulate as follows:

$$\begin{aligned} B_{i,0} &= U \\ B_{i,1} &= B_{i,0}R + \omega \\ &\vdots \\ B_{i\tau} &= B_{i,\tau-1}R + \omega \end{aligned}$$

Plugging these into each other and iterating,

$$B_{i\tau} = R^\tau U + \omega \frac{R^\tau - 1}{R - 1}$$

The project exits after τ periods and starts making regular payments to lenders and shareholders. The project regularly pays off share λ of their current period's obligation; however, while they do so, they incur additional interest and dividend costs on the remaining obligations. Therefore,

$$B_{i,\tau+1} = B_{i\tau}R - \lambda B_{i\tau}R$$

The parameter λ is defined in a convenient way so that projects always have the same debt-to-capital ratio,

$$\frac{B_{i,\tau+1}}{B_{i\tau}} = R(1-\lambda) = 1-\delta$$

So λ is,

$$\lambda = \frac{R+\delta-1}{R}$$

After n periods of producing, the project's remaining obligation in period n is,

$$B_{i,\tau+n} = B_{i,\tau}(1-\delta)^n$$

The present value of all future payments is the discounted sum of all future payments:

$$D_{i\tau} = \sum_{n=0}^{\infty} \beta^n \lambda B_{i,\tau+n} R = B_{i\tau}\, p_d$$

where p_d is the added cost of carrying obligations while producing,

$$p_d = \frac{\beta(1-\delta)(R-1+\delta)}{1-\beta(1-\delta)} \tag{17}$$

These payments start when the project exits the queue. The project is only obligated to pay equity investors if the project exits successfully, and the project is obligated to pay the debt-financed portion for any type of exit. Therefore, the present value of these payments is

$$V_i^{\text{Upfront}} = \sum_{\tau=0}^{\infty} \beta^\tau D_{i,t+\tau} \Big[\underbrace{\epsilon\, g(\tau)}_{\text{Equity payout}} + \underbrace{(1-\epsilon) f(\tau)}_{\text{Debt payout}} \Big]$$

Simplifying, this is $V_i^{\text{Upfront}} = U p_u + \omega p_\omega$, where coefficients p_u and p_ω are equal to,

$$p_u = p_d\, \boldsymbol{a}_0' (\mathbf{I} - \beta R \mathbf{A}^{(l)})^{-1} \big(\mathbf{e}^G + (1-\epsilon)\mathbf{e}^B \big) \tag{18}$$

$$p_\omega = p_d\, \boldsymbol{a}_0' \Big[(\mathbf{I} - \beta R \mathbf{A}^{(l)})^{-1} - (\mathbf{I} - \beta \mathbf{A}^{(l)})^{-1} \Big] (R^{-1} - 1) \big(\mathbf{e}^G + (1-\epsilon)\mathbf{e}^B \big) \tag{19}$$

These coefficients represent the added cost of financing the sunk (U) and per-period costs (ω).

5.2 Present Value of Capital Cost

Successful projects will buy capital at the aggregate-good price that is normalized to 1. Tax and depreciation deductions are applied to the capital purchase. τ is the marginal tax rate on income, Z represents the value of depreciation deductions, and ξ is used in calibration and represents other deductions or taxes that may apply. The price of capital after these deductions is $p^k = (1 - \tau_{\text{tax}} Z)\xi$.

Similar to the upfront costs, the capital is financed by debt and equity. Using the definition of p_d given by (17), the future value of payments on productive capital k_{it} is given by $p_d p^k k_{it}$. The term $V_i^{\text{Capital Cost}}$ in Equation (2) is

$$V_i^{\text{Capital Cost}} = \sum_{\tau=0}^{\infty} \beta^\tau p_d p^k k_{it}\, g(\tau) = k_{it}\, \tilde{p}^k \tag{20}$$

where $\tilde{p}^k = p_d p^k \boldsymbol{a}_0' (\mathbf{I} - \beta \mathbf{A}^{(l)})^{-1} \mathbf{e}^G$.

5.3 Present Value of Total Revenues

Successful projects produce with constant returns to scale. They buy k_{it} upon a successful exit and capital starts depreciating at rate δ. Then $j \geq 1$ periods after exit, the project produces y_{ijt} using the $(1-\delta)^j k_{it}$ remaining capital:

$$y_{ijt} = (1 - \delta)^j k_{it}$$

The project sells output at price P_t in time t and there are no marginal costs, $\Pi_t(y_{ijt}) = y_{ijt}P_t$. Given a successful exit from the queue at time t, the present value of future profits is

$$V_{it}^{\Pi} = \beta\Pi_{t+1}(y_{ijt}) + V_{i,t+1}^{\Pi} = k_{it}\sum_{z=1}^{\infty} P_{t+z}\beta^z(1-\delta)^{z-1}$$

Written recursively, the term x_t is defined to represent the present value of capital after a successful exit,

$$x_t = \beta P_{t+1} + \beta x_{t+1}(1-\delta) \tag{21}$$

so that $V_{it}^{\Pi} = k_{it}x_t$. Therefore $V_i^{\text{Total Revenue}}$ in Equation (2) can be written

$$V_i^{\text{Total Revenue}} = \sum_{\tau=0}^{\infty} \beta^{\tau} V_{i,t+\tau}^{\Pi}\, g(\tau) = k_{it}\sum_{\tau=0}^{\infty} \beta^{\tau} x_{t+\tau}\, g(\tau) = k_{it}\, \boldsymbol{a}_0'\boldsymbol{\chi}_t$$

where the term $\boldsymbol{\chi}_t$ is,

$$\boldsymbol{\chi}_t = \mathbf{e}^G x_t + \beta\mathbf{A}^{(l)}\boldsymbol{\chi}_{t+1} \tag{22}$$

5.4 Project Entry and Investment

Projects enter until the value of entering is zero,

$$V_{it} \geq 0 \tag{23}$$

and using the results from the sections above,

$$k_{it}\, \boldsymbol{a}_0'\boldsymbol{\chi}_t - k_{it}\, \tilde{p}^k - Up_u - \omega p_\omega \geq 0 \tag{24}$$

In equilibrium, projects choose capital so that (24) holds with equality. After rearranging,

$$k_{it} = \frac{Up_u + \omega p_\omega}{\boldsymbol{a}_0'\boldsymbol{\chi}_t - \tilde{p}^k} \tag{25}$$

This is the level of investment in each period of the model that ensures that the free entry condition in Equation (23) holds.

5.5 Aggregation

Project sponsors are price takers upon exiting the queue, there are sunk entry costs, and returns to scale in production are constant—this set of features is not consistent with an endogenously determined number of entrants and a finite, non-degenerate distribution of project sizes.[5] Therefore the mass of entrants each period is set at a constant number, $\mathcal{N}$:

$$k_t = \mathcal{N} k_{it} = \int_{i \in \Omega_t} k_{it}\, di$$

where Ω_t is the set of i-entrants that enter in time t. The number of projects in stage v of m possible stages is

$$N_t^v = \sum_{q=1}^{m} N_{t-1}^q \alpha_{qv} + a_{0v}\mathcal{N}$$

with a_{0v} being the v-th row of $\boldsymbol{a}_0$, and α_{qv} the entry in row q and column v of $\mathbf{A}^{(l)}$. The average amount of investment for each project in stage v, k_t^v, is the weighted average of investment for all entering projects:

$$k_t^v = \sum_{q=1}^{m} \phi_{qv}^t k_{t-1}^q + a_{0v} k_{it}$$

where each weight is $\phi_{qv}^t = N_{t-1}^q \alpha_{qv} / N_t^v$. Re-writing this and stacking across m states,

$$\boldsymbol{\Delta}_t' = \boldsymbol{\Delta}_{t-1}' \mathbf{A}^{(l)} + \boldsymbol{a}_0' k_{it} \tag{26}$$

[5] See Dixit and Stiglitz (1977) for the relevant discussion of firm entry, sunk costs, and price taking.

and $\mathbf{\Delta}_t$ is an $m \times 1$ vector with values in each row v equal to $\Delta_t^v = N_t^v k_t^v$. Successfully exiting projects purchase capital and invest, so total investment each period is

$$I_t = \mathbf{\Delta}_t' \times \mathbf{e}^G \tag{27}$$

Depreciation starts once the project starts producing, so total output each period is

$$Q_t = Q_{t-1}(1-\delta) + I_t \tag{28}$$

The demand curve for Q_t is represented by a CES function that satisfies goods market clearing,

$$P_t = \gamma(Q_{t-1})^{-1/\sigma} \tag{29}$$

and $\sigma > 0$, $\gamma > 0$ are parameters.

5.6 The User Cost of Capital

This section derives several expressions that are related to the calibration of the investment elasticity with respect to the User Cost of Capital (UCC). Since this section only derives conditions for the steady state of the model, t-subscripts are dropped for convenience. The free entry condition in Equation (23) implies that, for a project entering in any period, the present value of the returns must equal the present value of the capital and upfront costs:

$$V_i^{\text{Total Revenue}} = V_i^{\text{Capital Cost}} + V_i^{\text{Upfront}} = V_i^{\text{Total Cost}} \tag{30}$$

Scaling both sides by k_i, the rate of return at which the marginal project is indifferent to entering is defined as V_i^{Enter}. Also, for convenience, V_i^{Succeed} is defined as the return on a

project if it is successful:

$$\frac{V_i^{\text{Total Revenue}}}{k_i} = \frac{V_i^{\text{Total Cost}}}{k_i} \equiv V_i^{\text{Enter}} = \sum_{\tau=0}^{\infty} \beta^{\tau} g(\tau)\, V_i^{\text{Succeed}} \tag{31}$$

Because some projects fail and all revenues are delayed, the value of a project before entry is lower than its value after a successful exit.

Suppose the project successfully exits and starts earning a consistent return of $\rho + \delta$ each period. These returns are enough to compensate the sponsor for capital depreciation, δ, while providing a market rate of return net of depreciation, ρ. The UCC is defined as $\rho + \delta$. The present value of those returns after adjusting for taxes (τ_{tax}) is

$$\begin{aligned} V_i^{\text{Succeed}} &= \beta(1-\delta)(\rho+\delta)(1-\tau) + \beta^2(1-\delta)^2(\rho+\delta)(1-\tau) + \cdots \\ &= \frac{(\rho+\delta)\,\beta(1-\delta)(1-\tau_{\text{tax}})}{1-\beta(1-\delta)} \end{aligned} \tag{32}$$

Using the definition for $g(\tau)$, V_i^{Enter} is,

$$\begin{aligned} V_i^{\text{Enter}} &= \sum_{\tau=0}^{\infty} \beta^{\tau} g(\tau)\, V_i^{\text{Succeed}} \\ &= \frac{(\rho+\delta)\,\beta(1-\delta)(1-\tau_{\text{tax}})}{1-\beta(1-\delta)}\, \boldsymbol{a}_0'(\mathbf{I} - \beta\mathbf{A}^{(l)})^{-1}\mathbf{e}^G \end{aligned} \tag{33}$$

where the second line follows because $\sum_{\tau=0}^{\infty} \beta^{\tau}\mathbf{A}^{(l)\,\tau} = (\mathbf{I} - \beta\mathbf{A}^{(l)})^{-1}$. Finally, using the result that $V_i^{\text{Total Cost}} = k_i\tilde{p}^k + Up_u + \omega p_\omega$ and rearranging,

$$\frac{k_i\tilde{p}^k + Up_u + \omega p_\omega}{k_i} = \frac{(\rho+\delta)\,\beta(1-\delta)(1-\tau_{\text{tax}})}{1-\beta(1-\delta)}\, \boldsymbol{a}_0'(\mathbf{I} - \beta\mathbf{A}^{(l)})^{-1}\mathbf{e}^G$$

Then using Equation (25), the UCC is

$$\text{UCC} \equiv \rho + \delta = \frac{\boldsymbol{a}_0'\boldsymbol{\chi}\big[1-\beta(1-\delta)\big]}{\beta(1-\delta)(1-\tau)\,\boldsymbol{a}_0'(\mathbf{I} - \beta\mathbf{A}^{(l)})^{-1}\mathbf{e}^G} \tag{34}$$

In steady state, it can be shown that the elasticity of investment with respect to the UCC is equal to the price elasticity of demand for the good, which is the parameter $-\sigma$ used in the demand curve given in Equation (29):

$$\frac{dI}{I} \cdot \frac{\text{UCC}}{d\,\text{UCC}} = -\sigma$$

Referring to the literature that explores this UCC elasticity, the appropriate calibration for σ depends on the model setting. In this model, Q_t should be interpreted as the capital good in a partial equilibrium setting, where labor, wages, interest rates, and an aggregate good's quantity and price are decided outside of the model. Q_t is installed capital bought by firms with an elasticity of substitution between capital and labor (or some other input) equal to $-\sigma$. This interpretation of σ aligns with what Chodorow-Reich (2025) calls "the generalized Hall and Jorgenson (1967) user cost elasticity", or the elasticity of substitution between capital and labor in the production function where the capital is used. Therefore $\sigma = 0.7$, a value that is consistent with this interpretation and with CBO's modeling of this elasticity elsewhere.[6]

6 Calibration

The calibrated DPH distributions under Rapanos and Sackett are displayed below. Table 4 summarises the remaining structural parameters. Under the Rapanos regime, the estimated DPH distribution for solar projects is:

$$f_{\text{Rapanos}}(\tau) = \underbrace{\begin{pmatrix} \mathbf{0.590} \\ (0.578,\ 0.602) \\ \mathbf{0.410} \\ (0.398,\ 0.422) \end{pmatrix}^{\mathsf{T}}}_{\text{Initial } \boldsymbol{\pi}} \times \underbrace{\begin{bmatrix} \mathbf{0.941} & \mathbf{0.049} \\ (0.938,\ 0.943) & (0.047,\ 0.050) \\ \mathbf{0.002} & \mathbf{0.905} \\ (0.002,\ 0.003) & (0.902,\ 0.908) \end{bmatrix}}_{\text{Transition } \mathbf{A}^{(l)}} \times \left[\underbrace{\begin{pmatrix} \mathbf{0.002} \\ (0.002,\ 0.002) \\ \mathbf{0.092} \\ (0.089,\ 0.095) \end{pmatrix}}_{\text{Successful } \mathbf{e}^G} + \underbrace{\begin{pmatrix} \mathbf{0.009} \\ (0.008,\ 0.009) \\ \mathbf{0.001} \\ (0.001,\ 0.001) \end{pmatrix}}_{\text{Unsuccessful } \mathbf{e}^B} \right]$$

[6] See Congressional Budget Office (2018). Chodorow-Reich (2025) also provides a detailed discussion of CBO's use of this elasticity.

Under the Sackett supreme court decision, the estimated DPH distribution is the following. In the structural simulation, we treat the Rapanos distribution as the baseline and compute

$$f_{\text{Sackett}}(\tau) = \underbrace{\begin{pmatrix} \mathbf{0.590} \\ (0.578,\ 0.602) \\ \mathbf{0.410} \\ (0.398,\ 0.422) \end{pmatrix}^{\mathsf{T}}}_{\text{Initial } \boldsymbol{\pi}} \times \underbrace{\begin{bmatrix} \mathbf{0.938} & \mathbf{0.051} \\ (0.936,\ 0.940) & (0.049,\ 0.053) \\ \mathbf{0.002} & \mathbf{0.903} \\ (0.002,\ 0.002) & (0.900,\ 0.906) \end{bmatrix}}_{\text{Transition } \mathbf{A}^{(l)}} \times \left[\underbrace{\begin{pmatrix} \mathbf{0.002} \\ (0.002,\ 0.003) \\ \mathbf{0.094} \\ (0.091,\ 0.097) \end{pmatrix}}_{\text{Successful } \mathbf{e}^G} + \underbrace{\begin{pmatrix} \mathbf{0.009} \\ (0.008,\ 0.010) \\ \mathbf{0.001} \\ (0.001,\ 0.001) \end{pmatrix}}_{\text{Unsuccessful } \mathbf{e}^B} \right]$$

the steady-state equilibrium under this parameterization. We then implement the Sackett decision as an unanticipated permanent shock that replaces the Rapanos DPH with the Sackett DPH, and solve for the new equilibrium path using the perfect-foresight solver described in Section 6.1.

Table 4 lays out the key parameters of the model with descriptions and some context about how CBO arrived at the parameters' values.

6.1 Structural Model Solution

To summarize the model, there are seven variables x_t, $\boldsymbol{\chi}_t$, P_t, Q_t, I_t, k_t, $\boldsymbol{\Delta}_t$ and seven equations (21), (22), (25), (26), (27), (28), (29). The model is solved in Dynare 6.2 using the perfect foresight solver with default settings.[8]

When the dimension of $\mathbf{A}^{(l)}$ subject to an unexpected policy change is greater than 1 (i.e. $p > 1$), the model's dynamics depend on the initial conditions and there is a random walk component. Since the benchmark has $p = 2$, this paper uses an approach like those discussed in Schmitt-Grohé and Uribe (2003) to make an *ad hoc* adjustment in the model. The adjustment has a negligible effect on the solution but grants long-run stability by eliminating the model's random walk. The adjustment is to define an autoregressive process in the sunk cost U, so that sunk costs rise as expected capital demand increases relative to steady-state

[8]Documentation for Dynare is available at `https://www.dynare.org/`.

Table 4: Model Parameters

Parameter	Description	Source
δ	Depreciation rate, monthly	CBO's CapTax model: Burnham and Carloni (2022)
β	Long-run rate of return on risk-free debt, monthly	
τ	Marginal tax rate on earnings	
ρ	Return on investment net of depreciation, monthly	
R	Weighted return on debt and equity for each BEA sector, monthly	
Z	Share of depreciation deductions for each BEA sector	
ξ	Other taxes and deductions that may apply for each BEA sector	Chosen so ρ matches CapTax's return on investment after depreciation
σ	Price elasticity of demand for the project; also the long-run elasticity of investment with respect to the user cost	Congressional Budget Office (2018)
U	Sunk cost of entering the time-to-build queue	Calibrated to match the ratio of fixed assets to total assets for each BEA sector[7]
ω	Per-period constant cost of waiting in the time-to-build queue	Calibrated to results from Trebbi and Zhang (2022a)

capital. Specifically, U is allowed to vary over time according to the process

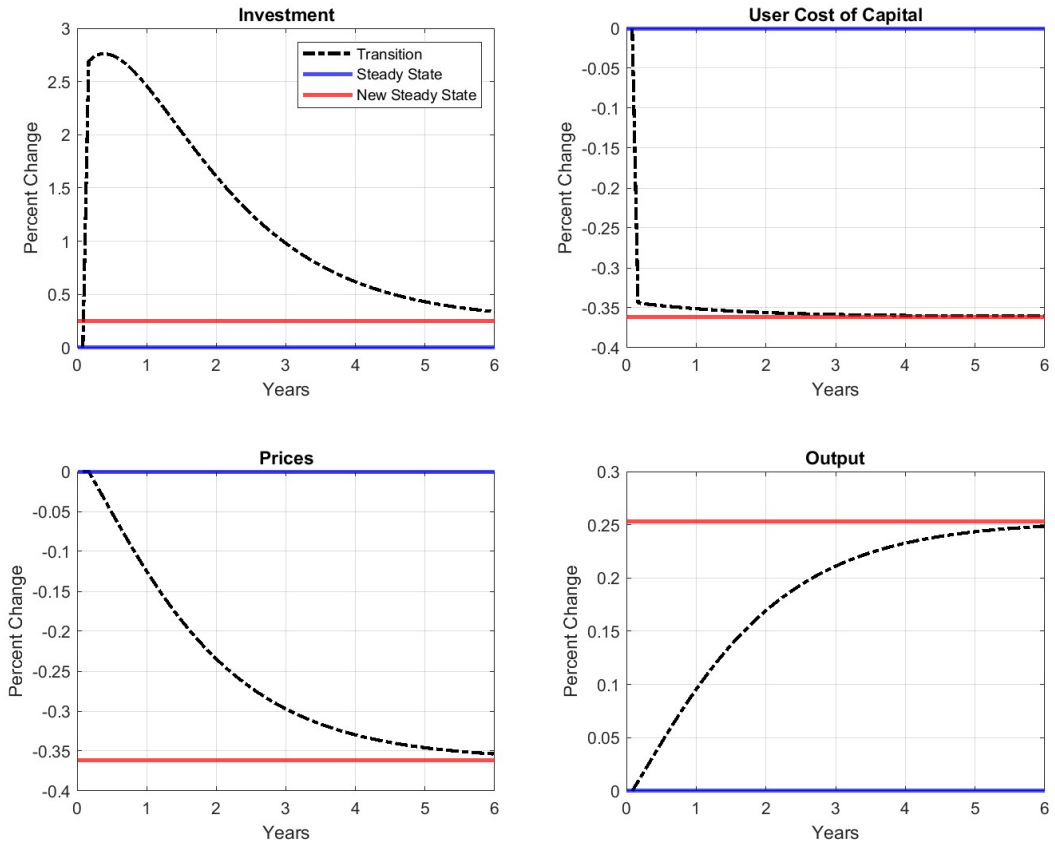
$$U_{t+1} = \bar{U}\,(1-\theta) + U_t\,\theta\,(Q_t/\bar{Q})$$

where $\bar{U}$ is the calibrated sunk cost parameter, $\bar{Q}$ is the steady-state value of capital output, and θ controls the deviation from $\bar{U}$. To minimize the effect of this *ad hoc* correction on the results, θ is set to 10^{-10}.

7 Structural Model Results

Figure 2 displays the impulse response of investment, output, and the user cost of capital to the Sackett shock for solar generation projects. Following the regime change, the user cost of capital falls as the expected time in the permitting queue shortens and the probability of successful exit rises. This lowers the threshold for profitable entry, attracting additional projects into the queue. Investment rises as projects that entered under the new regime begin exiting successfully. The long-run elasticity of investment with respect to the reduction in expected permitting duration is approximately -0.36, somewhat below the housing-supply elasticity estimated by Gabriel and Kung (2024) but consistent with the generally lower short-run demand elasticity for electricity infrastructure.

Figure 2: Dynamic Response to the Sackett Decision



8 Conclusion

This paper develops a framework for measuring how changes in federal environmental permitting requirements affect investment. The principal mechanism is the time-to-build channel: permitting delays raise the effective cost of capital, reduce the expected return to entry, and discourage new investment. We estimate this channel using variation in Clean Water Act jurisdictional stringency generated by a sequence of Supreme Court decisions and administrative rules, matched to monthly data on the development status of planned electric generators.

Our empirical findings indicate that CWA stringency affects project timelines through two

channels: the per-period exit hazard from the permitting and construction phases, and the probability of transitioning from permitting to construction. The routing channel is the more precisely identified of the two, while the effect on the probability of successful completion is less precisely estimated, in part because approximately one-third of the projects in our sample are right-censored. Evaluated at the transition from Rapanos to Sackett, the model implies that the Sackett decision reduced expected project duration by roughly 1.75 months and increased expected approval rates.

Feeding these estimates into a calibrated structural model of capital investment, we simulate the investment effects of the Sackett decision for solar energy projects. We find a long-run elasticity of investment with respect to a reduction in permitting duration of approximately -0.36. This estimate is smaller in magnitude than comparable estimates for housing (Gabriel and Kung, 2024), consistent with the lower short-run demand elasticity for electricity infrastructure.

Several important extensions are left for future work. First, we estimate only the direct effect of CWA stringency on project duration, holding project scale fixed. If stricter regulation also induces developers to downsize projects, the total effect on investment would be larger. Second, our structural model assumes constant returns to scale and a fixed mass of entrants, abstracting from extensive-margin entry and technology choice. Third, the CWA is one of several federal and state permitting requirements facing electric generators; the framework developed here could in principle be applied to other statutes, such as the National Environmental Policy Act, where quantitative stringency measures become available.

9 Appendix

9.1 Estimation Algorithm

This section describes the scalar-path algorithm used to evaluate the log-likelihood for each project. Because project phases $(x_1, \ldots, x_{d_i})$ are observed, each transition contributes a single scalar entry $\alpha^{(l)}_{x_t,x_{t+1}}$ rather than a full matrix power, reducing the per-project computational cost.

Step 1 — First-block log-product. The project is locked in regime l_1 for all b_1 periods of the first calendar block. The contribution of this certain block is accumulated as a single log-sum:

$$L_1 = \log \pi_{x_1} + \sum_{t=1}^{b_1} \log \alpha^{(l_1)}_{x_t,x_{t+1}}. \tag{35}$$

Step 2 — Suffix log-products for each target regime. After the first block, the project may switch permanently to any target regime $l_{\text{tgt}} \neq l_1$ that the calendar offers. For each such target, define the suffix log-product

$$\text{suf}_{l_{\text{tgt}}}(\kappa) = \sum_{t=b_1+\kappa+1}^{d_i-1} \log \alpha^{(l_{\text{tgt}})}_{x_t,x_{t+1}}, \quad \kappa = 0, 1, \ldots, s_{\text{after}}, \tag{36}$$

where $s_{\text{after}} = d_i - b_1 - 1$ is the number of uncertain transition steps after the first block. The suffix $\text{suf}_{l_{\text{tgt}}}(\kappa)$ is the log-product of l_{tgt}-regime entries from period $b_1 + \kappa + 1$ through period $d_i - 1$ — the portion of the path that would be spent in the target regime if the switch occurred at step κ. This array is built right-to-left in a single pass, with $\text{suf}_{l_{\text{tgt}}}(s_{\text{after}}) = 0$ (empty product when no target-regime steps remain).

Step 3 — Switch-time forward loop. A running prefix $\text{pre}_1(\kappa) = \sum_{t=b_1+1}^{b_1+\kappa} \log \alpha^{(l_1)}_{x_t,x_{t+1}}$ accumulates the additional l_1-regime steps taken after the first block before switching. This prefix is shared across all target regimes, so it is computed only once. Setting the log-reference

$L_{\text{ref}} = L_1$, the likelihood sum is accumulated in a numerically stable log-sum-exp form:

$$S = \exp\big(L_1 + \text{pre}_1(s_{\text{after}}) - L_{\text{ref}} + \log e^{(l_1,\omega_i)}_{x_{d_i}}\big) \qquad \text{(no-switch)}$$
$$+ \sum_{l_{\text{tgt}}} \sum_{\kappa=\kappa_{\min}}^{s_{\text{after}}-1} \exp\big(L_1 + \text{pre}_1(\kappa) + \text{suf}_{l_{\text{tgt}}}(\kappa) - L_{\text{ref}} + \log e^{(l_{\text{tgt}},\omega_i)}_{x_{d_i}}\big), \qquad \text{(switch terms)}$$

where $e^{(l,\omega_i)}_{x_{d_i}}$ is the exit-probability vector entry defined in equation (6). The log-likelihood contribution of project i is then $\log \mathcal{L}_i = L_{\text{ref}} + \log S$. The lower bound $\kappa_{\min} = \max(0, s_{\text{after}} + 1 - B_{\text{tgt}})$ enforces the calendar capacity constraint: the project cannot spend more total periods in l_{tgt} (including the terminal period) than the calendar offered.

Right-censored observations. For a right-censored project (one still active at the end of the sample), no absorption event is observed and the terminal factor $e^{(l,\omega_i)}_{x_{d_i}}$ is dropped entirely. Setting this term to $e^0 = 1$ in the formulas above gives the survival probability $P(D_i > d_i)$ as the likelihood contribution, which is the correct censored likelihood.

Optimisation. The full log-likelihood $\ell(\theta) = \sum_{i=1}^{n} \log \mathcal{L}_i(\theta)$ is maximised numerically using `fminunc` with the trust-region algorithm. Because the likelihood surface for phase-type models can be multimodal (Faddy and McClean, 1999), the optimisation is run from multiple random starting points and the solution with the highest log-likelihood value is retained. Standard errors are computed from the inverse of the numerical Hessian at the optimum: $\widehat{\text{Cov}}(\hat{\theta}) = \hat{H}^{-1}$, with Wald 95% confidence intervals $\hat{\theta}_k \pm 1.96\sqrt{[\hat{H}^{-1}]_{kk}}$.

References

Omar Al-Ubaydli and Patrick A McLaughlin. Regdata: A numerical database on industry-specific regulations for all united states industries and federal regulations, 1997–2012. *Regulation & Governance*, 11(1):109–123, 2017.

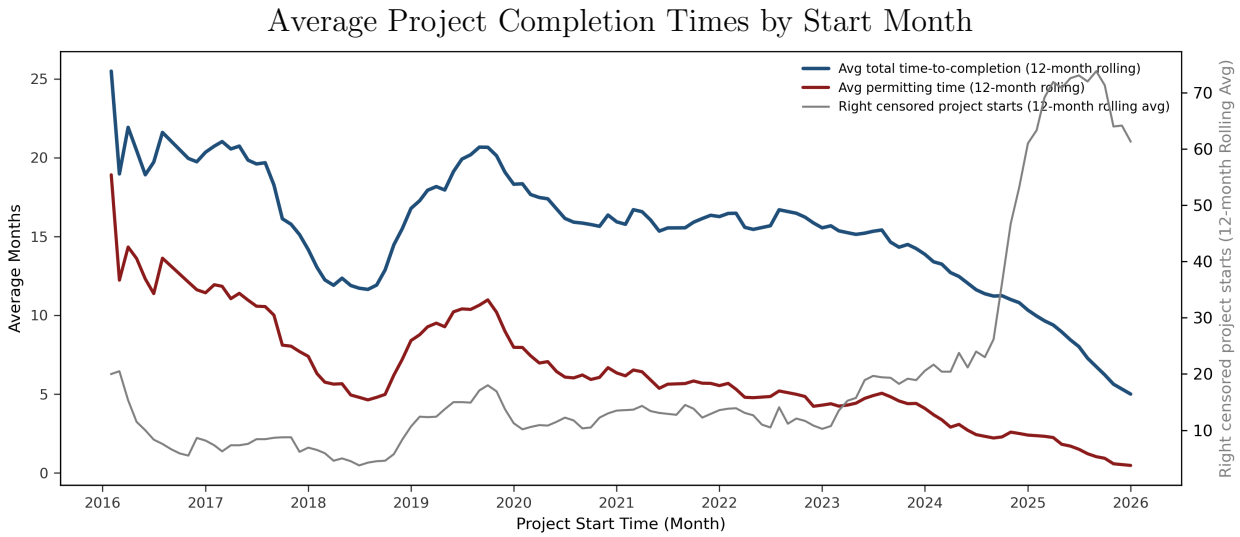


Figure 3: This is the caption for my PNG figure.

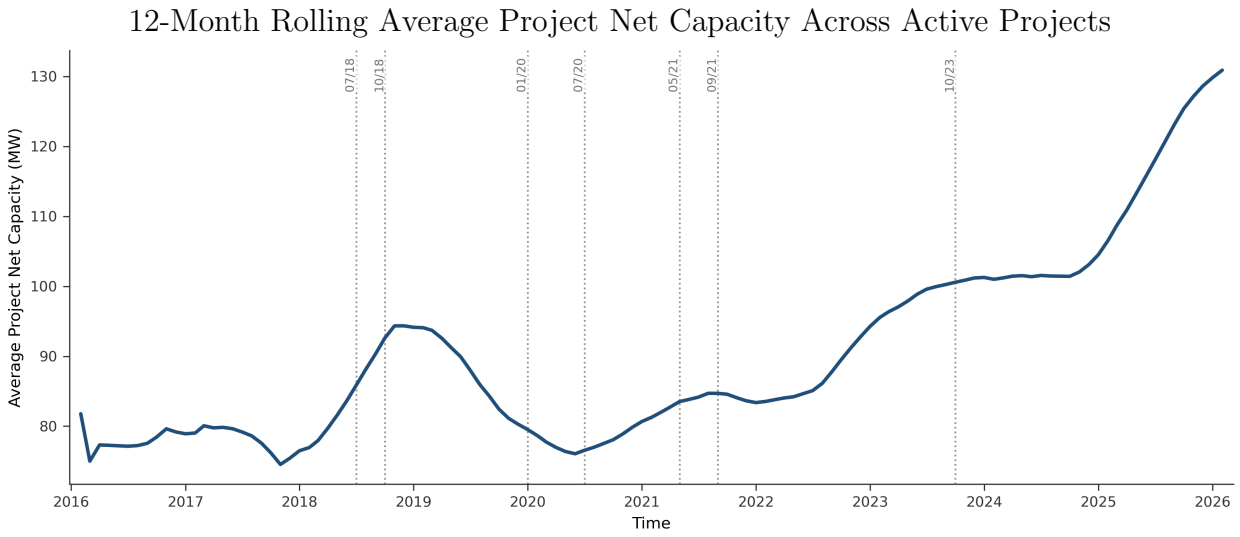


Table 5: Clean Water Act Regulatory Regime Changes in Empirical Coding

Date	Regime	Description	States Affected
Pre-2015	Rapanos	Baseline regime. Federal jurisdiction over Waters of the United States governed by the 2008 EPA/Army Corps guidance implementing the split Rapanos v. United States (2006) decision.	All States
Aug. 28, 2015	Clean Water Rule	Clean Water Rule (CWR) takes effect, broadening WOTUS jurisdiction under the significant-nexus standard. Thirteen states obtained injunctions and remained under Rapanos guidance throughout this period.	All but AK, AZ, AR, CO, ID, MO, MT, NE, NV, NM, ND, SD, WY
Jun. 8, 2018	Rapanos	Court injunction granted; eleven additional states revert from the Clean Water Rule to the Rapanos framework for the remainder of the CWR window (through Dec. 23, 2019).	AL, FL, GA, IN, KS, KY, NC, SC, UT, WV, WI
Sep. 12, 2018	Rapanos	Court injunction granted; four additional states revert from the Clean Water Rule to the Rapanos framework for the remainder of the CWR window (through Dec. 23, 2019).	IA, LA, MS, TX
Dec. 23, 2019	Rapanos	2019 Repeal Rule takes effect nationwide, formally rescinding the Clean Water Rule and returning all states to the Rapanos/2008 guidance framework pending a replacement rule.	All States
Jun. 22, 2020	NWPR	Navigable Waters Protection Rule (NWPR) takes effect, narrowing WOTUS jurisdiction relative to both the CWR and Rapanos guidance. Colorado was blocked by separate litigation and remained under Rapanos until April 2021.	All but CO
Apr. 26, 2021	NWPR	NWPR takes effect in Colorado after state-level litigation resolved, completing nationwide application of the rule.	CO only
Aug. 30, 2021	Rapanos	Federal court vacates the NWPR; on Sep. 3, 2021 agencies confirm nationwide effect. All states revert to the pre-2015 Rapanos/2008 guidance framework.	All States
Sep. 8, 2023	Sackett	Conforming rule implementing Sackett v. EPA (2023) takes effect, codifying the Supreme Court's narrowed definition of WOTUS and removing the significant-nexus test.	All States

M. Bennon, D. De La Hormaza, and R. R. Geddes. A hazard analysis of federal permitting under the National Environmental Policy Act of 1970. *Journal of Regulatory Economics*, 65(1):154–183, 2024. doi: 10.1007/s11149-023-09465-y.

P. Burnham and D. Carloni. CBO's model for estimating the effect that federal taxes have on capital income from new investment. Working Paper 2022-01, Congressional Budget Office, Washington, D.C., 2022. URL `https://www.cbo.gov/publication/57429`.

M. Cacciatore and G. Fiori. The macroeconomic effects of goods and labor markets deregulation. *Review of Economic Dynamics*, 20:1–24, 2016. doi: 10.1016/j.red.2015.10.002.

G. Chodorow-Reich. The neoclassical theory of firm investment and taxes: A reassessment. Working Paper w33922, National Bureau of Economic Research, Cambridge, Mass., 2025. URL `https://www.nber.org/papers/w33922`.

Congressional Budget Office. Key methods that CBO used to estimate the macroeconomic effects of the 2017 Tax Act: Supplemental material for *The Budget and Economic Outlook: 2018 to 2028*. Technical report, Congressional Budget Office, Washington, D.C., 2018. URL `https://www.cbo.gov/publication/53651`.

A. K. Dixit and J. E. Stiglitz. Monopolistic competition and optimum product diversity. *American Economic Review*, 67(3):297–308, 1977.

M. J. Faddy and S. I. McClean. Analysing data on lengths of stay of hospital patients using phase-type distributions. *Applied Stochastic Models in Business and Industry*, 15(4):311–317, 1999. doi: 10.1002/(SICI)1526-4025(199910/12)15:4⟨311::AID-ASMB395⟩3.0.CO;2-S.

Arthur Fraas, Valkyrie Buffa, and Lindsay Rich. Establishing utility-scale solar projects: Federal involvement. Technical report, Resources for the Future, 2021.

Stuart A. Gabriel and Ed Kung. Development approval timelines, approval uncertainty, and new housing supply: Evidence from Los Angeles. Working Paper, 2024.

D. Glancy, R. J. Kurtzman, and L. Loewenstein. On commercial construction activity's long and variable lags. FEDS Working Paper 2024-016, Federal Reserve Board, Washington, D.C., 2024. URL https://www.clevelandfed.org/publications/working-paper/2024/wp-2414-commercial-construction-activity-lags.

Simon Greenhill, Brant J Walker, and Joseph S Shapiro. Deep learning projects jurisdiction of new and proposed clean water act regulation. Working Paper 34947, National Bureau of Economic Research, March 2026. URL http://www.nber.org/papers/w34947.

Michael Greenstone. The impacts of environmental regulations on industrial activity: Evidence from the 1970 and 1977 clean air act amendments and the census of manufactures. *Journal of political economy*, 110(6):1175–1219, 2002.

R. E. Hall and D. W. Jorgenson. Tax policy and investment behavior: Reply and further results. *American Economic Review*, 59(3):388–401, 1967.

Stephen P. Jenkins. Easy estimation methods for discrete-time duration models. *Oxford Bulletin of Economics and Statistics*, 57(1):129–138, 1995. doi: 10.1111/j.1468-0084.1995.tb00031.x.

Sarah Johnston, Yifei Liu, and Chenyu Yang. An empirical analysis of the interconnection queue. Technical report, National Bureau of Economic Research, 2023.

M. Kalouptsidi. Time to build and fluctuations in bulk shipping. *American Economic Review*, 104(2):564–608, 2014. doi: 10.1257/aer.104.2.564.

Tony Lancaster. *The Econometric Analysis of Transition Data*. Number 17 in Econometric Society Monographs. Cambridge University Press, Cambridge, 1990. doi: 10.1017/CCOL0521265967.

Eric Lewis. Patchwork policies, spillovers, and the search for oil and gas. *American Economic Journal: Economic Policy*, 11(1):380–405, 2019.

Robi Nilson, Lanbing Tao, Will Maddock, Juniper Katz, Natalie Baillargeon, and Ben Hoen. Outcomes and timelines for state-based energy facility permitting in the united states. *Frontiers in Sustainable Energy Policy*, 5:1715811, 2026.

R. L. Prentice and L. A. Gloeckler. Regression analysis of grouped survival data with application to breast cancer data. *Biometrics*, 34(1):57–67, 1978. doi: 10.2307/2529588.

Joseph Rand, Rose Strauss, Will Gorman, Joachim Seel, Julie Mulvaney Kemp, Seongeun Jeong, Dana Robson, and Ryan H. Wiser. Queued up: Characteristics of power plants seeking transmission interconnection as of the end of 2022. Technical Report LBNL-2001558, Lawrence Berkeley National Laboratory, 2023.

Stephen P Ryan. The costs of environmental regulation in a concentrated industry. *Econometrica*, 80(3):1019–1061, 2012.

S. Schmitt-Grohé and M. Uribe. Closing small open economy models. *Journal of International Economics*, 61(1):163–185, 2003. doi: 10.1016/S0022-1996(02)00056-9.

F. Trebbi and M. B. Zhang. The cost of regulatory compliance in the United States. Working Paper w30691, National Bureau of Economic Research, Cambridge, Mass., 2022a. URL `https://www.nber.org/papers/w30691`.

Francesco Trebbi and Miao Ben Zhang. The cost of regulatory compliance in the united states. Technical report, National Bureau of Economic Research, 2022b.