

Decision-Relevant Information in Generator Interconnection: Value, Calibration, and Public-Data Limits

Benjamin Xie* Jonathan Tsang*

University of California, Berkeley

benx@berkeley.edu jy_tsang@berkeley.edu

October 2, 2026

Abstract

When is a better interconnection cost estimate worth paying for? We study a developer that can buy optional information before deciding whether to proceed with a generation project. In a risk-neutral model, information has value when it improves this decision; predictive accuracy alone need not measure that value. We specialize classical value-of-information theory to the project’s economic threshold and show that a cost distribution alone permits any information value between zero and the perfect-information bound. We then distinguish a predicted cost score’s implied value from the expected payoff of following its recommendation. When predicted and actual costs have the same mean, their value discrepancy is at most half the mean absolute conditional-calibration error; the bound is sharp. In an exact four-state example, a signal with one-eighth the mean-squared prediction error has lower decision value. A separate retrospective analysis of public PJM records yields 162 comparable impact–facilities study pairs. The median absolute estimate change is 14.11 real-2020 dollars/kW; 28 pairs (17.3 percent) cross an illustrative 100 dollars/kW threshold. These selected historical estimates describe revisions, but do not establish realized-cost forecast errors or the value of a particular advisory service. Together, the analysis and data audit identify the decision, timing, calibration, and outcome evidence needed to evaluate optional interconnection information.

Keywords: generator interconnection; value of information; decision-focused prediction; calibration; PJM; cost uncertainty.

1 Introduction

A generator developer needs to know whether an interconnection bill will leave a project worth pursuing. An additional estimate is useful if it helps the developer make that decision before committing more resources. Even a very accurate estimate can have little decision value when every plausible result leads to the same action. A less accurate estimate can be valuable if it reliably separates projects that should proceed from those that should stop.

Interconnection costs and queue processes are consequential barriers to generation development. Recent work documents these barriers across the United States and examines detailed PJM cost records [1, 4]. Berkeley Lab’s PJM cost study provides an accessible public dataset, while its continuing queue reports place individual projects within a changing national development pipeline

*Equal contribution; both authors are first authors.

[2, 3]. These sources motivate careful evaluation of information, but neither a large cost spread nor a high withdrawal rate by itself establishes willingness to pay for a particular forecast, engineering screen, or advisory service.

We study one decision: whether to proceed with a project after receiving optional information about its interconnection cost. The developer compares expected cost with a known economic threshold—the amount the project can afford to spend on interconnection. The model assumes risk neutrality and a fixed cost environment. It isolates the value of learning before a remaining commitment; it does not model the full queue process.

Main findings. Information can be valuable even when it leaves substantial cost uncertainty, provided it identifies the better action. Conversely, a forecast can get average cost right and still recommend harmful decisions. Public study records can show how often estimates cross illustrative thresholds, but evaluating an advisory service also requires evidence about when its information was available, the decision it could change, and the resulting payoff.

1.1 Relationship to prior work

The underlying theory is established. Blackwell’s comparison of experiments formalizes information ordering, and classical decision analysis values information by its effect on optimized expected utility [7, 8]. Decision-focused prediction likewise distinguishes statistical prediction loss from downstream optimization loss [9, 10]. Uncertainty and information have also been studied in electricity investment, transmission planning, and constrained energy-system design [11, 12, 13].

In the interconnection setting, Johnston, Liu, and Yang model learning and project continuation in an empirical analysis of the queue [5]. Sun’s available conference abstract describes a multistage Bayesian-learning model and counterfactual provision of earlier cost information in PJM [6]. These studies already connect earlier information to continuation decisions. Our focus is the evaluation of an optional signal at a specified decision, the gap between a score’s implied and attained value, and the limits of public study estimates. We do not estimate a dynamic queue model or causal savings from an information intervention.

1.2 Contributions and scope

The paper makes three application-specific contributions:

1. **A decision-based evaluation framework.** Threshold identities, an identification limit, and bounds on score-implied value all use the same no-purchase baseline. The elementary information-value results specialize established theory; they are not claims of new general decision theory.
2. **Exact examples separating accuracy from value.** Constructed signals show why prediction error can misrank information and why matching average cost is insufficient for useful decisions.
3. **A reproducible public-data audit.** We distinguish a latest-estimate snapshot from a study-stage panel and construct comparable study pairs using explicit cost components, dates, and capacity restrictions.

The empirical exercise is intentionally descriptive. For 162 retained impact–facilities pairs, the median absolute difference between study estimates is 14.11 real-2020 dollars/kW. At an illustrative threshold of 100 dollars/kW, 17.3% switch sides. These numbers show how a threshold-oriented description differs from a single average revision. They are conditional on receiving both studies and satisfying our filters; they do not show that the later estimate is truth, that a developer changed its action, or that an optional advisor would have supplied the same information earlier.

Readers interested primarily in the application can begin with the examples in Section 5, the public-data findings in Section 6, and the evaluation questions in Section 7. Sections 3 and 4 give the formal model and guarantees. All proposition and corollary proofs appear in Appendix A.

2 The interconnection decision

Interconnection studies assess the network and connection changes associated with adding a generator. A project’s relevant cost may include both point-of-interconnection facilities and allocated network upgrades. Those components must be distinguished from the total cost of a set of upgrades shared with other requests. Our empirical section uses the sum of point-of-interconnection cost and the request’s allocated network cost, where both are explicitly recorded.

Study and decision structures vary across market operators and over time. FERC Order No. 2023 introduced major changes to generator interconnection procedures, including cluster-study and readiness provisions [14]. The PJM cycle-process manual available at the time of this analysis is Revision 07, effective August 19, 2026, and organizes the process around study phases and decision points [15]. The historical impact and facilities records analyzed below should not be interpreted as observations from that current process.

The modeled acquisition is an *optional advisory signal*: for example, an additional engineering screen or an independent interpretation of available network information before a permitted development decision. It does not replace mandatory operator studies, required agreements, or official cost allocation. The model asks whether a developer would benefit from buying additional information, given the actual decision and obligations it already faces.

At that decision, accounting must separate sunk expenditure, avoidable future expenditure, and recoverable payments. A deposit is not automatically an economic cost equal to its face value; its treatment depends on refund and forfeiture rules. Likewise, delay can reduce a project’s value only if development continues, or impose an unconditional acquisition expense, or change the project’s queue circumstances. These cases have different payoffs. We initially model only an unconditional information-acquisition charge. Appendix B treats a deterministic loss incurred only if the project proceeds.

The target variable must also be defined before evaluating forecasts. In the model, C is the cost entering the project’s payoff. In the public-data analysis, the observed variables are successive *study estimates*, denoted Q_I and Q_F . We do not identify either estimate with C or assume it is a calibrated conditional expectation. This distinction is essential: a change in a study estimate may reflect new engineering information, a changed allocation group, a changed project, or a changed network.

3 A single-decision information model

3.1 Payoffs and assumptions

Let $C \geq 0$ be the uncertain interconnection cost, with finite expectation, and let Y be an optional signal about that cost. All expectations condition on the project’s current information, which we suppress in the notation. A cost distribution across different projects is therefore not automatically the prior distribution for any one project. Table 1 collects the main notation.

Let $b \in \mathbb{R}$ be the known incremental discounted benefit available to cover interconnection cost, measured relative to abandoning. It includes the other project revenues and avoidable costs relevant

Table 1: Main notation. All payoffs and costs use a common monetary unit and date.

Symbol	Meaning
$C, \mu = \mathbb{E}C$	Uncertain interconnection cost and its current expected value.
b	Net project benefit before interconnection cost: the economic threshold.
$Y, M = \mathbb{E}[C \mid Y]$	Optional information and expected cost after receiving it.
$B = (b - \mu)_+$	Optimal expected payoff without the optional information.
$V(Y), h$	Gross information value under optimal use, and the acquisition fee.
$Z, g(Z) = \mathbb{E}[C \mid Z]$	Predicted cost score and the conditional expected cost given that score.
J, G	Score-implied value and attained expected gain, both relative to B .
ε_{cal}	Mean absolute calibration error, $\mathbb{E} g(Z) - Z $.

to this comparison. For an action $a \in \{0, 1\}$, where 1 means proceed, the incremental payoff is

$$u(a, C) = a(b - C). \quad (1)$$

All monetary quantities use the same date and units. A dollars-per-kW formulation requires normalizing both benefit and cost by the same capacity; it does not make different definitions of capacity interchangeable.

Write $\mu = \mathbb{E}C$ and $x_+ = \max\{x, 0\}$. A risk-neutral developer without the optional signal obtains

$$B = (b - \mu)_+. \quad (2)$$

After observing the signal, expected cost becomes the posterior mean $M = \mathbb{E}[C \mid Y]$. The optimal action is $a^*(Y) = \mathbf{1}\{M < b\}$, where $\mathbf{1}\{\cdot\}$ is one when the stated condition holds and zero otherwise. Throughout the model, ties are resolved by abandoning. The gross expected value of sample information is

$$V(Y) = \mathbb{E}[(b - M)_+] - B. \quad (3)$$

If acquisition imposes an unconditional total charge $h \geq 0$, its net incremental value is $V(Y) - h$. The project buys the information when this quantity is positive and is indifferent at zero. The fee is incurred whether the subsequent choice is proceed or abandon. The value of having the option to buy, relative to no access, is $\max\{0, V(Y) - h\}$.

This formulation assumes risk neutrality, a fixed cost-generating environment during the modeled choice, no strategic response by other projects, no financing constraint, and an action that remains feasible after acquisition. The signal changes beliefs rather than engineering outcomes. Additional downstream learning and endogenous waiting are outside the model. If revenue is uncertain, the target should instead be the conditional expectation of the entire incremental net payoff; a cost-only threshold need not suffice.

3.2 Threshold identities and information bounds

The first result expresses information value in terms of posterior expected costs that cross the project's threshold. This makes the economic role of the signal explicit.

Proposition 3.1 (Threshold representation). *The gross value in (3) satisfies*

$$V(Y) = \frac{\mathbb{E}|b - M| - |b - \mu|}{2} = \begin{cases} \mathbb{E}[(M - b)_+], & \mu \leq b, \\ \mathbb{E}[(b - M)_+], & \mu \geq b. \end{cases} \quad (4)$$

Both cases agree when $\mu = b$.

When the prior favors proceeding, information is valuable through states in which posterior expected cost exceeds b . When the prior favors abandoning, the relevant states are those in which it falls below b . If all posterior means remain on the same side of the threshold, gross value is zero in this model. This does not exclude other uses of information in engineering, financing, contracting, or risk management.

Proposition 3.2 (Bounds and refinement). *Define the perfect-information value by*

$$V_{\text{PI}} = \mathbb{E}[(b - C)_+] - (b - \mu)_+. \quad (5)$$

Then $0 \leq V(Y) \leq V_{\text{PI}}$. If $C \rightarrow Y_2 \rightarrow Y_1$ is a Markov chain, so that Y_1 is a garbling of Y_2 , then $V(Y_1) \leq V(Y_2)$. If $v = \text{Var}(M) < \infty$, then

$$V(Y) \leq \frac{\sqrt{(b - \mu)^2 + v} - |b - \mu|}{2}. \quad (6)$$

The variance bound is useful for screening out signals whose value cannot cover their fee. When $b \neq \mu$, rationalizing its numerator gives the weaker bound $V(Y) \leq v/(4|b - \mu|)$. At $b = \mu$, the bound is $\sqrt{v}/2$. These are upper bounds, not estimates of value. With finite, fixed $\text{Var}(C)$, better squared-error prediction by a true conditional mean raises $\text{Var}(M)$, but actual decision value need not rise. Section 5 gives an exact example.

3.3 What a marginal cost distribution cannot identify

Proposition 3.3 (Sharp nonidentification from the marginal). *Fix the distribution of C and a threshold b . If the conditional signal law is unrestricted, the set of possible values of $V(Y)$ is exactly $[0, V_{\text{PI}}]$.*

Even a perfectly measured marginal cost histogram cannot identify the value of a particular information product without assumptions about its relationship to cost. It can provide a perfect-information ceiling under an assumed project prior and benefit threshold. Historical cross-sectional heterogeneity supplies neither assumption automatically. In particular, fitting a distribution across withdrawn and completed requests and treating it as one prospective project's uncertainty introduces a substantive modeling assumption, not a measurement of that project's beliefs.

4 Predicted cost scores and actual decision value

An advisory product supplies a predicted cost, which may differ from the true expected cost conditional on its information. Evaluating the recommendation therefore requires separating the payoff the score implies from the payoff the developer can expect to attain. Calibration relative to downstream decisions is an established theme [16].

4.1 Implied value and attained gain

Let $Z = Z(Y)$ be an integrable predicted cost score. A user proceeds when $a_Z = \mathbf{1}\{Z < b\}$. Treating the score as expected cost gives the *score-implied gross value*

$$J = \mathbb{E}[(b - Z)_+] - B. \quad (7)$$

The same recommendation yields the *attained gross gain*

$$G = \mathbb{E}[(b - C)a_Z] - B = \mathbb{E}[(b - M)a_Z] - B. \quad (8)$$

Both are measured relative to the optimal no-purchase payoff B . Here, “attained” refers to expected payoff under the score’s policy, not a realized payoff in one project. Replacing B with $(b - \mathbb{E}Z)_+$ would change the comparison. The quantity J can overstate or understate gain; the policy’s net expected gain from acquisition is $G - h$.

4.2 How calibration bounds the value discrepancy

A score is conditionally calibrated when projects assigned that score have the same conditional expected cost. Write this expected cost as $g(Z) = \mathbb{E}[C \mid Z]$. The calibration error $\varepsilon_{\text{cal}} = \mathbb{E}|g(Z) - Z|$ measures the average absolute discrepancy. Matching the overall mean, $\mathbb{E}Z = \mathbb{E}C$, is a weaker requirement: errors at different scores can cancel.

Proposition 4.1 (Attained-policy certificate). *Let $g(Z) = \mathbb{E}[C \mid Z]$, $\varepsilon_{\text{cal}} = \mathbb{E}|g(Z) - Z|$, and $\delta = \mu - \mathbb{E}Z$. Then*

$$J - G = \mathbb{E}[(g(Z) - Z)\mathbf{1}\{Z < b\}], \quad (9)$$

and

$$J - \frac{\varepsilon_{\text{cal}} + \delta}{2} \leq G \leq J + \frac{\varepsilon_{\text{cal}} - \delta}{2}. \quad (10)$$

If the scores are mean-matched, $\mathbb{E}Z = \mu$, then

$$|J - G| \leq \frac{\varepsilon_{\text{cal}}}{2}. \quad (11)$$

The factor $1/2$ cannot be reduced uniformly over this class of problems.

Consequently, under population mean matching, a justified bound $\varepsilon_{\text{cal}} \leq \bar{\varepsilon}$ gives the sufficient acquisition condition

$$J > h + \frac{\bar{\varepsilon}}{2}. \quad (12)$$

Thus, the score-implied value must cover both the acquisition fee and a margin for calibration error. This guarantee depends on population quantities. To use it empirically, an evaluation must justify the error bound, quantify uncertainty in J , and address estimation of μ . Training error is not a substitute for a valid error bound, and finite-sample mean matching does not establish population mean matching.

4.3 Calibration and loss from discarding information

Calibration addresses whether the score’s value calculation is credible. A second question is whether the score uses all the decision-relevant information in Y . The following result separates these questions.

Table 2: A more accurate signal can have lower decision value. The threshold is $b = 50$ and the fee is $h = 13.5$, in illustrative dollars/kW. MSE has squared units. Net gain assumes purchase; choosing no purchase gives incremental gain zero.

Signal	MSE	Gross value V	Net gain	Buy?
A	50	12.5	-1.0	No
B	400	15.0	1.5	Yes
Perfect revelation	0	15.0	1.5	Yes

Corollary 4.2 (Calibration, discarded information, and regret). *Let $\varepsilon_Y = \mathbb{E}|M - Z|$. Then $\varepsilon_{\text{cal}} \leq \varepsilon_Y$ and*

$$0 \leq V(Y) - G = \mathbb{E}[|M - b| \mathbf{1}\{a_Z \neq a^*(Y)\}] \leq \varepsilon_Y. \quad (13)$$

Also $\varepsilon_{\text{cal}} \leq \mathbb{E}|C - Z|$. If the score is conditionally calibrated, $g(Z) = Z$ almost surely, then $J = G$; this need not imply $G = V(Y)$.

The bound ε_{cal} controls the gap between score-implied and attained value. The bound ε_Y controls regret—the expected payoff lost relative to optimal use of all information in Y . In particular, ε_{cal} cannot replace ε_Y in the regret bound. Held-out cost-score pairs can help assess calibration, although raw absolute prediction error also includes irreducible cost variation and may give a conservative bound. Valid finite-sample inference requires assumptions on sampling, tails, and distribution shift that the descriptive data below do not supply.

5 Worked examples: accuracy, calibration, and value

These examples are constructed to isolate mechanisms. They are not fitted to PJM data, do not represent observed advisor performance, and use illustrative dollars/kW without an empirical price-year interpretation. Let

$$\mathbb{P}(C = c) = \frac{1}{4}, \quad c \in \{0, 40, 60, 100\}, \quad b = 50. \quad (14)$$

Thus $\mu = 50$, $B = 0$, and $V_{\text{PI}} = 15$.

5.1 Prediction accuracy can reverse the decision ranking

Signal A identifies costs of 0 and 100 exactly, but cannot distinguish 40 from 60. Signal B identifies whether cost is in the lower pair, $\{0, 40\}$, or the upper pair, $\{60, 100\}$. Their posterior means, in cost-state order, are $(0, 50, 50, 100)$ and $(20, 20, 80, 80)$. Both are true conditional expectations. Table 2 compares mean-squared error (MSE), $\mathbb{E}(C - M)^2$, with decision value.

Signal A has one-eighth the MSE of B, yet its gross decision value is smaller. It distinguishes the extreme costs exactly but pools the two states on opposite sides of the decision threshold. Signal B leaves substantial within-group cost uncertainty but reveals the preferred action in every state. It therefore attains perfect *decision* value without revealing the exact cost. Neither signal is a garbling of the other, so the example does not contradict Blackwell monotonicity.

The ordering also depends on the project’s benefit threshold. Figure 1 varies b while holding the cost prior and signal partitions fixed. A is more valuable at some thresholds, while B is more valuable near 50. A universal information-product ranking based only on one accuracy number or one project’s value therefore misses the decision context.




Figure 1: Exact value curves for the constructed signals. Curves are piecewise linear; displayed vertices suffice to reconstruct them. No parameter is estimated from interconnection records.

5.2 Matching average cost can still produce harmful decisions

Suppose Y reveals C perfectly, but an intentionally misspecified score reports $Z = 100 - C$. Then $\mathbb{E}Z = \mathbb{E}C = 50$, yet the score proceeds precisely in the $C = 60$ and $C = 100$ states. Its implied value is $J = 15$, while its attained gain is

$$G = \frac{(50 - 60) + (50 - 100)}{4} = -15.$$

Because C is a function of Z , $g(Z) = C$ and

$$\varepsilon_{\text{cal}} = \frac{100 + 20 + 20 + 100}{4} = 60.$$

The lower bound in (11) is attained: $G = J - \varepsilon_{\text{cal}}/2 = -15$. This is an adversarial stress test of mean matching, not a proposed model of realistic errors. Perfect information remains worth 15 when used correctly; harm comes from following the wrong score.

For a less extreme sensitivity calculation, suppose a mean-matched product reports $J = 15$ and acquisition costs $h = 13.5$. An independently justified calibration-error bound of 0, 1, 2, or 3 implies lower bounds on net gain of 1.5, 1.0, 0.5, or 0, respectively. Strict profitability is certified only when the bound is below 3. This calculation illustrates error accounting; it does not estimate an error radius for a real product.

5.3 One prior, a continuum of possible values

For the reveal-or-null signal in Proposition 3.3, the same four-state prior gives $V(Y) = 15r$. At a fee of 5, acquisition is strictly worthwhile if $r > 1/3$. The marginal cost distribution is identical for every r . Any inference that the marginal alone establishes one of these acquisition decisions therefore requires an unstated assumption about the signal. Independence of the reveal event from cost is essential to this particular construction; otherwise silence itself can be informative.

The accompanying standard-library Python calculation uses exact rational arithmetic. It evaluates all 15 partitions of the four states, 101 integer thresholds, and three strictly positive priors, giving 4,545 base cases, with six score variants per case giving 27,270 error-bound cases. These

checks support implementation accuracy; the proofs establish the general claims. Appendix C gives the finite-state formulas and a minimal standalone calculation.

6 Historical evidence from public PJM records

The empirical exercise asks a descriptive question: how much do comparable study estimates change, and how often do they cross illustrative economic thresholds? It evaluates historical records, not the performance of an advisory product.

6.1 Two distinct public data structures

We audited two public sources accessed October 2, 2026. Berkeley Lab’s PJM workbook contains 1,127 distinct requests with one latest available study estimate per request [2]. Its codebook defines the study date accordingly. That structure supports a cross-sectional description but not a sequence of within-project revisions. Its costs are expressed in real-2022 dollars/kW of nameplate capacity.

The public study-stage data released by Johnston, Liu, and Yang provide a different structure [17]. The base file contains 7,492 records for 4,174 distinct requests, unique by request and study type. The documentation describes requests submitted in 2008–2020 and engineering documents available as of August 2024. It retains the most recent version of each study stage, rather than every historical version. These records permit selected study-stage comparisons, but they do not provide a complete reconstruction of what a developer knew on each decision date.

We use only separate system-impact and facilities records, excluding combined-stage classifications from the paired analysis. This choice avoids assigning ambiguous combined reports to a particular transition; it is not a judgment that combined reports are invalid. We also avoid a general feasibility-to-impact comparison because early records commonly report a combined upgrade amount rather than the individual request’s allocated network amount. Differencing these fields would mix cost definitions.

6.2 A conservative comparable-pair sample

Table 3 reports the cumulative restrictions. We use base-file costs and dates rather than the augmented file’s zero-filled costs or imputed dates. Requested capacity is the maximum of the reported energy MW and capacity MW, following the source construction. It is a study-request size, not an independently verified installed nameplate capacity. We require it to be positive and unchanged across the retained stages, and exclude uprates.

The final reconciliation tolerance is 10^{-6} in the source’s nominal-million-dollar units. We exclude discrepancies rather than correct them. A discrepancy does not show that the original engineering study is wrong; it shows that the extracted record fails our deliberately narrow consistency check. No missing component is interpreted as zero.

For stage $s \in \{I, F\}$, let T_s be the source total in nominal millions of dollars, A_s the source CPI factor converting that year’s dollars to 2020 dollars, and K_s the requested MW. We calculate

$$Q_s = \frac{1000 T_s A_s}{K_s} \quad \text{real-2020 dollars/kW.} \quad (15)$$

The factor 1000 converts millions of dollars per MW to dollars per kW. We borrow only inflation factors and descriptive snapshot labels from the augmented file, and check the inflation year against each retained raw study date. We do not combine these values numerically with the Berkeley Lab workbook’s 2022-dollar observations.

Table 3: Cumulative construction of the retrospective impact–facilities sample. Filters were developed during this exploratory audit, not preregistered.

Restriction	Pairs
Separate impact and facilities records for the same request	778
Both raw total cost estimates present	723
Valid raw year–month dates; facilities strictly later	624
Positive requested capacity, unchanged between stages	544
No uprate at either stage	496
Available positive source inflation factors	496
Explicit POI and allocated network components at both stages	213
Total agrees with POI plus network within one nominal dollar at both stages	162

Table 4: Selected paired study estimates, $n = 162$. All costs are real-2020 dollars/kW of requested capacity. Quantiles use linear interpolation. A paired-change median is not the difference of the two marginal medians.

Quantity	10th pct.	25th pct.	Median	75th pct.	90th pct.
Impact estimate Q_I	14.83	33.29	71.67	121.72	167.87
Facilities estimate Q_F	11.84	28.86	56.01	105.68	139.89
Signed change Δ	−78.05	−26.41	−6.50	4.57	23.94
Absolute change $ \Delta $	1.55	5.90	14.11	39.96	89.98

6.3 Study-estimate differences and threshold crossings

The primary sample contains 118 requests labeled solar, 19 wind, 17 natural gas, and eight storage in the source primary-fuel field. These labels need not describe every component of a hybrid project. Requested capacity has a median of 80 MW, an interquartile range of 50–130 MW, and a range of 12–1,681 MW. Retained impact dates span 2010–2023 and facilities dates span 2011–2023. This composition further limits extrapolation to other project types or current queue cohorts.

For the primary sample of 162 pairs, define $\Delta_i = Q_{F,i} - Q_{I,i}$. Table 4 shows marginal and paired summaries. The median signed change is −6.50 dollars/kW; the median absolute change is 14.11. The signed 10th and 90th percentiles are −78.05 and 23.94. There are 59 increases, 102 decreases, and one unchanged estimate using a numerical tolerance of 10^{-8} dollars/kW. The median interval between the retained study dates is 14 months. These are intervals between recorded study dates, not estimates of the delay caused by optional information acquisition.

To illustrate threshold sensitivity without assuming project revenues, we report a retrospective crossing statistic at five arbitrary benefit thresholds:

$$H_b = \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{\mathbf{1}\{Q_{I,i} \leq b\} \neq \mathbf{1}\{Q_{F,i} \leq b\}\}. \quad (16)$$

The descriptive classification treats equality as at-or-below. It is a statistic of two estimates, not a measured developer policy. The model’s strict action convention therefore makes no empirical decision claim here.

Table 5: Retrospective crossings at illustrative thresholds, $n = 162$. “Up” means at-or-below to above; “down” means above to at-or-below. Thresholds are real-2020 dollars/kW and are not observed project valuations.

Threshold b	Up	Down	Total crossings	Share H_b
25	10	12	22	13.6%
50	3	21	24	14.8%
100	7	21	28	17.3%
200	4	10	14	8.6%
400	0	4	4	2.5%

At $b = 100$, 28 of 162 pairs switch sides, including seven upward and 21 downward crossings. A crossing frequency alone contains no payoff magnitude, acquisition fee, or project-specific benefit distribution, and thus cannot be converted into $V(Y)$ or G . It illustrates why the same engineering revisions can matter differently at different economic thresholds.

6.4 Sensitivity and interpretation limits

Restricting the primary sample to source version 1 at both stages leaves 81 pairs. The median signed change is -6.79 dollars/kW, the median absolute change is 14.06, and 13 of 81 pairs cross $b = 100$ (16.0%). The median interval is 19 months. This subset addresses one aspect of version replacement; a version label still does not establish the exact date a developer received or acted on information.

Constant-dollar changes include the effect of annual CPI normalization. On the same 162 pairs, comparing nominal dollars/kW at each study’s own price year yields 66 increases, 93 decreases, and three ties, compared with 59, 102, and one in constant dollars. The nominal median signed and absolute changes are -4.11 and 13.82 dollars/kW, respectively. These nominal differences mix price years and are provided only as a sensitivity check; a decline in a real estimate need not mean the nominal engineering estimate declined.

Before requiring complete, reconciled components, the broader sample contains 496 pairs. Its median signed change is approximately -0.48 dollars/kW, compared with -6.50 in the primary sample. This difference reinforces the dependence of descriptive summaries on sample construction. We foreground the consistent-component sample because its cost definitions are clearer, not because its revisions represent the population. The full filter flow and broader summaries are retained in the analysis outputs.

Five limitations prevent interpreting these revisions as the causal value of information:

1. **Selection.** Both stages are observed only for projects that survived long enough to obtain a facilities estimate. Corresponding downstream outcomes for earlier withdrawals are unavailable.
2. **Timing.** The latest version within a stage may differ from the version available at the historical decision date.
3. **Changing targets.** Unchanged requested MW does not fix location, network topology, allocation group, or engineering scope.
4. **Unobserved realized cost.** Both observations are estimates; neither is established as realized expenditure.

Table 6: Evidence needed to evaluate a specified optional information product.

Evidence	Evaluation role
Timestamped document versions and actual availability	Reconstruct the baseline and prevent hindsight leakage.
Stable target definition, project scope, and cost allocation	Distinguish information updates from changes in the underlying target.
Action deadline, available alternatives, and project benefit	Establish whether a signal can change a remaining decision and locate its threshold.
Acquisition price, time burden, and contingent expenses	Convert gross gain into the appropriate net payoff.
Held-out scores and consistently defined outcomes	Assess calibration, policy gain, and uncertainty out of sample.
Continued tracking of withdrawals and censored outcomes	Expose selection and missing counterfactuals instead of coding missing costs as zero.

5. **Missing decision context.** The data omit project benefits, optional-product fees, action deadlines, and the decision-time information set.

Consequently, $Q_F - Q_I$ is a study-estimate revision, not an identified forecast error against the model’s C . The empirical exercise does not estimate ε_{cal} , ε_Y , or an information-acquisition treatment effect. It also does not establish current PJM revision behavior under subsequent process changes. What it supplies is a reproducible description of a selected historical comparison and a concrete account of the evidence needed to go further.

7 Implications for prospective evaluation

The model suggests evaluating optional information around a specified decision, rather than starting with an aggregate accuracy target. A prospective study should fix the target cost, the permissible action, the benefit threshold or benefit distribution, the baseline information set, and the time at which a signal would have been available. A later estimate is a defensible prediction target only if it is explicitly labeled as such; it should not silently stand in for realized cost.

Table 6 connects missing evidence to the quantity it supports. These are requirements for future validation, not a claim that the present paper has collected them. Workflow interviews can test whether the binary approximation is meaningful and identify the relevant records, but interviews alone cannot identify a causal return to information.

For predictive validation, the training and evaluation split must respect information availability and project identity. Alternative versions of the same request should not leak across a supposed prospective split. Calibration and decision loss should be reported across relevant thresholds and project groups, with uncertainty and sensitivity to missing outcomes. If the intended claim is that access to the product changes behavior or payoff, a credible intervention design is needed; retrospective predictive performance alone is insufficient.

Our framework also distinguishes stakeholder objectives. The project’s private information value need not equal a grid operator’s value or the social value of a queue reform. A private decision can alter other projects’ allocations, administrative workload, or network investment. Those externalities are omitted from (1). An information service could also change engineering design directly, in which

case its benefit includes an intervention effect in addition to learning about a fixed cost.

Questions for developer and operator interviews. The framework suggests asking which upcoming decision is difficult, what information could still change it, which cost components are uncertain, what happens to deposits if the project stops, and which records could establish that an estimate was available and actionable. These questions connect the model to actual workflows without presuming demand for a particular service.

8 Conclusion

Additional information about interconnection cost has private value when it improves a decision that remains available. In a risk-neutral binary benchmark, that value is governed by posterior means crossing a project-specific economic threshold. Marginal cost uncertainty alone cannot identify it, and prediction accuracy need not rank signals in the same order as decision value.

For estimated scores, distinguishing implied value from attained gain exposes a separate requirement: calibration. Under mean matching, one-half of the mean absolute conditional-calibration error bounds their discrepancy, and the factor is sharp. The result provides a transparent acquisition screen when its population inputs are defensible, while highlighting what a retrospective training statistic cannot certify.

The public-data audit makes the empirical boundary equally explicit. A latest-estimate snapshot and a latest-version-per-stage panel answer different questions. Our conservative PJM sample supports descriptive study-estimate comparisons and threshold-crossing counts, with substantial selection and timing limitations. Establishing the value of an optional information product requires a prospective decision definition and evidence connecting the signal, available action, acquisition costs, and a consistently defined outcome.

Data, code, and research transparency

The empirical inputs are public third-party releases cited below. No interviews or nonpublic developer records are used in the reported analysis. Source fingerprints, extraction rules, and exact finite-state calculations are documented in Appendix C. The reproducibility supplement accompanying the manuscript contains `information_value.py`, `audit_umd_pjm.py`, and `audit_pjm_data.py`, dependency versions, input fingerprints, and aggregate outputs. Original data are obtained from the cited providers; third-party raw files are not redistributed in the supplement.

OpenAI Codex was used extensively for literature discovery, derivation checks, code generation, and manuscript drafting, including separate AI-assisted review passes. The finite-state examples were evaluated with exact rational arithmetic, and the descriptive calculations were reproduced from the cited public inputs. AI-assisted review is not external peer review. The named authors are responsible for the manuscript’s claims, calculations, sources, and text.

Funding. This research received no external funding.

Competing interests. The authors declare no competing interests relevant to this work.

A Proofs

A.1 Threshold representation

Proof of Proposition 3.1. Use $x_+ = (x + |x|)/2$ and $\mathbb{E}M = \mu$ to cancel the linear terms in (3). For the first case, apply $(b - M)_+ = b - M + (M - b)_+$ and subtract $b - \mu$. In the second case the baseline is zero. $\square$

A.2 Bounds and refinement

Proof of Proposition 3.2. The function $c \mapsto (b - c)_+$ is convex. Jensen's inequality gives the lower bound, while conditional Jensen gives $(b - M)_+ \leq \mathbb{E}[(b - C)_+ \mid Y]$ and hence the upper bound. Under the Markov condition, $\mathbb{E}[\mathbb{E}(C \mid Y_2) \mid Y_1] = \mathbb{E}(C \mid Y_1)$; another application of conditional Jensen proves refinement monotonicity. Finally, Cauchy–Schwarz gives $\mathbb{E}|b - M| \leq \sqrt{\mathbb{E}(b - M)^2} = \sqrt{(b - \mu)^2 + v}$, which substituted into (4) proves (6). $\square$

A.3 Identification from a marginal distribution

Proof of Proposition 3.3. Proposition 3.2 excludes values outside the interval. For any $r \in [0, 1]$, let R be Bernoulli with probability r , independent of C , and reveal $Y = (R, RC)$. Including R distinguishes a genuine zero-cost observation from a null report. Conditional on $R = 1$, the posterior mean is C ; conditional on $R = 0$, it is μ . Therefore

$$\mathbb{E}[(b - M)_+] = r\mathbb{E}[(b - C)_+] + (1 - r)(b - \mu)_+,$$

so $V(Y) = rV_{\text{PI}}$. Varying r attains the full interval, including the degenerate case $V_{\text{PI}} = 0$. $\square$

A.4 Attained-policy certificate

Proof of Proposition 4.1. Since $(b - Z)_+ = (b - Z)\mathbf{1}\{Z < b\}$ and the action is measurable with respect to Z , subtracting (8) from (7) proves (9). Put $e = g(Z) - Z$, $e_- = (-e)_+$, and $A = \{Z < b\}$. Pointwise,

$$-e_- \leq e\mathbf{1}_A \leq e_+.$$

Moreover, $\mathbb{E}e = \delta$, so $\mathbb{E}e_+ = (\varepsilon_{\text{cal}} + \delta)/2$ and $\mathbb{E}e_- = (\varepsilon_{\text{cal}} - \delta)/2$. Taking expectations and rearranging proves (10); setting $\delta = 0$ gives (11). Section 5.2 gives an example attaining its lower endpoint with positive error, establishing sharpness. $\square$

A.5 Calibration and regret

Proof of Corollary 4.2. Because Z is measurable with respect to Y , $g(Z) = \mathbb{E}[M \mid Z]$. Conditional Jensen gives $\varepsilon_{\text{cal}} \leq \mathbb{E}|M - Z|$; applying the same argument directly to C gives $\varepsilon_{\text{cal}} \leq \mathbb{E}|C - Z|$. A wrong action loses $|M - b|$ relative to the oracle using Y . If the actions differ and $M \neq b$, the threshold lies between M and Z , so $|M - b| \leq |M - Z|$. At $M = b$ the loss is zero. This proves (13). Conditional calibration makes the integrand in (9) zero. A constant score $Z = \mu$ is calibrated but can discard all useful variation in Y , proving the final distinction. $\square$

B Delay that is incurred only if the project proceeds

The main acquisition charge h is unconditional. If instead an optional investigation costs k unconditionally and reduces surplus by a deterministic $d \geq 0$ only when the project proceeds, its net value under correct posterior use is

$$N_d = \mathbb{E}[(b - d - M)_+] - (b - \mu)_+ - k. \quad (17)$$

This is generally different from $V(Y) - k - d$. The threshold after acquisition is $b - d$. Define

$$D_d = \mathbb{E}[(b - M)_+ - (b - d - M)_+].$$

Then $N_d = V(Y) - D_d - k$, and the pointwise integral representation of the positive-part function gives

$$D_d = \int_{b-d}^b \mathbb{P}(M < t) dt, \quad d\mathbb{P}(M \leq b - d) \leq D_d \leq d\mathbb{P}(M < b). \quad (18)$$

Tonelli's theorem justifies exchanging expectation and integration because the integrand is non-negative. The endpoint conventions handle atoms; changing a value at one integration endpoint has no effect on the integral. In particular, $0 \leq D_d \leq d$. This extension still omits stochastic deadlines, queue-position changes, financing constraints, and strategic waiting. Such features require a dynamic model rather than an unconditional subtraction from information value.

C Reproduction specification

C.1 Finite-state calculation

For costs c_i , probabilities $p_i > 0$ summing to one, and a deterministic partition with labels ℓ_i , the signal-group mass and posterior mean are

$$w_j = \sum_{i:\ell_i=j} p_i, \quad m_j = \frac{\sum_{i:\ell_i=j} p_i c_i}{w_j}.$$

Compute $\mu = \sum_i p_i c_i$, $B = (b - \mu)_+$, and

$$V = \sum_j w_j (b - m_j)_+ - B, \quad \text{MSE} = \sum_i p_i (c_i - m_{\ell_i})^2.$$

Given estimated group scores z_j , compute

$$J = \sum_j w_j (b - z_j)_+ - B, \quad G = \sum_j w_j (b - m_j) \mathbf{1}\{z_j < b\} - B.$$

These formulas also explain why MSE for the ideal posterior and MSE for a misspecified score are different quantities. The supplement reports both separately. The reversed score in Section 5.2 has score MSE 5200, although the underlying perfectly revealing signal has posterior MSE zero.

The following complete Python fragment reproduces Table 2 using exact rational arithmetic and no third-party package:

```
from fractions import Fraction as F
c = [F(x) for x in (0, 40, 60, 100)]
p = [F(1, 4)] * 4
b, fee = F(50), F(27, 2)
```

```

mu = sum(q*x for q, x in zip(p, c))
base = max(b-mu, 0)
signals = {"A": [0,1,1,2], "B": [0,0,1,1],
          "perfect": [0,1,2,3]}
for name, lab in signals.items():
    w, m = {}, {}
    for j in sorted(set(lab)):
        ids = [i for i in range(4) if lab[i] == j]
        w[j] = sum(p[i] for i in ids)
        m[j] = sum(p[i]*c[i] for i in ids) / w[j]
    value = sum(w[j]*max(b-m[j], 0) for j in w)-base
    mse = sum(p[i]*(c[i]-m[lab[i]])**2 for i in range(4))
    print(name, float(mse), float(value), float(value-fee))

```

For the broader verification, the three probability vectors are $(1, 1, 1, 1)/4$, $(1, 2, 3, 4)/10$, and $(4, 3, 2, 1)/10$. Enumerate each set partition once using restricted-growth labels: $\ell_1 = 0$ and $\ell_i \leq 1 + \max_{j < i} \ell_j$. For each threshold $0, 1, \dots, 100$, check the threshold identity, perfect-information bounds, and zero oracle regret. The six score vectors are $m + q$ for $q \in \{-50, -10, 0, 10, 50\}$ and the reversed vector of group means. Check the general one-sided certificate using both ε_{cal} and ε_Y , the attained-policy identity, and the regret bound by exact fractions. Scores shared by multiple groups are pooled when calculating $g(Z)$. Repeated numerical checks do not replace or extend the scope of the proofs.

C.2 Public-source extraction and fingerprints

The UMD release is available from its institutional landing page [17]. The audited archive is `pjm-data-and-code-for-posting.zip`. The analysis reads the base file `pjm_main_data_for_public.csv` and joins source CPI factors from `pjm_main_data_for_public_with_additional_vars.csv` on the unique pair of `queue_number` and `study`. Raw costs and dates always come from the base file. Only raw dates in year-month form are parsed; no missing date is filled.

Within each source row, set $K = \max(\text{energy_mw}, \text{capacity_mw})$ over available values. Join pure `Imp` and `Fac` rows by request. Apply the restrictions in Table 3 in that order, using an absolute capacity tolerance of 10^{-8} MW and zero relative tolerance. Source flags must mark both rows as non-updates. Convert costs with (15), where $A = \text{base_cpi_2020}/\text{cpilegns}$. Enforce complete POI and allocated-network values and the stated component-reconciliation tolerance. Summaries are unweighted across retained requests, not capacity-weighted. The two stage medians and all change quantiles are calculated separately; percentages are rounded after computing integer counts.

SHA-256 fingerprints of the exact audited inputs are listed below. They identify the retrieved files, not a certification of their accuracy. Line breaks in the typeset hashes are cosmetic.

UMD archive

6cf9f966808bd5608c67b679ce12f5823e7f7be57905b637134e0de4cfcbcbdbd

UMD base CSV

381b614ec2963125c29b927f941686f4e60573f673185787b6d5368a5006bf17

UMD augmented CSV

403cf8c83ae69708fce4202daa0feef8f073d8e343192056d80148ddcdd6cc86

Berkeley Lab PJM workbook

50a363318b18e36418ac1af4f7afb923998c0863d990a73b59816d014c0af245

The Berkeley Lab audit reads its `data` and `codebook` sheets and verifies unique request IDs and the definition of the latest-study date. It is used here to establish the data structure, not to pool observations with the UMD panel. Third-party raw files are not required to compile this manuscript and are not embedded in its source.

References

- [1] Will Gorman, Julie Mulvaney Kemp, Joseph Rand, Joachim Seel, Ryan Wiser, Nick Manderlink, Fredrich Kahrl, Kevin Porter, Will Cotton. Grid connection barriers to renewable energy deployment in the United States. *Joule*, 9(2):101791, 2025. <https://doi.org/10.1016/j.joule.2024.11.008>.
- [2] Joachim Seel, Joseph Rand, Will Gorman, Dev Millstein, Ryan H. Wiser, Will Cotton, Katherine Fisher, Olivia Kuykendall, Ari Weissfeld, Kevin Porter. Interconnection Cost Analysis in the PJM Territory. Lawrence Berkeley National Laboratory, 2023. <https://emp.lbl.gov/publications/interconnection-cost-analysis-pjm>.
- [3] Joseph Rand, Anna Cheyette, Chris Talley, Steven Zhang, Will Gorman, Ryan H. Wiser, Joachim Seel, Seongeun Jeong, Fredrich Kahrl. Queued Up: 2026 Edition, Characteristics of Power Plants Seeking Transmission Interconnection As of the End of 2025. Lawrence Berkeley National Laboratory, 2026. <https://emp.lbl.gov/publications/queued-2026-edition-characteristics>.
- [4] Sarah Johnston, Yifei Liu, Chenyu Yang. Grid connection costs as a barrier to new generation: Evidence from PJM and implications for transmission policy. *Energy Economics*, 161:109540, 2026. <https://doi.org/10.1016/j.eneco.2026.109540>.
- [5] Sarah Johnston, Yifei Liu, Chenyu Yang. An Empirical Analysis of the Interconnection Queue. National Bureau of Economic Research, Working Paper 31946, 2023. <https://doi.org/10.3386/w31946>.
- [6] Xiaochen Sun. Impacts of Uncertainty in Transmission Interconnection on Energy Transition. Working paper, conference abstract, undated. Undated three-page abstract; working paper listed on the author’s website. Accessed October 2, 2026. https://usaee.org/aws/USAEE/asset_manager/get_file/905148?ver=1.
- [7] David Blackwell. Equivalent Comparisons of Experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953. <https://doi.org/10.1214/aoms/1177729032>.
- [8] Ronald A. Howard. Information Value Theory. *IEEE Transactions on Systems Science and Cybernetics*, 2(1):22–26, 1966. <https://doi.org/10.1109/TSSC.1966.300074>.
- [9] Adam N. Elmachtoub, Paul Grigas. Smart “Predict, then Optimize”. *Management Science*, 68(1):9–26, 2022. <https://doi.org/10.1287/mnsc.2020.3922>. Preprint: [arXiv:1710.08005](https://arxiv.org/abs/1710.08005).
- [10] Heyuan Liu, Paul Grigas. Risk Bounds and Calibration for a Smart Predict-then-Optimize Method. In *Advances in Neural Information Processing Systems*, vol. 34, pp. 22083–22094, 2021. <https://proceedings.neurips.cc/paper/2021/hash/b943325cc7b7422d2871b345bf9b067f-Abstract.html>. Preprint: [arXiv:2108.08887](https://arxiv.org/abs/2108.08887).

- [11] Jean-Daniel Saphores, Eric Gravel, Jean-Thomas Bernard. Regulation and Investment under Uncertainty: An Application to Power Grid Interconnection. *Journal of Regulatory Economics*, 25(2):169–186, 2004. <https://doi.org/10.1023/B:REGE.0000012288.66438.f5>.
- [12] Adriaan Hendrik van der Weijde, Benjamin F. Hobbs. The economics of planning electricity transmission to accommodate renewables: Using two-stage optimisation to evaluate flexibility and the cost of disregarding uncertainty. *Energy Economics*, 34(6):2089–2101, 2012. <https://doi.org/10.1016/j.eneco.2012.02.015>.
- [13] Max Langtry, Ruchi Choudhary. Quantifying the benefit of load uncertainty reduction for the design of district energy systems under grid constraints using the value of information. *Applied Energy*, 400:126549, 2025. <https://doi.org/10.1016/j.apenergy.2025.126549>. Preprint: [arXiv:2412.16105](https://arxiv.org/abs/2412.16105).
- [14] Federal Energy Regulatory Commission. Fact Sheet: Improvements to Generator Interconnection Procedures and Agreements. 2023. Order No. 2023, July 27, 2023. <https://ferc.gov/news-events/news/fact-sheet-improvements-generator-interconnection-procedures-and-agreements>.
- [15] PJM Interconnection. PJM Manual 14H: New Service Requests Cycle Process. Revision 07, 2026. Effective August 19, 2026. Downloaded October 2, 2026. <https://www.pjm.com/-/media/DotCom/documents/manuals/m14h.ashx>.
- [16] Shengjia Zhao, Michael P. Kim, Roshni Sahoo, Tengyu Ma, Stefano Ermon. Calibrating Predictions to Decisions: A Novel Approach to Multi-Class Calibration. In *Advances in Neural Information Processing Systems*, vol. 34, pp. 22313–22324, 2021. <https://proceedings.neurips.cc/paper/2021/hash/bbc92a647199b832ec90d7cf57074e9e-Abstract.html>. Preprint: [arXiv:2107.05719](https://arxiv.org/abs/2107.05719).
- [17] Sarah Johnston, Yifei Liu, Chenyu Yang. Data on Interconnection Queues: PJM data, documentation, and cleaning code. University of Maryland, Department of Economics, public data release. Accessed October 2, 2026; source documentation covers engineering studies available as of August 2024. <https://www.econ.umd.edu/research/data-interconnection-queues>.