Fuzzy Clock

Decision-Specific Timing Information

Chae-Yeon Xon

Yonsei University

September 2026

Abstract

How much timing information does a policy decision require? I study an additive recording clock under a common-mean response restriction. Recorded-response moments and clock moments recover latent-response moments recursively. Within the raw-moment hierarchy, observationally equivalent environments can share every lower-order clock moment and differ at the next response moment. For polynomial timing loss, this difference preserves scheduled-date rankings and changes the minimum loss, allowing the preferred provision mode to reverse. Quadratic loss separates the mean displacement used to choose a date from the variance used to value flexibility. I describe which validation inputs support these queries, how response-dependent recording changes the required moments, and how sampling uncertainty enters the cost comparison. Gas-disruption calculations illustrate the distinctions among an estimated response, a maintained clock, and an executable delivery rule.

JEL Codes: C14, C18, C21, C32, F51, Q41

Keywords: event time, event study, treatment timing, inverse problems, counterfactual identification, partial identification, natural gas

---

Yonsei University. Email: eunixon@yonsei.ac.kr. This draft is circulated for discussion; comments and suggestions are welcome.

1. Introduction

A buyer may know enough about an event date to schedule a delivery while still being uncertain whether flexible provision is worth its cost. A recorded date locates the event imperfectly, and different features of the recording error enter these two decisions. I ask which timing summaries preserve the decision made from a recorded response.[1] The decision determines the required clock information, while full recovery of the latent response path can require richer validation.

Write the event date as G and its recorded date as $W = G + U$. The observation problem has two inputs: the recorded-date mean response m and a disclosed summary $s(p)$ of the displacement law. Under a common-mean restriction, $m = A_p\tau$, where A_p averages translations of the onset-indexed response τ.[2] Identification asks whether the query agrees across all response–clock pairs that give these inputs. A separate causal comparison supplies m when the response is interpreted as a treatment effect.

A two-environment example shows why the query matters. Both environments generate the same recorded response and have mean displacement one. Their optimal onset-referenced delivery date is also one. In one environment the need is concentrated at that date; in the other it is split equally across dates zero and two. Their normalized minimum quadratic losses are zero and one. A flexibility cost between these values produces different provision modes from the same recorded response and mean timing information. The two clock laws remain observationally equivalent under the recorded summary.

Theorem 1 gives a sharp timing-information hierarchy on the stated classes. The triangular recursion gives sufficiency: recorded moments and clock moments through order r recover response moments through order r. The adjacent-order construction gives necessity by preserving the entire recorded response and every clock moment through $r-1$ while changing response moment r. The reflected-clock construction holds the disclosed information fixed as the query changes.

Theorem 2 converts the hierarchy into a minimum-information result for polynomial timing loss. Clock moments through the polynomial degree recover the loss function and its comparison with a fixed flexibility cost; the adjacent lower-order summary can leave the provision mode unidentified. The highest-order response moment enters as a common level across scheduled dates, so lower-order information can still fix their ranking. For quadratic loss, mean displacement locates the date and displacement variance determines the flexibility threshold.

[1] The response is indexed by recorded event time. Its population, support, and causal interpretation are declared inputs to the clock calculation.

[2] The maintained restriction is mean independence of each response coordinate from displacement under the fixed target weights. Full independence is sufficient and stronger.

The information comparison also keeps the execution rule explicit. A rule referenced to a verified onset requires that onset to be available when the rule is carried out. With need mass T_0 and quadratic mismatch coefficient $\gamma > 0$, optimized recorded-origin and verified-origin scheduled losses differ by

$$\Delta_{\text{exec}} = \gamma T_0 \operatorname{Var}(U). \tag{1}$$

This calculation holds the need profile and loss fixed. Validation of the clock law supplies information for evaluating the two technologies; operational access to the onset makes the second technology feasible.[3]

Implementation begins with the same inputs and the population they describe. A finite moment vector determines the quadratic scheduling queries away from zero response mass. Its covariance combines the response and validation designs. If recording delays covary with response size, the relevant correction also uses joint response–timing moments. Input uncertainty then maps into the provision-mode comparison. More general queries use recorded-space bridges, with range and domain conditions stated separately from the moment hierarchy. Partial validation produces sets of decisions. The associated regret calculation is relative to its declared clock region and action class.

European gas disruptions illustrate the distinction in three settings. An exposure design supplies a recorded-date network response, which documentary and route-flow dates reindex. A fixed seven-market residual-need path and a maintained displacement law illustrate successive clock disclosures. A German storage-option contract supplies a price benchmark for a specified flexibility technology. These calculations use distinct populations and uncertainty statements.[4] The first asks how uncertain dating changes an estimated response; the second isolates which clock moments change a scheduling decision; the third compares the resulting loss with a priced flexibility option. Together they show the successive information and action requirements of the timing calculation.

Related literature. Integer-moment recovery under additive error is developed by Hall and Lahiri (2008, Section 2.2). I use that recursion as the sufficiency argument. The additional comparison preserves all lower-order clock moments and the full recorded response while changing the next response moment. Convolution commutativity also underlies the ambiguity constructions of Choudhary and Mitra (2018, 2014). Here the preserved summaries and the selected decision specify which part of that ambiguity matters. Measurement-error identification with repeated measurements, instruments, and auxiliary samples supplies related routes to learning the latent object or error law

[3] A validation sample can reveal the displacement distribution while leaving the realized onset unavailable when an operating decision must be executed.

[4] The exposure estimates, seven-market scheduling calculation, and contract benchmark are three linked exercises with separate populations.

(Schennach 2004, 2007; Hu and Schennach 2008; Chen et al. 2005). Those observation structures also specify how the auxiliary information relates to the population of interest. Chesher and Schluter (2002) show that different welfare measures respond differently to measurement error, using small-error approximations. The timing results here give exact polynomial-loss identities under the additive common-mean restriction and an adjacent-order comparison that holds the full recorded mean path fixed. Linear-functional identification supplies the adjoint-bridge language used for the query implementation (Severini and Tripathi 2006; Santos 2011; Severini and Tripathi 2012; Carrasco et al. 2007).

Duration models study how treatment timing relates to potential outcomes and selection (Abbring and van den Berg 2003). Event-date uncertainty, treatment classification, and staggered-adoption methods study what can be recovered from observed histories (Ball and Torous 1988; Negi and Negi 2025; Callaway and Sant'Anna 2021; Sun and Abraham 2021; Roth et al. 2023; Borusyak et al. 2024). I take the resulting response, its population weights, and its observation domain as explicit inputs to the clock calculation. This follows the distinction between a counterfactual functional and the conditions supporting its causal interpretation (Chernozhukov et al. 2013). Decision rules and regret under partial identification provide the framework for the subsequent action comparison (Manski 2004; Stoye 2009, 2012; Montiel Olea et al. 2026).

The paper moves from the response and clock assumptions to scheduling, moment information, validation, and the gas application. The online appendix collects proofs, operator extensions, and empirical details.

2. Response, observation, and clock assumptions

2.1. Latent responses and counterfactual queries

The response is defined on a fixed population and evaluated with fixed analysis weights. A causal interpretation uses the untreated-outcome comparison below; the subsequent clock calculation uses an additional recording restriction.

Time is discrete and indexed by $t \in \mathbb{Z}$. Unit i has a latent event time $G_i \in \mathbb{Z} \cup \{\infty\}$. The clock stage is defined on units with $G_i < \infty$; units with $G_i = \infty$ can enter the untreated-outcome construction in the causal first stage. For each finite event date g, let $Y_{it}(g)$ denote the potential outcome under an absorbing treatment that begins at g, and let $Y_{it}(\infty)$ denote the untreated potential outcome. For $G_i < \infty$, define the unit-level response at latent horizon h by

$$\tau_{i,h} = Y_{i,G_i+h}(G_i) - Y_{i,G_i+h}(\infty). \tag{2}$$

A cohort-specific dynamic effect is

(3) $$\tau_{g,h} = \mathbb{E}[Y_{i,g+h}(g) - Y_{i,g+h}(\infty) \mid G_i = g].$$

The additive-clock sections use the common mean response $\tau_h = \mathbb{E}[\tau_{i,h} \mid G_i < \infty]$. The operator sections allow a vector or function $\tau \in \mathcal{H}$. The heterogeneous-cohort section returns to $\tau_{g,h}$. A counterfactual query is a functional $Q : \mathcal{X} \to \mathcal{Z}$ of a latent response path. The operator sections use the continuous linear specialization

(4) $$\psi_a(\tau) = \langle a, \tau \rangle_{\mathcal{H}}.$$

In a finite event window, $\tau \in \mathbb{R}^H$ and $\psi_a = a^\top \tau$. Choices of a produce an impact response, an average over a horizon, a cumulative response, a temporal moment, or a weighted policy loss.

2.2. Clocks and translations

On the finite-event population, a clock is a conditional distribution $K(\cdot \mid G_i, X_i)$ for a recorded timing signal W_i, given latent event time G_i and covariates X_i. A clock has point timing on a maintained support when G_i is measurable with respect to (W_i, X_i). A clock is fuzzy on that support when a recorded signal corresponds to at least two latent event times with positive conditional probability. The additive benchmark is

(5) $$W_i = G_i + U_i,$$

where U_i is an integer-valued clock error. Positive U_i records the event after its latent onset. Negative U_i records the event before its latent onset. Let $\mathcal{X}$ be a path space closed under integer translations and define

(6) $$(S_u x)(h) = x(h + u), \qquad u \in \mathbb{Z}.$$

The operators satisfy $S_u S_v = S_{u+v}$ and $S_u^{-1} = S_{-u}$. A latent-origin path x^G and its recorded-origin version satisfy

(7) $$x^W = S_U x^G.$$

For a clock law $p_u = \mathbb{P}(U = u)$, define the average clock operator

(8) $$A_p x = \mathbb{E}_p[S_U x] = \sum_{u \in \mathbb{Z}} p_u S_u x$$

whenever the displayed expectation is finite in the response space.

ASSUMPTION 1 (Common-mean clock condition). *On the fixed finite-event target population,* $\mathbb{E}[\tau_{i,h} \mid U_i = u] = \tau_h$ *for every horizon* h *and every displacement* u *in the support, with expectations using the stated analysis weights.*

Assumption 1 requires mean independence of each response coordinate from recording displacement, using fixed target weights. Independence of the entire response vector and U implies this restriction. Conditional variances and higher-order dependence may still vary with U. The condition also allows a nonzero mean displacement and dependence between U and G. Measurement-error models distinguish mean restrictions of this kind from independence of the full latent object (Schennach 2016; Hu and Schennach 2008); Schennach (2019) gives a separate subindependence characterization for convolution. Here the restriction is stated directly for conditional response means. A conditional version applies within covariate or cohort cells, with their weights retained in aggregation.[5] Cohort mixing preserves these weights, while association between timing and response magnitude requires additional cross-moments.

2.3. A comparison design supplying the causal contrast

Let $\mathcal{E}$ denote the target population with finite event dates. Expectations in the clock stage use this same population and the same aggregation weights at every horizon. The recorded-horizon causal contrast is

$$m_\ell = \mathbb{E}\Big[Y_{i,W_i+\ell}(G_i) - Y_{i,W_i+\ell}(\infty) \mid i \in \mathcal{E}\Big]. \tag{9}$$

This object requires an untreated-outcome design. An observed outcome profile enters the clock stage as m when that design identifies the displayed contrast. The following result is a sufficient conditional-comparison construction using the cohort and calendar-time logic of modern difference-in-differences (Callaway and Sant'Anna 2021; Sun and Abraham 2021; Roth et al. 2023; Borusyak et al. 2024). Its role is to supply the input contrast; estimator construction and treatment-comparison inference follow the chosen first-stage design.

PROPOSITION 1 (Interval-anchored causal contrast). *For each target unit, suppose* $G_i \in [\underline{G}_i, \overline{G}_i]$ *and choose an observed baseline* $r_i < \underline{G}_i$. *Assume consistency and* $Y_{i,r_i}(G_i) = Y_{i,r_i}(\infty)$. *Let* $\mathcal{C}$ *be a comparison population whose observed outcomes equal their untreated potential outcomes at every evaluated calendar pair* (r, t). *For pre-event covariates* X, *define*

$$g_0(x; r, t) = \mathbb{E}[Y_{jt} - Y_{jr} \mid j \in \mathcal{C}, X_j = x]. \tag{10}$$

[5] The conditional construction uses the same cell definition for the response and validation inputs before aggregation.

Suppose target covariate support is contained in comparison support at every evaluated pair, and for every target cell (x, w, r), at $t = w + \ell$,

$$\mathbb{E}[Y_{it}(\infty) - Y_{ir}(\infty) \mid i \in \mathcal{E}, X_i = x, W_i = w, r_i = r] = g_0(x; r, t). \tag{11}$$

With finite displayed expectations and observed outcomes on this support,

$$m_\ell = \mathbb{E}[Y_{i,W_i+\ell} - Y_{i,r_i} - g_0(X_i; r_i, W_i + \ell) \mid i \in \mathcal{E}]. \tag{12}$$

PROOF. Comparison consistency identifies g_0 on target support. In each target cell, observed change is the treatment response at t plus the untreated change from r to t, because baseline consistency gives $Y_{ir} = Y_{ir}(\infty)$. Equation (11) identifies the conditional mean of the latter change. Subtract g_0 and average over the target cells to obtain (12). □

An implementation estimates the comparison change at each required calendar pair, predicts it for target covariates, and averages the resulting target residuals with the fixed analysis weights. The recorded response is this weighted estimand. A regression representation uses it as the clock input when its own weighting and comparison restrictions recover that estimand. Estimated untreated-outcome predictions contribute to the first-stage uncertainty under that design. The comparison-population condition includes treatment status and exposure to spillovers at those dates. Equation (11) states parallel untreated changes within recorded-clock cells; intervals locate an admissible baseline, and the comparison design supplies the untreated change. Anticipatory behavior is accommodated by choosing a baseline at which the stated pre-event consistency condition holds. An announcement and an implementation can also define separate interventions; the intervention represented by G is fixed before selecting its recorded clock.

Under Assumption 1, conditioning the identified contrast on U gives

$$m_\ell = \sum_u p_u \tau_{\ell+u}, \qquad m = A_p \tau. \tag{13}$$

Clock estimation or validation supplies information about p; the comparison design supplies the untreated counterfactual in (9). For a sensitivity analysis, define d_ℓ as the target-population mean of the untreated change from r_i to $W_i + \ell$ minus its comparison prediction g_0, and set $\delta_h(u) = \mathbb{E}[\tau_{i,h} \mid U = u] - \tau_h$. The mean observed design score then equals $(A_p\tau)_\ell + d_\ell + \sum_u p_u \delta_{\ell+u}(u)$. Each deviation has a separate identifying restriction and data requirement.

3. A scheduling decision with an uncertain event date

Consider a buyer choosing a dated delivery or paying for flexible provision. The object to be evaluated is a nonnegative onset-indexed need profile. Before deriving the general information hierarchy, this section makes the distinction between the date and the provision mode explicit.

For this example write $T_r = \sum_h h^r \tau_h$, $M_r = \sum_\ell \ell^r m_\ell$, and $\mu_r = \mathbb{E}[U^r]$. The normalized second central moments are $V_\tau = T_2/T_0 - (T_1/T_0)^2$ and $V_m = M_2/M_0 - (M_1/M_0)^2$. The general transport calculation applies to these quantities.

3.1. The scheduled date and the flexibility threshold

For the quadratic specialization, let $\tau_h \geq 0$ have $T_0 > 0$ and a finite second moment. A scheduled program incurs mismatch loss with coefficient $\gamma > 0$, and flexible provision matches the profile at cost $K \geq 0$:

$$L(a;\tau) = \gamma \sum_h \tau_h (h-a)^2, \qquad L(F;\tau) = K, \qquad \lambda = \frac{K}{\gamma T_0}. \tag{14}$$

The date a is measured from an operational origin available to the execution system. The formulas allow a real-valued date; a constrained date set uses the corresponding projection.

PROPOSITION 2 (Scheduling and flexibility). *Write $\bar{h} = T_1/T_0$. Scheduled loss satisfies*

$$\frac{L(a;\tau)}{\gamma T_0} = V_\tau + (a-\bar{h})^2, \qquad a^* = \bar{h}, \qquad \frac{\min_a L(a;\tau)}{\gamma T_0} = V_\tau. \tag{15}$$

Flexible provision is strictly preferred when $\lambda < V_\tau$, scheduling is strictly preferred when $\lambda > V_\tau$, and both are optimal at equality. Under Assumption 1,

$$a^* = \frac{M_1}{M_0} + \mu_1, \qquad V_\tau = V_m - \mathrm{Var}(U). \tag{16}$$

PROOF OF PROPOSITION 2. Write $h - a = (h-\bar{h}) + (\bar{h}-a)$. Expanding the square and summing against τ_h eliminates the cross term because $\sum_h \tau_h (h-\bar{h}) = 0$, which gives (15). Comparing the scheduled minimum with K gives the mode rule. Corollary 2 gives $\bar{h} = M_1/M_0 + \mu_1$ and $V_\tau = V_m - \mathrm{Var}(U)$. □

Mean displacement determines the verified-origin date and displacement variance determines the flexibility threshold. The example isolates the latter distinction, and a first-order construction supplies the corresponding date witness.

3.2. The same record, different provision modes

Let δ_j denote unit mass at horizon j. Consider the two environments

$$(p, \tau) = \left(\tfrac{1}{2}\delta_0 + \tfrac{1}{2}\delta_2, \delta_1\right), \qquad (p', \tau') = \left(\delta_1, \tfrac{1}{2}\delta_0 + \tfrac{1}{2}\delta_2\right). \tag{17}$$

The recorded path is $(\delta_{-1} + \delta_1)/2$ in both environments. Both clocks have mean one, and both optimal verified-origin offsets equal one. The need profile in the first environment has variance zero; the second has variance one. Each response is nonnegative and supported at or after onset. Within each environment it can be taken as common across units, satisfying Assumption 1.

TABLE 1. A common recorded response and different provision modes

Environment	$\mathbb{E}U$	$\mathrm{Var}(U)$	a_G^*	V_τ	Mode at $\lambda = 1/2$
A	1	1	1	0	Scheduled
B	1	0	1	1	Flexible

Exact quantities from the two finite-support environments in Section 3.2. Total mass is one. Loss is divided by γT_0 and flexibility costs $\lambda = 1/2$. The same reversal holds for every $0 < \lambda < 1$. The table is an algebraic example.

The common mean disclosure settles the scheduled date. The variance disclosure settles which provision mode minimizes cost. The same comparison extends to any raw-moment order.

3.3. Information and feasible execution origins

Execution at the recorded origin uses the response m directly. For nonnegative unit profiles with finite absolute cross-moments,

$$\mathbb{E}\left[\sum_h \tau_{i,h}\{h - (U_i + a)\}^2\right] = \sum_\ell m_\ell (\ell - a)^2. \tag{18}$$

The optimal recorded-origin offset is M_1/M_0. Write L_W^* and L_G^* for minimum scheduled losses under recorded-origin and verified-origin execution. Under Assumption 1,

$$L_W^* = \gamma T_0 V_m, \qquad L_G^* = \gamma T_0 V_\tau, \qquad \Delta_{\mathrm{exec}} = L_W^* - L_G^* = \gamma T_0 \,\mathrm{Var}(U). \tag{19}$$

This comparison holds the need profile and loss specification fixed. An execution-system investment can be evaluated by comparing its cost with Δ_{exec}. Validation information affects the decision through the admissible clock summaries.

The execution-origin comparison concerns two feasible classes of operating rules. Estimating the displacement law supplies information for evaluating those classes. Implementing a rule relative to each unit's verified onset additionally requires that onset to be available at the chosen execution date. A deterministic change of the analyst's event-time label merely changes coordinates; the variance term in (19) arises from heterogeneous displacements under the maintained common-mean condition.

TABLE 2. Decision queries and clock information

Target	Clock information	Recovery equation
Recorded-origin scheduling	Recorded response and execution rule	$a_W^* = M_1/M_0$
Verified-origin scheduling	Mean displacement	$a_G^* = M_1/M_0 + \mu_1$
Verified-origin flexibility	Displacement variance	$V_\tau = V_m - \mathrm{Var}(U)$
Cumulative response	Probability normalization	$T_0 = M_0$
Raw temporal moment T_r	Clock moments through r	Proposition 3
Polynomial timing loss, degree d	Clock moments through d	Theorem 2
Finite linear action ordering	Pairwise common bridges	Proposition 4

$M_r = \sum_\ell \ell^r m_\ell$. The scheduling rows use quadratic mismatch, positive total response mass, and the common-mean clock condition. Finite-window queries also use boundary information.

4. Identification by moment order

A timing summary records features of the clock law that accompany the recorded response in a query calculation.

DEFINITION 1 (Query-identifying timing summary). *Let $\mathcal{P}$ be a class of clock laws, let $\mathcal{T}$ be a class of response paths, and let Q be a query. A map $s : \mathcal{P} \to \mathcal{S}$ identifies Q over $\mathcal{P} \times \mathcal{T}$ when, for every $(p, \tau), (p', \tau') \in \mathcal{P} \times \mathcal{T}$,*

$$A_p\tau = A_{p'}\tau' \quad \text{and} \quad s(p) = s(p') \quad \Longrightarrow \quad Q(\tau) = Q(\tau'). \tag{20}$$

4.1. Temporal moments

Assume $\sum_h (1 + |h|^r)|\tau_h| < \infty$ and $\mathbb{E}|U|^r < \infty$. Define

$$T_j(x) = \sum_{h\in\mathbb{Z}} h^j x_h, \qquad T_j = T_j(\tau), \qquad M_j = T_j(m), \qquad \mu_j = \mathbb{E}[U^j], \tag{21}$$

with $T_0 = \sum_h \tau_h$, $M_0 = \sum_\ell m_\ell$, and $\mu_0 = 1$. Collect moments through order r as $\mathbf{T}_r = (T_0, \dots, T_r)^\top$ and $\mathbf{M}_r = (M_0, \dots, M_r)^\top$. For $j, k \in \{0, \dots, r\}$, define the Pascal-action matrix

$$[\Pi_r(u)]_{jk} = \begin{cases} \binom{j}{k}(-u)^{j-k}, & k \le j, \\ 0, & k > j. \end{cases} \tag{22}$$

The shift identity is

$$\mathbf{T}_r(S_u x) = \Pi_r(u)\mathbf{T}_r(x). \tag{23}$$

The following recursion is the additive measurement-error moment calculation under the sign convention $W = G + U$. Hall and Lahiri (2008, Section 2.2) gives the integer-moment recovery formula in a deconvolution model. Here the recursion establishes the sufficiency side of the timing-summary comparison.

PROPOSITION 3 (Moment transport recursion). *For every integer $r \geq 0$,*

$$\mathbf{M}_r = \mathbb{E}[\Pi_r(U)]\mathbf{T}_r. \tag{24}$$

Its r-th component is

$$M_r = \sum_{k=0}^{r} \binom{r}{k} (-1)^k \mu_k T_{r-k}. \tag{25}$$

Hence m and $s_r(p) = (\mu_1, \ldots, \mu_r)$ determine $T_0, \ldots, T_r$ recursively; set $s_0(p)$ equal to the empty tuple.

The moment-interchange argument justifies the recursion under the stated integrability condition.

4.2. The information retained by adjacent summaries

THEOREM 1 (Sharp timing-information hierarchy). *Fix $r \geq 1$. Suppose the clock class contains two bounded-support laws p and p' with*

$$\mu_j(p) = \mu_j(p') \quad \text{for } j = 1, \ldots, r-1, \qquad \mu_r(p) \neq \mu_r(p'), \tag{26}$$

and suppose the response class contains every nonnegative finite-support path with total mass one. Then there exist response paths τ, τ' in the class such that

$$A_p\tau = A_{p'}\tau', \qquad s_{r-1}(p) = s_{r-1}(p'), \qquad T_r(\tau) \neq T_r(\tau'). \tag{27}$$

DETAILED PROOF OF THEOREM 1. Let p and p' be the two clock laws in the theorem. Define reflected laws $\check{p}(j) = p(-j)$ and $\check{p}'(j) = p'(-j)$. Bounded support permits an integer c such that $\check{p}(h-c)$ and $\check{p}'(h-c)$ have support in the nonnegative integers. Define environment one by clock p and response $\tau_h = \check{p}'(h-c)$. Define environment two by clock p' and response $\tau'_h = \check{p}(h-c)$. Each response is nonnegative, has finite support, and has total mass one. The recorded responses satisfy

$$m_\ell = (\check{p} * \check{p}')(\ell - c), \tag{28}$$

$$m'_\ell = (\check{p}' * \check{p})(\ell - c). \tag{29}$$

Commutativity gives $m = m'$. Equation (26) gives $s_{r-1}(p) = s_{r-1}(p')$. If $U \sim p$ and $U' \sim p'$, the response moments are

$$T_r(\tau) = \mathbb{E}[(c - U')^r], \tag{30}$$

$$T_r(\tau') = \mathbb{E}[(c - U)^r]. \tag{31}$$

Expanding each moment gives

$$T_r(\tau) - T_r(\tau') = \sum_{k=0}^{r} \binom{r}{k} c^{r-k} (-1)^k \{\mu_k(p') - \mu_k(p)\} \tag{32}$$

$$= (-1)^r \{\mu_r(p') - \mu_r(p)\}. \tag{33}$$

The final quantity is nonzero by (26). □

Theorem 1 compares adjacent summaries s_{r-1} and s_r on the displayed clock and response classes. Its scope is the raw-moment hierarchy; a different query or a restricted response class has its own information requirement.[6] Choudhary and Mitra (2018) characterize rank-two null-space ambiguity for blind convolution; the associated working paper develops almost-everywhere pair ambiguity (Choudhary and Mitra 2014, Section IV). Those comparisons allow the clock's lower moments to change. Here the construction enforces their equality through $r - 1$ and changes the selected r-th response moment. Commutativity is the shared algebraic ingredient; the preserved summaries specify the information comparison. Corollary 2 gives summaries for the mean and variance queries.

Role of the conditions. The transport step uses absolute moments to interchange the response and clock sums. The necessity construction uses bounded clock support and a response class rich enough to contain the reflected, shifted laws. The common shift places both responses after onset while preserving the information comparison. Thus the theorem also has witnesses with zero response at negative latent horizons. Restricted shape classes and summaries chosen outside the raw-moment hierarchy require their own comparison. The necessity conclusion is a statement over the maintained class: it supplies two observationally equivalent environments. A particular record or a fixed flexibility

[6] The theorem characterizes the timing information needed for the declared query family. The full clock distribution lies outside the declared target. For broader measurement-error and auxiliary-data identification results, see Schennach (2004, 2007); Chen et al. (2005).

cost can settle a decision with fewer moments. The loss function below determines which recovered moments enter a particular decision.

4.3. Location, dispersion, and the observation window

COROLLARY 1 (Clock conservation). *If* $\tau \in \ell^1(\mathbb{Z})$*, then*

$$\sum_{\ell \in \mathbb{Z}} m_\ell = \sum_{h \in \mathbb{Z}} \tau_h. \tag{34}$$

COROLLARY 2 (Temporal centroid and dispersion). *The first two causal moments satisfy*

$$T_1 = M_1 + \mu_1 M_0, \tag{35}$$

$$T_2 = M_2 + 2\mu_1 T_1 - \mu_2 M_0. \tag{36}$$

If $\tau_h \geq 0$ *and* $T_0 > 0$*, define the response-mass distribution* $q_h = \tau_h / T_0$ *and the recorded-mass distribution* m_ℓ / M_0*. Their variances* V_τ *and* V_m *satisfy*

$$\frac{T_1}{T_0} = \frac{M_1}{M_0} + \mu_1, \qquad V_\tau = V_m - \mathrm{Var}(U), \qquad V_\tau = \frac{T_2}{T_0} - \left(\frac{T_1}{T_0}\right)^2. \tag{37}$$

The clock summaries are μ_1 *for mean location and* $\mathrm{Var}(U)$ *for temporal variance.*

Full-axis moment statements require the displayed absolute-moment conditions. A persistent level response is handled through a specified finite-window target or an integrable response transformation. For a finite window, the recorded total weights each latent horizon by the probability that its translated coordinate remains inside the window.[7] The observation domain and its boundary information therefore belong to the query specification.

4.4. Response-dependent recording

The common-mean condition links marginal clock information to the target. The following identity shows what changes when response size and displacement are associated. This is a query-specific version of the distinction between classical and nonclassical measurement error emphasized by Hu and Schennach (2008). Let $m_\ell = \mathbb{E}[\tau_{i,\ell+U_i}]$, $T_{ji} = \sum_h h^j \tau_{i,h}$, and $M_j = \sum_\ell \ell^j m_\ell$, using the fixed target population and weights. For integer $r \geq 0$, assume $\mathbb{E}[\sum_h (1 + |h| + |U_i|)^r |\tau_{i,h}|] < \infty$.

[7] Window truncation is part of the observation operator. Restoring response mass translated beyond an unobserved boundary requires boundary information in addition to clock moments.

LEMMA 1 (Query cross-moment identity). *The recorded moment satisfies*

$$M_r = \sum_{k=0}^{r} \binom{r}{k} (-1)^k \mathbb{E}[U_i^k T_{r-k,i}]. \tag{38}$$

PROOF. For each unit, reindex $\sum_\ell \ell^r \tau_{i,\ell+U_i} = \sum_h (h-U_i)^r \tau_{i,h}$. Expand the power and take expectations. The absolute cross-moment condition justifies both exchanges. At $r = 0$ this gives $M_0 = \mathbb{E}[T_{0i}]$. At $r = 1$, write $\mathbb{E}[U_i T_{0i}] = \mu_1 T_0 + \mathrm{Cov}(U_i, T_{0i})$, yielding (39). At $r = 2$, write the two cross-moments as products plus covariances and rearrange to obtain (40). □

The marginal clock moments and the joint response–timing moments coincide in the products used by Proposition 3 under Assumption 1. For an individual query, the displayed cross-moment equation states the corresponding restriction directly. A two-cell calculation illustrates the distinction.

The first two queries. Let $T_{ji} = \sum_h h^j \tau_{i,h}$, $T_j = \mathbb{E}[T_{ji}]$, and $M_j = \sum_\ell \ell^j m_\ell$ on the same target population. The first-order pathwise identity can be averaged whenever $\mathbb{E}[\sum_h (1 + |h| + |U_i|)|\tau_{i,h}|] < \infty$, giving

$$M_0 = T_0, \qquad T_1 = M_1 + \mu_1 M_0 + \mathrm{Cov}(U_i, T_{0i}). \tag{39}$$

The total is invariant to the association between displacement and response. For the first moment, mean-displacement correction uses the restriction $\mathrm{Cov}(U_i, T_{0i}) = 0$; Assumption 1 supplies this restriction as part of its pathwise conditional-mean requirements. Delays associated with the total size of the response contribute the covariance in (39). Information about that joint association can be included explicitly in the query equation.

$$T_2 = M_2 + 2\mu_1 T_1 - \mu_2 T_0 + 2\,\mathrm{Cov}(U_i, T_{1i}) - \mathrm{Cov}(U_i^2, T_{0i}). \tag{40}$$

Equation (40) uses the corresponding second-order absolute cross-moments. These equations distinguish information about the marginal clock law from information about recording and response jointly. A timing-only validation file measures the clock moments. The covariance corrections additionally involve the unit-level response under its causal or measurement model.

For a positive total response, the centroid correction is $\mu_1 + \mathrm{Cov}(U_i, T_{0i})/T_0$. In the two-cell example, using the marginal mean alone gives $-1/6$ day, whereas the exact centroid is zero. The extra term is $1/6$ day. Thus validation of the marginal clock and validation of response–clock association answer different parts of the query. Conditional implementations retain each cell's response and clock before pooling with the target weights.

5. Higher-order and partially informed decisions

The quadratic example separates the information needed for a date from the information needed for the cost of flexibility. The moment-order result now extends that distinction to polynomial losses. Finite action menus and partial validation state the corresponding decision problems under other information sets.

5.1. Higher-order timing losses

Let $c(z) = \sum_{j=0}^{d} c_j z^j$ be a nonnegative polynomial on $\mathbb{R}$, with even degree $d \geq 2$ and $c_d > 0$. A scheduled action a belongs to a compact set $\mathcal{A} \subset \mathbb{R}$, and flexible provision costs $K \geq 0$. Define

$$L_c(a;\tau) = \sum_h \tau_h c(h-a), \qquad L_c(F;\tau) = K. \tag{41}$$

THEOREM 2 (Minimum timing information for polynomial decisions). *Suppose $\tau_h \geq 0$, $T_0 > 0$, the moments through order d are finite, and Assumption 1 holds. The recorded response m and clock moments $s_d(p)$ identify $a \mapsto L_c(a;\tau)$, its minimum on $\mathcal{A}$, and the comparison with K. The coarser summary $(m, s_{d-1}(p))$ identifies every scheduled-date loss difference and the scheduled minimizer correspondence. Under the clock and response classes of Theorem 1, suppose two clock laws agree through moment $d-1$ and differ at moment d. There exist two environments with the same recorded response and the same s_{d-1}, and therefore the same scheduled minimizers, for which minimum scheduled losses differ. Every K strictly between these minima yields different scheduled-versus-flexible choices. Thus order $d-1$ is sufficient for timing within the scheduled menu, while order d is necessary and sufficient on the stated classes for the full scheduled-versus-flexible decision family.*

PROOF OF THEOREM 2. For each scheduled date a, expand $c(h-a) = \sum_{j=0}^{d} b_j(a)h^j$, with $b_d(a) = c_d$. Proposition 3 recovers $T_0, \dots, T_d$ from m and $s_d(p)$, which determines $L_c(a;\tau)$ over the action set and its comparison with K. Apply Theorem 1 at order d. The two constructed environments share m and clock moments through $d-1$, so their latent moments through $d-1$ agree by the triangular recursion. Their scheduled losses differ at every a by $c_d\{T_d(\tau) - T_d(\tau')\}$. The minima inherit this difference, and every flexibility cost between the two minima produces different provision modes. □

The same expansion gives a useful distinction within the scheduled menu. Write $c(h-a) = \sum_{j=0}^{d} b_j(a)h^j$ and fix a reference date $a_0 \in \mathcal{A}$. Since $b_d(a) = c_d$ at every date,

$$L_c(a;\tau) - L_c(a_0;\tau) = \sum_{j=0}^{d-1} \{b_j(a) - b_j(a_0)\} T_j. \tag{42}$$

Hence m and $s_{d-1}(p)$ suffice for all scheduled-date comparisons and their minimizing correspondence. This gives a sufficient timing summary for the scheduled menu. The level of minimum scheduled loss also contains $c_d T_d$, which enters the comparison with flexibility. In the theorem's two environments the scheduled minimizers agree, and the provision mode differs for costs between their minimum losses. Different losses or restricted response classes can use still coarser summaries.

5.2. Finite linear action sets

Let a finite action set $\mathcal{A}$ have losses $L_a(\tau) = k_a + \langle q_a, \tau \rangle$. For recorded response m and clock summary s, let $\mathcal{E}(m, s)$ collect every (τ, p) in the maintained response and clock classes satisfying $m = A_p \tau$ and $s(p) = s$. A decision is identified when its optimal-action correspondence agrees across this environment fiber.

PROPOSITION 4 (Decision fiber and common bridges). *The optimal-action correspondence is identified at (m, s) exactly when $\arg\min_{a \in \mathcal{A}} L_a(\tau)$ is the same for every $(\tau, p) \in \mathcal{E}(m, s)$. A sufficient condition is that for every action pair (a, b) there is a recorded-space vector $v_{ab,s}$ satisfying*

$$A_p^* v_{ab,s} = q_a - q_b \qquad \text{for every clock law } p \text{ with } s(p) = s. \tag{43}$$

Under this condition, $L_a(\tau) - L_b(\tau) = k_a - k_b + \langle v_{ab,s}, m \rangle$, so (m, s) determines the action ordering. If two environments in one fiber reverse the sign of a pairwise variable-loss contrast, an interval of relative fixed costs in the associated two-action problem produces different optimal actions.

Proposition 4 applies the linear-functional representation to payoff contrasts. Components shared by every q_a cancel from those contrasts.

5.3. Partial validation and regret

Partial validation supplies a set of admissible profiles. With loss divided by γT_0, deterministic minimax regret is

$$R^*(\lambda; \mathcal{H}) = \inf_{d \in \mathbb{R} \cup \{F\}} \sup_{(\bar{h}, v) \in \mathcal{H}} \left\{ \ell_d(\bar{h}, v) - \min(\lambda, v) \right\}, \qquad \ell_a = v + (a - \bar{h})^2, \qquad \ell_F = \lambda. \tag{44}$$

The application uses the closed-form solution for these moment regions.

The population moment regions are outer relaxations of the full response–clock fiber. The resulting regret calculation is relative to those declared regions and deterministic action classes. A finite construction shows that feasible corrected mean and variance can coexist with a recorded profile that rules out the candidate clock. The reported regret therefore concerns the stated relaxation. Sampling uncertainty enters through the estimation design and uses a specified coverage target (Imbens and Manski 2004; Romano and Shaikh 2010).[8]

6. Estimation and validation inputs

The population results specify which information enters the decision. An implementation also states how that information is obtained. The causal sample estimates the recorded response using its chosen comparison design. A validation sample, a maintained law, or a set of admissible displacements supplies the timing input. These three timing inputs lead to different uncertainty calculations.

6.1. The population represented by validation

The timing input refers to the displacement $U = W - G$ in the same finite-event population used for m. Linked dates for that population provide a direct validation design. With a differently selected validation sample, the required relation between the two populations becomes an additional assumption. Chen et al. (2005) make this issue explicit for measurement-error models with auxiliary data.

For example, let X index finitely many pre-event cells with target shares π_x. If the validation sample covers every target cell and its conditional clock moments $\mu_{j,v}(x) = \mathbb{E}_v[U^j \mid X = x]$ equal the target-cell moments, then $\mu_j = \sum_x \pi_x \mu_{j,v}(x)$. This averaging recovers the marginal timing input under the stated transfer restriction. Assumption 1 licenses using this marginal input with the pooled response. Under cell-specific response restrictions, each cell retains its own response and clock; query recovery precedes aggregation with the target weights. Estimating the shares and conditional moments contributes to their sampling covariance. Named-event date checks supply information for those events; a population moment uses a selection and weighting design.

A maintained clock fixes its moments as modeling inputs. A sensitivity region allows them to vary over declared values. A validation confidence set records sampling uncertainty around population moments. Each can enter the same transport equations, with its own interpretation of the resulting range.

6.2. A finite moment implementation

For a quadratic timing decision, the calculation has three stages. The response design supplies recorded mass, first moment, and second moment for a fixed population and

[8] For treatment choice under sampling uncertainty and minimax regret, see Manski (2004); Stoye (2009, 2012). Clock ambiguity generates the decision set used here; treatment-effect heterogeneity generates the sets in those applications.

horizon window. Validation or a maintained clock supplies mean displacement and its second moment for that population.[9] The transport equations then give the scheduled date and minimum mismatch loss, which are compared with the cost of flexibility. Reporting these inputs separately lets a reader vary timing knowledge while keeping the causal comparison and action menu fixed.

For a fixed moment order, estimate the required recorded moments with the weights and observation domain used to define m. Combine them with the corresponding clock moments through Proposition 3. With $z = (M_0, M_1, M_2, \mu_1, \mu_2)$, the quadratic scheduling pair is the smooth map

$$g(z) = \left(\frac{M_1}{M_0} + \mu_1,\ \frac{M_2}{M_0} - \left(\frac{M_1}{M_0} \right)^2 - \mu_2 + \mu_1^2 \right). \tag{45}$$

Its entries are the verified-origin date and normalized minimum scheduled loss. The scheduling interpretation uses nonnegative response mass and a positive M_0. The finite-window definition and any zero extension are part of the target.

If a joint central limit theorem gives $\sqrt{n}(\widehat{z} - z)$ covariance Σ_z, the delta-method covariance is $Dg(z)\Sigma_z Dg(z)^\top$, locally where M_0 is bounded away from zero. The covariance matrix retains dependence among the recorded-moment estimates and any association with validation estimates. For example, suppose the response estimates use n units and the independent validation estimates use n_v units. Let Σ_M and Σ_μ be their limiting covariance matrices under their own square-root sample-size normalizations. If $n/n_v \to \kappa \in (0, \infty)$, the covariance under the common $\sqrt{n}$ normalization is

$$\Sigma_z = \operatorname{diag}(\Sigma_M, \kappa\Sigma_\mu). \tag{46}$$

A joint design retains its cross-covariance blocks. These are sampling assumptions on the supplied estimators. The estimated-operator calculation and the application-specific response profile use the same covariance accounting.

6.3. From moment uncertainty to the provision mode

Under Assumption 1 and the unrestricted scheduled-date menu in Proposition 2, define the cost advantage of flexibility by

$$\Delta(z; \gamma, K) = \gamma \left\{ M_2 - \frac{M_1^2}{M_0} - M_0(\mu_2 - \mu_1^2) \right\} - K. \tag{47}$$

[9] Clock moments from an external validation sample require a stated transport condition to the response population.

A positive value favors flexibility; a negative value favors the optimally dated scheduled program. This expression is the existing minimum-loss formula in recorded moments. A constrained date menu adds its projection loss. Its gradient propagates the covariance of the recorded inputs, including uncertainty in estimated γ and K.

A confidence interval for Δ lying above zero supports flexibility, and one lying below zero supports scheduling. An interval containing zero retains both modes. This is a statement about a scalar cost contrast. Simultaneous coverage of an entire identified set is a different inferential target (Imbens and Manski 2004; Romano and Shaikh 2010). When the response and validation designs jointly give a confidence set for the moment inputs, one can retain every mode optimal at some feasible input in that set. On the event that the input set contains the true vector, the resulting mode set contains every mode optimal at that vector. This projection implication inherits the input set's coverage; its statistical justification is the supplied first-stage procedure.

The smooth-query calculation applies at a fixed positive response mass. Uniform validity near vanishing mass or a changing rank requires additional control. Small cost contrasts also make a reported mode sensitive to moment uncertainty. Local Gaussian and projection calculations keep the clock operator and its sampling law explicit. When inputs arrive with only marginal standard errors, cross-moment dependence remains part of the uncertainty calculation; Cocci and Plagborg-Møller (2025) provide worst-correlation bounds for calibrated parameters. The smooth contrast also admits an elementary variance bound.

6.4. Linear-query bridges

The linear-functional identification literature provides the following characterization (Severini and Tripathi 2006; Santos 2011; Severini and Tripathi 2012; Carrasco et al. 2007). Let $A : \mathcal{H} \to \mathcal{Y}$ be bounded, let $m = A\tau$, and consider $\langle a, \tau\rangle$. The query is constant on each observational fiber when $a \in \ker(A)^{\perp} = \overline{\mathrm{Range}(A^*)}$. A bounded linear representation in the observed response has the form

$$\langle a, \tau\rangle = \langle b, m\rangle, \qquad A^* b = a, \tag{48}$$

and exists exactly when $a \in \mathrm{Range}(A^*)$. This operator geometry supports the timing calculation; Theorem 1 addresses the additional information question when the clock law itself is described through nested moment summaries.

The spectral specialization yields a generalized least-squares bridge (Zinde-Walsh 2014). When a validation sample estimates the clock, constant-rank pseudoinverse calculus supplies the bridge path for differentiable plug-in inference (Golub and Pereyra 1973). The covariance calculation separates first-order compatibility from local query identification.

Gaussian projection and identified-set constructions then supply two implementations under partial information.

6.5. What a timing exercise can validate

A deterministic candidate date reindexes a realized response, and a set of such dates gives a pathwise sensitivity set. A population clock law describes recording displacement across a specified population. Estimating that law requires linked timing information and a sampling or transfer model for those links. Moment-compatible outer regions and full response–clock identified sets lead to different decision sets.

These distinctions guide the application. Documentary dates are used as event-specific checks. A maintained candidate law supports the population timing-information exercise. The standard errors for the network response concern its fifteen exposure units and stated comparison design. Each reported interval or range is labeled by the uncertainty it represents.

7. Application: European gas disruptions

The gas calculations use three distinct objects. A fifteen-market exposure design estimates a recorded-date network response. A fixed seven-market residual-need path and a maintained displacement law illustrate what successive timing disclosures recover. A German storage option prices a specified callability benchmark. The first exercise has sampling uncertainty; the second conditions on the path and clock law; the third conditions on a coverage rule and cost comparison. Documentary dates supply event-specific timing checks. ENTSOG flow records, Eurostat balances, and the cited regulator and operator records provide the underlying inputs. The information hierarchy is illustrated under the maintained common-mean model. The three exercises retain their own populations and reference dates; the timing law used for the seven-market calculation is a specified input.

7.1. Estimated network response and deterministic timing sensitivity

Let N_{it} be transit-adjusted cross-border and LNG service divided by 2021 baseline service, and let E_i be pre-event Nord Stream exposure. The local design estimates

$$N_{it} = \alpha_i + \delta_t + \beta E_i \mathbf{1}\{t \geq \text{June 14, 2022}\} + g(\text{weather}_{it}) + \varepsilon_{it}. \tag{49}$$

Under exposure-specific parallel counterfactual changes, predetermined exposure and scale, and consistency of the measured network-service outcome, β is the June 14 network-reallocation gradient. Inference uses fifteen exposure units.

TABLE 3. June 14 network-reallocation contrast

Window	Effect	Delete-one 95% interval	Familywise exact p	Lead joint p
21 days	25.5 pp	[11.5, 39.5]	0.0272	0.966
28 days	27.7 pp	[15.1, 40.2]	0.0086	0.966

Effects are percentage points of baseline daily service for a 10-percentage-point increase in pre-event Nord Stream exposure. Delete-one intervals use a market jackknife. Familywise score inference enumerates shared market sign flips across the two displayed windows. Its exact reference law assumes that the joint null score array is invariant to those flips. The weekly lead test covers event weeks −3 and −2, with week −1 omitted.

The 21-day estimate is 25.53 percentage points and the 28-day estimate is 27.68 points. The shared-sign maximum-statistic procedure gives familywise adjusted p-values 0.0272 and 0.0086 for the 21- and 28-day contrasts, respectively. The weekly lead joint p-value is 0.97. The post-event interpretation retains the stated counterfactual-trend and exposure assumptions. The lead test assesses the displayed pre-event contrasts. Its calibration and the post-event trend restriction have separate roles; the exposure design also requires an appropriate model for dependence across markets. Alternative scaling choices retain positive estimates in both windows.

Let $\widehat{\beta}_k$ be the exposure interaction for event week k, relative to week −1. Define the recorded three-week response profile by

$$\widehat{m}^N_\ell = \widehat{\beta}_k, \quad \ell = 7k, \dots, 7k+6, \quad k = 0, 1, 2, \qquad \widehat{m}^N_\ell = 0 \text{ outside days } 0, \dots, 20. \tag{50}$$

Its mass, centroid, and variance are

$$\widehat{M}^N_0 = \sum_\ell \widehat{m}^N_\ell, \qquad \widehat{a}^N_W = \frac{\sum_\ell \ell \widehat{m}^N_\ell}{\widehat{M}^N_0}, \qquad \widehat{V}^N_W = \frac{\sum_\ell (\ell - \widehat{a}^N_W)^2 \widehat{m}^N_\ell}{\widehat{M}^N_0}. \tag{51}$$

Table 4 reports dispersion with each weekly mass placed at its midpoint. The daily step profile in (50) has the same mass and centroid and adds four squared days of within-week variance, giving 29.95 squared days. A timing-loss calculation uses the resolution of its delivery rule, with the weekly coefficients and delta calculation evaluated at that resolution.

Bundesnetzagentur's 2022 monitoring report records Nord Stream 1 at about 60 percent of maximum capacity on June 14 and 40 percent on June 15 (Bundesnetzagentur and Bundeskartellamt 2022). June 14 therefore supplies the regulator reference for the first reduction day. A route-flow screen around that date gives June 13 and June 16 gas-day candidates, corresponding to $U = W - G \in \{1, -2\}$. For each deterministic candidate,

$$a^N_G = \widehat{a}^N_W + U, \qquad G + a^N_G = W + \widehat{a}^N_W, \qquad V^N_G = \widehat{V}^N_W. \tag{52}$$

The regulator reference sets $U = 0$. The route candidates give a finite sensitivity set for physical onset. They shift the origin of the same estimated path and preserve its calendar centroid. The variance reduction in (19) concerns heterogeneous population displacements under a maintained clock law.

TABLE 4. Clock queries from the response profile

Query	Point or clock set	Sampling interval / outer set	Delete-one range
Three-week response mass	515.4 pp-days	[132.2, 898.6]	[251.9, 584.8]
Recorded/regulator centroid	12.19 days	[10.70, 13.68]	[11.71, 13.93]
Physical-candidate centroid	[10.19, 13.19] days	[8.70, 14.68]	[9.71, 14.93]
Weekly-midpoint variance	25.95 day^2	[17.47, 34.42]	[11.97, 27.45]

Mass uses percentage-point days per ten-percentage-point exposure. Centroids use days relative to the displayed origin. The physical-centroid row takes the outer union over the two route candidates. The variance row reports weekly-midpoint dispersion; daily-step dispersion adds four squared days to each displayed variance value. Either resolution is invariant under deterministic translation.

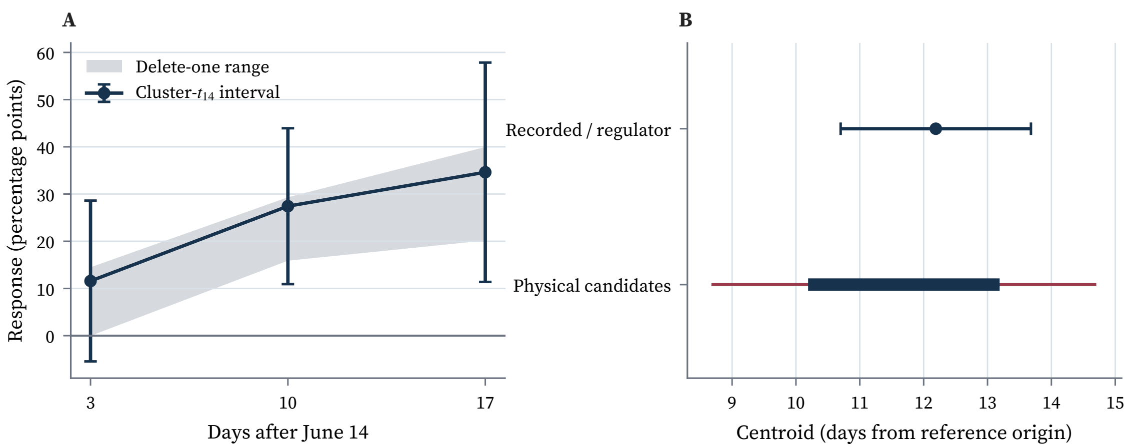

FIGURE 1. Network response and event-time translation

The left panel gives weekly coefficients and sampling intervals. The right panel gives the regulator-referenced centroid and the route-flow sensitivity set. Each origin maps to the same calendar centroid through (52).

The estimated mass is 515.43 percentage-point days and the regulator-referenced centroid is 12.19 days. The route candidates give a physical-origin centroid range of 10.19–13.19 days. The corresponding sampling intervals are reported in Table 4.

7.2. A maintained population-clock calculation

For the policy exercise, let R_{it} be allocated Russian service and let A_{it}^{j} be the signed observed-minus-baseline deviation for external supply, production, or storage withdrawal.

Define the projected residual service need

$$Q^R_{it} = (\widehat{R}^0_{it} - R_{it})_+, \qquad A^\Sigma_{it} = \sum_{j \in J} A^j_{it}, \qquad b_{it} = \frac{(Q^R_{it} - A^\Sigma_{it})_+}{B_{it}}, \tag{53}$$

where B_{it} sums the component baselines. The seven markets with complete 2022–2024 paths define the common target used in the population-clock calculation. The candidate displacement law is applied to that windowed target, yielding $m = A_p\tau$ by construction. Here τ is the fixed residual-need path. This exercise holds the path and candidate law fixed while varying the disclosed clock summary; its percentages describe that information comparison.

The first three disclosure levels recover

$$T_0 = M_0, \qquad T_1 = M_1 + \mu_1 M_0, \qquad T_2 = M_2 + 2\mu_1 T_1 - \mu_2 M_0. \tag{54}$$

TABLE 5. Clock disclosure and deterministic regret for a maintained need path

Clock information	Scheduled date a^*	V_τ/V_m	$\sup_{\lambda\geq 0} R^*/V_m$
Support	[218.9, 491.9]	[82.3, 100.0]%	17.71%
+ mean	399.0	[84.1, 99.2]%	7.55%
+ variance	399.0	92.3%	0.00%

The rows use one recorded target and one candidate displacement law. Dates use days. Variance and regret are percentages of the observed V_m. The first two rows use moment-compatible outer regions conditional on the target and displacement support. Regret is maximized over all nonnegative flexibility costs for each disclosure level; it is measured on the common V_m scale.

Cumulative mass is 172.85 share-days. Support information gives a scheduled-date range of 218.9–491.9 days; mean displacement fixes the date at 399.0. Variance disclosure fixes the flexibility threshold at 92.29 percent of V_m. Maximum normalized regret moves from 17.71 percent under support information to 7.55 percent after mean disclosure and reaches zero after variance disclosure. In the same clock construction, verified-origin execution changes minimum scheduled loss by 7.71 percent of the recorded-origin benchmark. The regret values summarize ambiguity over the declared moment regions. The execution comparison prices a change in the feasible rule, conditional on access to the verified origin. Their common scale permits comparison while their underlying decisions remain explicit.

7.3. A contract-price benchmark for flexibility

German storage options provide a price and execution window for the flexibility comparison. The second 2022 SSBO tender contracted 35.7175 TWh, including 7.1435 TWh callable from October 2022 through January 2023 (Trading Hub Europe 2022a). Weighting the awarded bids by lot counts gives EUR 232.114 million in storage fees and EUR 249.130 million in availability fees; the latter equals EUR 34.88 per callable MWh (Trading Hub Europe 2022b). The award announcement reports approximately EUR 481 million in combined fees (Trading Hub Europe 2022c). Contract terms assign the availability charge to continuous call-off capacity and price commodity delivery separately (Trading Hub Europe 2022d).

ENTSOG assigns Germany a disruption value of EUR 114.1 per MWh for project appraisal (European Network of Transmission System Operators for Gas 2020). Applying the callable quantity uniformly across the 123-day contract window covers 6.31 TWh of the projected residual path. The corresponding gross covered-disruption value is EUR 720.44 million, and combined contract fees imply EUR 76.22 per covered MWh.

TABLE 6. German storage-option benchmark

Quantity	Value
Covered projected shortfall	6.314 TWh
ENTSOG German CoDG	EUR 114.1/MWh
Gross covered-disruption value	EUR 720.4 million
Combined-fee break-even	EUR 76.22/MWh
CoDG headroom after combined fees	EUR 37.88/MWh

The calculation uses the projected residual path, uniform daily callability over October 1, 2022 through January 31, 2023, and ENTSOG's German disruption value. Commodity call charges, storage constraints, realized withdrawals, and incidence enter a welfare calculation through additional inputs.

The covered contract-window profile has centroid 59.79 days from October 1 and variance 1199.83 squared days. Under the quadratic benchmark, a scheduled plan has minimum loss gQ_0V_q. Equating that loss to the EUR 249.13 million availability charge gives

$$g^*_{\text{avail}} = \frac{F_{\text{avail}}}{Q_0 V_q} = \text{EUR } 0.0329 \text{ per MWh per squared day.} \tag{55}$$

This value is the timing-curvature threshold implied by the contract price for the specified need profile. It treats storage-service and commodity terms as common to the fixed-date comparator. In a smooth operating-cost model, the third-derivative remainder in (A99) determines the accuracy of the quadratic approximation over the evaluated mismatch range. The reported threshold conditions on that cost technology.

7.4. Documentary timing checks

Public operator records provide event-date checks on named interruptions. Gasum and Gasgrid document the Imatra interruption, GRTgaz records the Germany–France physical interruption, and Gasgrid records the Balticconnector pressure-loss and isolation episode (Gasum 2022; Gasgrid Finland 2022; GRTgaz 2022; Gasgrid Finland 2023). Within the 2021–2024 Gasgrid English-bulletin frame and the retained ENTSOG routes, the eligible abrupt shutdown events are Imatra and Balticconnector. The reference route-flow rule matches both civil dates. Calendar-date precision maps each event to the adjacent compatible gas-day labels; the screening rule is held fixed in the associated sensitivity calculations.

The documentary matches are event-specific checks. Transferring a displacement law from these events to another population requires an explicit event-selection and transfer model.

8. Conclusion

A timing record is informative relative to a response query and a feasible decision. Under the common-mean additive clock, the moment recursion recovers response moments from the corresponding clock summaries. The adjacent-order construction keeps the full recorded response and all lower-order clock moments fixed while changing the next response moment. Its polynomial-loss counterpart preserves scheduled-date rankings and changes the level that governs the comparison with flexible provision.

Quadratic loss makes the distinction concrete. Mean displacement locates the verified-origin schedule; displacement variance determines its minimum loss. Access to that origin at execution is an additional requirement. A causal comparison, a validation sample, and an operating technology therefore supply different inputs to the same decision.

The gas calculations retain these roles. Sampling intervals describe an exposure-based recorded response. Event-specific dates reindex that response. A maintained population clock illustrates moment disclosure, and a contract price gives a conditional curvature benchmark. The resulting calculations state what each timing summary can resolve and which additional data and operating assumptions support its use.

References

Abbring, Jaap H. and Gerard J. van den Berg, "The Nonparametric Identification of Treatment Effects in Duration Models," *Econometrica*, 2003, *71* (5), 1491–1517.

Ball, Clifford A. and Walter N. Torous, "Investigating Security-Price Performance in the Presence of Event-Date Uncertainty," *Journal of Financial Economics*, 1988, *22* (1), 123–153.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event-Study Designs: Robust and Efficient Estimation," *Review of Economic Studies*, 2024, *91*, 3253–3285.

Bundesnetzagentur and Bundeskartellamt, "Monitoring Report 2022: Key Findings and Summary," Technical Report, Bundesnetzagentur and Bundeskartellamt 2022. Published November 2, 2022.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Carrasco, Marine, Jean-Pierre Florens, and Eric Renault, "Linear Inverse Problems in Structural Econometrics: Estimation Based on Spectral Decomposition and Regularization," in James J. Heckman and Edward E. Leamer, eds., *Handbook of Econometrics*, Vol. 6B, Elsevier, 2007, chapter 77, pp. 5633–5751.

Chen, Xiaohong, Han Hong, and Elie Tamer, "Measurement Error Models with Auxiliary Data," *Review of Economic Studies*, 2005, *72*, 343–366.

Chernozhukov, Victor, Iván Fernández-Val, and Blaise Melly, "Inference on Counterfactual Distributions," *Econometrica*, 2013, *81* (6), 2205–2268.

Chesher, Andrew and Christian Schluter, "Welfare Measurement and Measurement Error," *Review of Economic Studies*, 2002, *69*, 357–378.

Choudhary, Sunav and Urbashi Mitra, "Fundamental Limits of Blind Deconvolution Part I: Ambiguity Kernel," 2014. arXiv:1411.3810v1, November 14; working-paper version associated with Choudhary and Mitra (2018).

___ **and** ___ , "On the Properties of the Rank-Two Null Space of Nonsparse and Canonical-Sparse Blind Deconvolution," *IEEE Transactions on Signal Processing*, 2018, *66* (14), 3696–3709.

Cocci, Matthew D. and Mikkel Plagborg-Møller, "Standard Errors for Calibrated Parameters," *Review of Economic Studies*, 2025, *92*, 2952–2978.

European Network of Transmission System Operators for Gas, "2nd ENTSOG Methodology for Cost-Benefit Analysis of Gas Infrastructure Projects: Annex D," Technical Report, ENTSOG 2020.

Gasgrid Finland, "Gas Supplies through the Imatra Entry Point Have Been Stopped," 2022. 21 May, operational bulletin.

___ , "Suspicion of a Leak in the Balticconnector Gas Pipeline between Finland and Estonia," 2023. 8 October, operational bulletin.

Gasum, "Natural Gas Imports from Russia under Gasum's Supply Contract Has Been Halted," 2022. 21 May, supplier statement.

Golub, Gene H. and Victor Pereyra, "The Differentiation of Pseudo-Inverses and Nonlinear Least Squares Problems Whose Variables Separate," *SIAM Journal on Numerical Analysis*, 1973, *10* (2), 413–432.

GRTgaz, "GRTgaz constate un arrêt du flux physique entre la France et l'Allemagne," 2022. Operational statement, June; retained on the successor NaTran website.

Hall, Peter and Soumendra N. Lahiri, "Estimation of Distributions, Moments and Quantiles in Deconvolution Problems," *Annals of Statistics*, 2008, *36* (5), 2110–2134.

Hu, Yingyao and Susanne M. Schennach, "Instrumental Variable Treatment of Nonclassical Measurement Error Models," *Econometrica*, 2008, *76* (1), 195–216.

Imbens, Guido W. and Charles F. Manski, "Confidence Intervals for Partially Identified Parameters," *Econometrica*, 2004, *72* (6), 1845–1857.

Manski, Charles F., "Statistical Treatment Rules for Heterogeneous Populations," *Econometrica*, 2004, *72* (4), 1221–1246.

Negi, Akanksha and Digvijay Singh Negi, "Difference-in-Differences with a Misclassified Treatment," *Journal of Applied Econometrics*, 2025, *40* (4), 411–423.

Olea, José Luis Montiel, Chen Qiu, and Jörg Stoye, "Decision Theory for Treatment Choice Problems with Partial Identification," *Review of Economic Studies*, 2026. Advance article rdag015, published February 17.

Romano, Joseph P. and Azeem M. Shaikh, "Inference for the Identified Set in Partially Identified Econometric Models," *Econometrica*, 2010, *78* (1), 169–211.

Roth, Jonathan, Pedro H. C. Sant'Anna, Alyssa Bilinski, and John Poe, "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature," *Journal of Econometrics*, 2023, *235* (2), 2218–2244.

Santos, Andres, "Instrumental Variable Methods for Recovering Continuous Linear Functionals," *Journal of Econometrics*, 2011, *161* (2), 129–146.

Schennach, Susanne M., "Estimation of Nonlinear Models with Measurement Error," *Econometrica*, 2004, *72* (1), 33–75.

___, "Instrumental Variable Estimation of Nonlinear Errors-in-Variables Models," *Econometrica*, 2007, *75* (1), 201–239.

___, "Recent Advances in the Measurement Error Literature," *Annual Review of Economics*, 2016, *8*, 341–377.

___, "Convolution without Independence," *Journal of Econometrics*, 2019, *211* (1), 308–318.

Severini, Thomas A. and Gautam Tripathi, "Some Identification Issues in Nonparametric Linear Models with Endogenous Regressors," *Econometric Theory*, 2006, *22* (2), 258–278.

___ **and** ___, "Efficiency Bounds for Estimating Linear Functionals of Nonparametric Regression Models with Endogenous Regressors," *Journal of Econometrics*, 2012, *170* (2), 491–498.

Stoye, Jörg, "Minimax Regret Treatment Choice with Finite Samples," *Journal of Econometrics*, 2009, *151* (1), 70–81.

___, "Minimax Regret Treatment Choice with Covariates or with Limited Validity of Experiments," *Journal of Econometrics*, 2012, *166* (1), 138–156.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Trading Hub Europe, "Ausschreibung zur Gewährleistung der Versorgungssicherheit: Strategic Storage-Based Options, Stufe 1," 2022. Second tender, May 30–June 13; service October 1, 2022–February 1, 2023.

___, "Results SSBO Tender Level 1: Bidding Period 30 May to 13 June 2022," 2022. Award workbook, Results sheet, cells A1:G370.

___, "Second SSBO Tender to Increase Security of Supply Completed," 2022. Press release, June 14.

___, "Terms and Conditions for the Contracting and Use of Stage 1 Strategic Storage-Based Options," 2022. English version dated May 4; sections 4, 6, 8, and 9.

Zinde-Walsh, Victoria, "Measurement Error and Deconvolution in Spaces of Generalized Functions," *Econometric Theory*, 2014, *30* (6), 1207–1246.