Minimal Economy over Worlds

Counterfactuals from Fragmented Evidence

Chae-Yeon Xon

Yonsei University

September 2026

Abstract

Economic theory, measurement rules, and evidence often describe different parts of one economy. Their intersection can contain many complete economies while placing a counterfactual in a short interval or at one value. A minimal economy over admissible worlds is a lowest-cost module collection that reaches a target resolution. At zero tolerance, target certificates compute the counterfactual from retained modules. In linear economies, certificate support solves least-cost acquisition, residual target rank gives minimum repair, covariance weighting selects the least-variable certificate, and certificate residuals trace an attained approximation frontier. Polyhedral modules yield completion and target certificates from one dual system. A joint confidence region provides target coverage after certificate selection. In the Federal Reserve's 2025 stress-test disclosures, revenue and loss modules carry the sign of cumulative pre-tax income for 20 of 22 banks. Minimum capital ratios additionally require quarterly path information, whose completion interval follows from change bounds and checkpoints.

JEL Codes: C10, C13, C14

Keywords: counterfactuals, data combination, identification, value of information

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

1. Introduction

An empirical argument often draws on records that were never designed to form one dataset. A survey measures choices, an administrative system records eligibility, an experiment identifies one behavioral response, and an economic model supplies equilibrium restrictions. Each source restricts a local part of an economic world. Their joint use raises two questions. Do the restrictions admit a common completion? Which distinctions in that completion determine the counterfactual under study?

The second question can be easier than recovery of the full economy. Consider a planner who wants the change in program expenditure after an eligibility rule changes. One record identifies the eligible population, another identifies take-up conditional on an offer, and a third identifies expenditure conditional on receipt. Many joint distributions can agree with those records. The expenditure counterfactual can nevertheless be common across all compatible joint distributions when the relevant conditional means and mixing weights are carried by the separate sources. A linkage that reconstructs every household history then solves a larger problem than the policy question requires.

A completion is a fully specified economy satisfying every retained module. A module may come from theory, a measurement map, an accounting relation, or a dataset.[1] Target ambiguity is the variation of the declared counterfactual across compatible completions. I call a lowest-cost module collection that keeps this ambiguity within a declared tolerance a minimal economy over admissible worlds.[2] It is the least complicated economy that gets the target right. Several collections can attain the same cost, so the object is generally a correspondence. At zero tolerance, a target certificate computes the counterfactual from module values and gives the same answer on every compatible completion.

The linear completion economy makes the idea transparent. Let $\theta \in \mathbb{R}^p$ index a complete economy. Module j reports $A_j\theta = \mu_j$, and a selected collection S reports $A_S\theta = \mu_S$. The counterfactual is $B\theta$. An exact certificate is a matrix Λ_S satisfying

$$B = \Lambda_S A_S. \tag{1}$$

It computes the target as $B\theta = \Lambda_S\mu_S$ while retaining the completion set.[3] Its block support gives a concrete least complicated economy: each used block names a theoretical or empirical distinction that survives the reduction. A failed equation returns a direction in $\ker A_S$ that changes the target while preserving every selected module. Covariance

[1] A module is any declared restriction on the common primitive. Its source may be empirical, institutional, accounting, or theoretical.

[2] The cost can represent collection expense, disclosure burden, storage, or a fixed cardinality budget. The optimization is conditional on the declared module menu and cost schedule.

[3] The row-space condition is the finite-dimensional form of identification of a linear functional from linear moments; the certificate records the corresponding representer.

weighting chooses among multiple certificates, while the residual $B - \Lambda_S A_S$ measures the target variation left in the completion fiber. The same equation connects economic resolution, acquisition, estimation, and approximation.

The first set of results describes the selected economic resolution. Exact least complicated economies are precisely the minimum-cost block supports of target certificates. The rank difference

$$\delta_B(A_S) = \operatorname{rank}\begin{bmatrix} A_S \\ B \end{bmatrix} - \operatorname{rank}(A_S) \tag{2}$$

is the minimum number of unrestricted scalar refinements required to carry the target. With a fixed module menu, evidence can instead have pure bundle value. Every upward-closed family of target-carrying collections can arise in a finite linear economy, and least-cost acquisition contains weighted set cover.[4] Tractable calculations therefore come from economic structure in the module menu: nesting, locality, separability, or another restriction supplied by the application.

The second set of results connects economic resolution to uncertainty. Covariance weighting gives the least-variable exact certificate. A weighted certificate residual gives the sharp worst-case error over an ellipsoidal completion class. A joint confidence region for all module values can then be passed through any data-selected exact certificate. Coverage holds uniformly over the selection rule because the population module vector is covered before the bundle is chosen. Certificate instability appears directly as interval length.

Equality and inequality modules make the local-to-global step explicit. A Farkas certificate can show that the fragments admit no common completion. When they do, a pair of dual target certificates gives the sharp target interval and detects whether its width meets the declared tolerance. Group support in this dual pair characterizes a least-cost target-carrying bundle. For a threshold decision, the interval endpoints also give the minimax regret, so the value of another module is stated in units of the policy problem.

The algebra behind several of these statements is familiar. Identification of linear functionals through null spaces and adjoint ranges appears in inverse problems and nonparametric instrumental variables; minimum-variance certificates are related to minimum distance, GMM, Riesz representers, and c-optimal design (Elfving 1952; Hansen 1982; Chamberlain 1987; Severini and Tripathi 2012; Chernozhukov, Newey, and Singh 2022). Kaji (2021) studies minimal sufficient underlying regular parameters for weakly regular objects. Data-combination research asks what separate samples jointly identify (Ridder and Moffitt 2007; Fan, Sherman, and Shum 2014; Komarova, Nekipelov, and Yakovlev 2018), while Gu, Russell, and Stringham (2026) constructs a finite latent representation

[4] This complexity statement concerns unrestricted finite module menus. Nested or separable menus can admit simpler acquisition rules.

preserving counterfactual identified sets in discrete outcome models. The discrete acquisition problem also meets the literature on optimal data collection and sensor selection (Carneiro, Lee, and Wilhelm 2020; Zhang, Ayoub, and Sundaram 2017). These connections place the certificate equation inside established econometric and statistical theory.

A neighboring literature studies correlated learning from an outcome map indexed by choices or attributes. In the Gaussian attribute-sampling model of Bardhi (2024) and the Brownian and Gaussian formulations surveyed by Bardhi and Callander (2026), a covariance structure transports information from sampled attributes to unsampled ones, and the policy functional determines which samples are valuable. A completion module can encode such a covariance or stochastic-process restriction when the application supports it. The certificate then reports the conclusion carried by that restriction together with the observed modules. With a broader completion class, the same report records the conclusion supported without probabilistic extrapolation. This separation keeps assumptions governing informational spillovers visible in the evidence ledger.

Together, these ingredients give a target-directed translation from local restrictions to an empirical world. The procedure asks whether theory, measurement, and data describe a common economy, measures target variation across their completions, and finds the least costly refinement that reaches the requested resolution. Its report contains a completion witness, a target certificate or ambiguity interval, an optimality record, and an uncertainty statement.

An application examines the public disclosures from the Federal Reserve's 2025 supervisory stress test. Revenue, provision, and loss components supply an accounting certificate for cumulative pre-tax net income. Interpreting each published tenth of a billion dollars as rounded, the component certificate carries the sign of cumulative pre-tax income for 20 of the 22 participating banks. The same cumulative disclosures leave the minimum capital ratio dependent on the unreported quarterly path. The comparison illustrates how a disclosure can carry one conclusion while leaving another to a direct target report or additional path evidence.

2. From local restrictions to empirical worlds

2.1. Evidence modules

Let Θ be a set of complete economies and let $J = \{1, \dots, q\}$ index modules. Module j is a map $M_j : \Theta \to \mathcal{M}_j$ together with an observed or imposed value μ_j. A module can encode a theoretical restriction, a measurement relation, a maintained stochastic law, or empirical evidence. These sources have different interpretations and enter the same ledger through the worlds they retain. For $S \subseteq J$, write $M_S = (M_j)_{j \in S}$ and $\mu_S = (\mu_j)_{j \in S}$. The completion set is

$$\mathcal{C}_S(\mu_S) = \{\theta \in \Theta : M_S(\theta) = \mu_S\}. \tag{3}$$

The economic target is a map $\tau : \Theta \to \mathcal{T}$. The state θ serves as an interface among the modules: each source can restrict a different projection, while shared coordinates state where their claims must agree.

DEFINITION 1 (Carrying). *A collection S carries τ at μ_S if $\mathcal{C}_S(\mu_S)$ is nonempty and τ is constant on $\mathcal{C}_S(\mu_S)$. It carries τ globally if there is a decoder d_S such that $\tau = d_S \circ M_S$ on Θ.*

Pointwise carrying concerns the evidence values at hand. Global carrying provides a rule that can be reused at other admissible module values. The linear benchmark below makes the two conditions coincide whenever a completion fiber is nonempty, because every fiber is a translation of the same null space.

Equip $\mathcal{T}$ with a metric $d_{\mathcal{T}}$. For a nonempty completion set, define target ambiguity by

$$\mathfrak{a}_\tau(S; \mu_S) = \sup_{\theta, \theta' \in \mathcal{C}_S(\mu_S)} d_{\mathcal{T}}\{\tau(\theta), \tau(\theta')\}. \tag{4}$$

Give each module a nonnegative cost c_j and let $c(S) = \sum_{j \in S} c_j$. At tolerance $\varepsilon \geq 0$, define

$$\mathcal{E}^\star_\varepsilon(\tau; c, \mu) = \operatorname*{arg\,min}_{S \subseteq J} \left\{ c(S) : \mathcal{C}_S(\mu_S) \neq \varnothing,\ \mathfrak{a}_\tau(S; \mu_S) \leq \varepsilon \right\}. \tag{5}$$

The case $\varepsilon = 0$ asks for exact carrying. Positive tolerance records the economic resolution requested by the analyst before acquisition begins. The object is a correspondence because several bundles can attain the same cost. Its element is the retained module ledger together with the full set of compatible worlds.

The report records a witness that the retained modules admit a common world, the target set across those worlds, a certificate or failure direction explaining that set, and an optimality record for the selected module cost. Sampling uncertainty and decision loss can then be attached to the same report.

2.2. Linear completion economies

For the benchmark, $\Theta = \mathbb{R}^p$, module j has a known matrix $A_j \in \mathbb{R}^{m_j \times p}$, and $M_j(\theta) = A_j\theta$. Stack selected modules into $A_S \in \mathbb{R}^{m_S \times p}$ and $\mu_S \in \mathbb{R}^{m_S}$. The target is $\tau(\theta) = B\theta$ for $B \in \mathbb{R}^{k \times p}$.

DEFINITION 2 (Target certificate). *A matrix $\Lambda_S \in \mathbb{R}^{k \times m_S}$ is an exact target certificate for (B, A_S) if*

$$B = \Lambda_S A_S. \tag{6}$$

Its module support contains the blocks of columns of Λ_S associated with nonzero module weights.

The certificate acts directly on reported evidence. If (6) holds, then every completion satisfies $B\theta = \Lambda_S\mu_S$. The weights may be nonunique even when the carried target is unique. This distinction becomes relevant when module estimates have different sampling precision.

2.3. Three orders

The framework uses three comparisons that answer different questions. Inclusion orders module collections by economic resolution. Cost orders collections by acquisition, linkage, modeling, or disclosure burden. Precision orders exact certificates by their sampling covariance. A collection can be cheaper and statistically weaker, and a larger collection can permit a lower-variance certificate without changing population identification. Keeping the three orders separate avoids calling an estimator efficient when the underlying module choice remains unsettled.

3. Target certificates

The first result gives equivalent descriptions of target identification in a linear completion economy.

THEOREM 1 (Certificate characterization). *Fix $A \in \mathbb{R}^{m\times p}$, $B \in \mathbb{R}^{k\times p}$, and $\mu \in \mathrm{range}(A)$. The following statements are equivalent.*

1. *$B\theta$ is constant on $\{\theta : A\theta = \mu\}$.*
2. $\ker A \subseteq \ker B$.
3. $\mathrm{row}(B) \subseteq \mathrm{row}(A)$.
4. *There is a matrix $\Lambda \in \mathbb{R}^{k\times m}$ satisfying $B = \Lambda A$.*
5. *There is a linear decoder $d(\mu) = \Lambda\mu$ satisfying $B\theta = d(A\theta)$ for every $\theta \in \mathbb{R}^p$.*

If these conditions fail, there are $\theta_0, \theta_1 \in \mathbb{R}^p$ such that $A\theta_0 = A\theta_1$ and $B\theta_0 \neq B\theta_1$.

The theorem turns a statement about all compatible economies into a finite equation. The failure witness is equally useful. It identifies an economic direction that every selected module leaves unchanged and along which the counterfactual moves. Such a direction describes what the next piece of evidence must distinguish.

COROLLARY 1 (Sharp target set on an affine fiber). *Let θ_0 satisfy $A\theta_0 = \mu$. Then the identified target set is*

$$\{B\theta : A\theta = \mu\} = B\theta_0 + B(\ker A). \tag{7}$$

It is a singleton exactly under the conditions of Theorem 1.

The affine formula clarifies why a normalization or bounded completion class is required for informative partial identification. When $B(\ker A)$ contains a nonzero vector and $\Theta = \mathbb{R}^p$, the target set is unbounded in that direction. The approximation analysis in Section 5.2 introduces an economically chosen metric and radius.

3.1. Relation to adjoint-range conditions

In a Hilbert-space version, an evidence operator $A : \mathcal{H} \to \mathcal{M}$ and scalar target $\tau(\theta) = \langle b, \theta \rangle_{\mathcal{H}}$ lead to the certificate equation $A^* \lambda = b$. Identification requires b to be orthogonal to $\ker A$, equivalently to belong to the closure of $\mathrm{range}(A^*)$. A regular exact representer requires membership in $\mathrm{range}(A^*)$. With a nonclosed range, an identified target can have ill-posed regular estimation (Severini and Tripathi 2012). The finite-dimensional benchmark has closed ranges, so identification and certificate existence coincide.

This relation sets a boundary for the paper. The adjoint equation is an established econometric object. The module structure adds questions about which restrictions enter the adjoint, whether the modules have a common completion, and how the target guides acquisition.

4. Minimal economic resolution

4.1. The target defect

Define the target defect after observing $A\theta$ by

$$\delta_B(A) = \mathrm{rank}\begin{bmatrix} A \\ B \end{bmatrix} - \mathrm{rank}(A). \tag{8}$$

It counts the independent target rows that remain outside the evidence row space.

THEOREM 2 (Minimum repair and attainment). *Let additional scalar linear measurements have rows collected in $C \in \mathbb{R}^{r \times p}$. The smallest r for which $[A^\top, C^\top]^\top \theta$ carries $B\theta$ is $\delta_B(A)$. A matrix C whose rows form a basis for the image of* $\mathrm{row}(B)$ *in the quotient space* $\mathbb{R}^p / \mathrm{row}(A)$ *attains the bound.*

The lower bound and construction refer to unrestricted scalar measurements. Existing instruments, record formats, or legal access constraints replace the unrestricted row choice with a menu of modules. The least-cost bundle then solves

$$\min_{S \subseteq J,\, \Lambda_S} c(S) \quad \text{subject to} \quad B = \Lambda_S A_S. \tag{9}$$

This program states the acquisition problem directly. It also shows why the target determines the relevant notion of data sufficiency.

Partition a certificate $\Lambda = (\Lambda_1, \dots, \Lambda_q)$ according to the row blocks of the full module matrix A_J and write $\mathrm{supp}_J(\Lambda) = \{j : \Lambda_j \neq 0\}$.

THEOREM 3 (Least-cost certificate support). *Suppose the full linear module menu carries $B\theta$. Then*

$$\min_{S:\, S \text{ carries } B\theta} c(S) = \min_{\Lambda:\, B=\Lambda A_J} c\{\mathrm{supp}_J(\Lambda)\}. \tag{10}$$

A collection S belongs to $\mathcal{E}_0^\star$ exactly when it contains the block support of an exact certificate and its cost equals the common value in (10). *If the full menu admits no exact certificate, every subcollection leaves a target-changing completion direction.*

The theorem gives the finite linear least-complication calculation. The certificate chooses which distinctions survive, its support gives the retained economy, and the cost comparison verifies minimality. The certificate verifies target carrying. The optimization record verifies the cost claim.

4.2. Evidence can be complementary

Take $\theta = (\theta_1, \theta_2)^\top$ and target $\tau(\theta) = \theta_1 + \theta_2$. Module 1 reports θ_1 and module 2 reports θ_2. The target defects are

$$\delta_\tau(\varnothing) = 1, \qquad \delta_\tau(\{1\}) = 1, \qquad \delta_\tau(\{2\}) = 1, \qquad \delta_\tau(\{1,2\}) = 0. \tag{11}$$

Each module has zero individual defect reduction, while their bundle identifies the target.

PROPOSITION 1 (Failure of one-step acquisition). *There is a linear completion economy in which every available module has zero individual target-defect reduction and a finite module bundle carries the target. Consequently, a rule that stops whenever all one-module defect reductions equal zero can stop before target identification.*

The construction shows how individual gains can miss bundle value. Guarantees for greedy selection require additional structure on the module menu. Nested modules, orthogonal target components, and submodular decision value are possible structures; each is an assumption to be verified in the application.

4.3. The computational boundary

For a scalar target, define the access family

$$\mathfrak{A}_b = \{S \subseteq J : b^\top \in \mathrm{row}(A_S)\}. \tag{12}$$

It is upward closed: evidence that carries the target continues to carry it after another module is added. Upward closure is the only universal restriction on this family.

THEOREM 4 (Evidence access and acquisition complexity). *The linear completion economy has the following properties.*

1. *For every upward-closed family $\mathfrak{A} \subseteq 2^J$ with $\varnothing \notin \mathfrak{A}$, there is a finite-dimensional scalar-target economy whose target-carrying collections are exactly the elements of $\mathfrak{A}$.*
2. *Finding a least-cost target-carrying collection is NP-hard, even when the target is scalar and every module reports a subset of coordinate measurements.*
3. *The carrying indicator $v(S) = \mathbf{1}\{S \in \mathfrak{A}_b\}$ is neither submodular nor supermodular over the class of linear completion economies.*

The first statement places evidence acquisition in the span-program tradition (Karchmer and Wigderson 1993). It says that linear target identification can exhibit any pattern of bundle complementarity consistent with retaining information. The second statement follows through a weighted set-cover embedding (Karp 1972). Closely related hardness appears in target-state sensor placement. These connections assign a clear role to economic structure: tractable acquisition requires restrictions on which modules can exist, how they overlap, or how their costs are formed. The unrestricted problem admits neither a universal greedy guarantee nor a closed-form ranking of individual sources.

For a finite menu, the certificate equation still gives a direct mixed discrete–continuous formulation. One may enumerate small menus, use mixed-integer optimization for larger ones, or exploit application-specific nesting. Each reported solution should include the selected bundle, a certificate, and a lower-bound or optimality record. The certificate verifies that the chosen evidence carries the target; the optimization record verifies the cost claim.

4.4. What least complication records

An element of $\mathcal{E}^\star_\varepsilon(\tau; c, \mu)$ records a target, a tolerance, an admissible module menu, a cost system, and a completion class. Changing any of these objects can change the minimizing bundle. "Least complicated" therefore refers to a declared economic problem. One target may require a path state, another only a cumulative flow, and a third an equilibrium restriction linking several local domains.

5. Precision and approximation

5.1. The least-variable exact certificate

Suppose the retained modules are estimated jointly and

$$\sqrt{n}(\widehat{\mu} - A\theta_0) \rightsquigarrow N(0, \Omega), \tag{13}$$

where Ω is positive definite. Any exact certificate $\Lambda A = B$ gives the regular linear estimator $\widehat{\tau}_\Lambda = \Lambda\widehat{\mu}$.

THEOREM 5 (Covariance-weighted certificate). *Assume* $\operatorname{row}(B) \subseteq \operatorname{row}(A)$ *and* $\Omega \succ 0$. *Let*

$$\Lambda_* = B(A^\top\Omega^{-1}A)^\dagger A^\top\Omega^{-1}. \tag{14}$$

Then $\Lambda_*A = B$. *For every other exact certificate* Λ,

$$\Lambda\Omega\Lambda^\top - \Lambda_*\Omega\Lambda_*^\top \succeq 0. \tag{15}$$

The minimum asymptotic covariance is

$$V_* = B(A^\top\Omega^{-1}A)^\dagger B^\top. \tag{16}$$

The population target can remain unchanged when an additional module is added, while its sampling variance falls because the larger module collection admits a more precise certificate. Population sufficiency and statistical precision therefore give different reasons to retain evidence.

The formula is the linear minimum-distance solution. Its role here is to complete the module calculus: the same certificate restriction used for identification defines the efficient combination of noisy module estimates. More general semiparametric claims require tangent-space and regularity analysis beyond (13).

5.2. Sharp approximation over bounded completions

Let the target be scalar, $b^\top\theta$, and let $G \succ 0$ set the units of economically plausible completion differences. Around a reference completion θ_0, consider

$$\mathcal{H}(R, G) = \{h \in \mathbb{R}^p : h^\top G h \le R^2\}. \tag{17}$$

An approximate certificate λ predicts a target change by $\lambda^\top Ah$.

THEOREM 6 (Sharp certificate frontier). *For every* $\lambda \in \mathbb{R}^m$,

$$\sup_{h \in \mathcal{H}(R,G)} \left| (b - A^\top \lambda)^\top h \right| = R\sqrt{(b - A^\top\lambda)^\top G^{-1}(b - A^\top \lambda)}. \tag{18}$$

The minimizing approximate certificate is any solution to

$$AG^{-1}A^\top \lambda = AG^{-1}b. \tag{19}$$

The minimum error is

$$R \, \mathrm{dist}_{G^{-1}}\big(b, \mathrm{row}(A)\big). \tag{20}$$

The bound is attained by completion differences proportional to $G^{-1}(b - A^\top\lambda)$ *whenever the residual is nonzero.*

The metric G is part of the economic specification. A metric based on policy loss gives large weight to completion directions that change the selected action. A metric based on a calibration neighborhood gives a local robustness statement. Reporting G and R states the units of the frontier and the completion differences it covers.

Gaussian correlated learning gives a probabilistic version of the same residual calculation. Suppose $\theta \sim N(\bar{\theta}, \Sigma)$ and a noisy module reports $Y = A\theta + \varepsilon$, where $\varepsilon \sim N(0, \Omega)$ is independent and $\Omega \succ 0$. Linear projection gives

$$\mathrm{Var}(b^\top\theta \mid Y) = \min_{\lambda} \left\{ (b - A^\top\lambda)^\top \Sigma (b - A^\top\lambda) + \lambda^\top \Omega \lambda \right\}, \tag{21}$$

with normal equation $(A\Sigma A^\top + \Omega)\lambda = A\Sigma b$. When $\Omega = 0$ and $\Sigma \succ 0$, the posterior variance is zero exactly when an exact certificate exists. Thus a covariance kernel can transport information across unsampled attributes, as in Gaussian attribute learning (Bardhi 2024); the first term in (21) records the target variation that this transport leaves behind. The ellipsoidal frontier uses a worst-case completion class in place of a Gaussian prior. The identity follows by completing the square in the mean-squared prediction error $\mathbb{E}[(b^\top\theta - \lambda^\top Y)^2]$ after centering θ and Y.

5.3. Weak certificates

Exact identification can coexist with unstable estimation in infinite-dimensional problems. As the target loading approaches the boundary of an adjoint range, certificate norms and asymptotic variances can grow. The finite-dimensional quantity in (16) is a first diagnostic. A full theory must give inference uniformly over sequences with changing rank or certificate support. The fixed-rank calculation identifies that remaining inference requirement.

5.4. Inference after certificate selection

There is a simple way to retain coverage when the evidence bundle and certificate are chosen after seeing the data. Let $\mu_0 = (A_j\theta_0)_{j\in J}$ collect the population module values, and let $\mathcal{Q}_n$ be a joint confidence region satisfying

$$\inf_{P\in\mathcal{P}} P\{\mu_0(P) \in \mathcal{Q}_n\} \geq 1 - \alpha - r_n, \qquad r_n \to 0. \tag{22}$$

A selection rule may use the same data to choose a collection $\widehat{S}$ and an exact certificate $\widehat{\Lambda}$ satisfying $B = \widehat{\Lambda}A_{\widehat{S}}$. Define

$$\mathcal{I}_n(\widehat{S}, \widehat{\Lambda}) = \{\widehat{\Lambda}\nu_{\widehat{S}} : \nu \in \mathcal{Q}_n\}. \tag{23}$$

If the rule finds no exact certificate, set $\mathcal{I}_n = \mathbb{R}^k$ and report the target as unresolved.

THEOREM 7 (Selection-robust certificate inference). *Under* (22)*, every measurable selection rule described above satisfies*

$$\inf_{P\in\mathcal{P}} P\{B\theta_0(P) \in \mathcal{I}_n(\widehat{S}, \widehat{\Lambda})\} \geq 1 - \alpha - r_n. \tag{24}$$

The result permits certificate support, rank, and the selected module collection to change with the sample.

The theorem uses simultaneous insurance across module values, in the same spirit as projection and simultaneous post-selection inference. Moment-inequality inversion, a high-dimensional simultaneous band, or a finite-sample concentration region can supply $\mathcal{Q}_n$ (Andrews and Soares 2010; Belloni et al. 2018). The construction may be conservative. Coverage follows from the joint region, while the acquisition rule can choose a sparse or precise certificate.

For a scalar target and the ellipsoid

$$\mathcal{Q}_n = \left\{\nu : (\nu - \widehat{\mu})^\top \Omega^{-1}(\nu - \widehat{\mu}) \leq q_{1-\alpha}^2/n\right\}, \tag{25}$$

an exact certificate λ gives the interval

$$\lambda^\top\widehat{\mu} \pm \frac{q_{1-\alpha}}{\sqrt{n}}\sqrt{\lambda^\top\Omega\lambda}. \tag{26}$$

Minimizing its half-length over $A^\top\lambda = b$ returns the covariance-weighted certificate in Theorem 5. The interval shrinks along a sequence exactly when the chosen certificate standard deviation is $o(\sqrt{n})$. Widening uncertainty therefore exposes a diverging certificate norm.

The same construction accommodates approximate certificates. If $\theta_0 - \bar{\theta} \in \mathcal{H}(R, G)$, then Theorem 6 enlarges (26) by

$$R\sqrt{(b - A^{\top}\lambda)^{\top}G^{-1}(b - A^{\top}\lambda)}. \tag{27}$$

Choosing λ trades sampling length against the declared completion error. This is the certificate form of bias-aware linear inference; optimality claims belong to the approximate-moment and minimax-inference literature (Armstrong and Kolesár 2021).

6. Global completion certificates

The equality model gives affine completion fibers. Many economic restrictions also contain inequalities: revealed preference, shape restrictions, accounting bounds, incentive constraints, and support restrictions. This section gives the global certificate system for a finite polyhedral economy.

Module j contributes equalities $A_j\theta = \mu_j$ and inequalities $G_j\theta \leq g_j$. For a module collection S, stack the corresponding blocks and define

$$\mathcal{C}_S = \{\theta \in \mathbb{R}^p : A_S\theta = \mu_S,\ G_S\theta \leq g_S\}. \tag{28}$$

The scalar target is $\tau(\theta) = b^{\top}\theta$. For any direction $d \in \mathbb{R}^p$, define the dual value

$$U_S(d) = \inf_{y,z}\left\{\mu_S^{\top}y + g_S^{\top}z : A_S^{\top}y + G_S^{\top}z = d,\ z \geq 0\right\}. \tag{29}$$

The infimum is $+\infty$ when the dual constraint has no solution.

THEOREM 8 (Polyhedral completion and target certificates). *The following statements hold.*

1. *The completion set $\mathcal{C}_S$ is empty if and only if there are y and $z \geq 0$ satisfying*

$$A_S^{\top}y + G_S^{\top}z = 0, \qquad \mu_S^{\top}y + g_S^{\top}z < 0. \tag{30}$$

2. *If $\mathcal{C}_S$ is nonempty and the target is bounded on it, its sharp identified interval is*

$$\tau(\mathcal{C}_S) = [-U_S(-b),\ U_S(b)]. \tag{31}$$

3. *If $\mathcal{C}_S$ is nonempty, the collection S carries τ if and only if there are y^+, y^- and $z^+, z^- \geq 0$ satisfying*

$$A_S^{\top}y^+ + G_S^{\top}z^+ = b, \tag{32}$$

$$A_S^\top y^- + G_S^\top z^- = -b, \tag{33}$$

$$\mu_S^\top (y^+ + y^-) + g_S^\top (z^+ + z^-) = 0. \tag{34}$$

When these conditions hold, the common target value is

$$b^\top \theta = \mu_S^\top y^+ + g_S^\top z^+ = -\mu_S^\top y^- - g_S^\top z^- \quad \textit{for every } \theta \in \mathcal{C}_S. \tag{35}$$

The theorem gives three outputs from one dual system. Equation (30) is a weighted contradiction among module values. Equations (32)–(34) certify a common target across all completions. When their combined objective is positive, the same programs return the sharp width of the target interval. These statements are consequences of finite-dimensional separation and linear-programming duality; support functions and convex geometry already organize sharp identification in moment-restriction models (Rockafellar 1970; Beresteanu, Molchanov, and Molinari 2011; Bontemps, Magnac, and Maurin 2012). The module support supplies the acquisition interpretation used here.

6.1. Minimum module support

Group the coordinates of (y, z) by evidence module. The support of a certificate pair contains module j when at least one dual block $(y_j^+, z_j^+, y_j^-, z_j^-)$ is nonzero.

COROLLARY 2 (Least-cost certificate support). *Suppose the full module system has a completion. The least-cost target-carrying collection solves the equivalent problem*

$$\min_{S, y^\pm, z^\pm} c(S) \quad \textit{subject to} \quad (32)\text{–}(34), \tag{36}$$

where the dual variables are supported on S. Thus every minimizing certificate pair constructs an exact least complicated economy, and every exact least complicated economy admits a minimizing certificate pair.

The corollary turns module selection into a group-support problem in the dual. It also exposes complementarity: the two sides of the target can require different modules, and the zero-width condition is achieved by their joint support. Computational guarantees therefore depend on structure in the module menu.

6.2. Two types of failure

Fragmented evidence can fail before the target is evaluated. The certificate report distinguishes

$$\mathcal{C}_S = \varnothing \qquad \text{from} \qquad \mathcal{C}_S \neq \varnothing \text{ and } |\tau(\mathcal{C}_S)| > 1. \tag{37}$$

The notation in (37) refers to the general completion definition in (3). The first expression records a failure to describe one economy. The second records target variation within a coherent set of economies.

In the exact linear benchmark, completion requires $\mu \in \mathrm{range}(A)$. A vector $z \in \ker(A^\top)$ with $z^\top \mu \neq 0$ certifies failure, because every feasible $\mu = A\theta$ must satisfy $z^\top \mu = 0$. When completion holds, Theorem 1 supplies a separate certificate for the target. The two certificates answer different questions and should be reported separately.

PROPOSITION 2 (Equality-model diagnostic). *For $A \in \mathbb{R}^{m\times p}$ and $\mu \in \mathbb{R}^m$, exactly one of the following holds: there is a completion θ with $A\theta = \mu$, or there is $z \in \ker(A^\top)$ with $z^\top \mu \neq 0$. Conditional on completion, exactly one of the following holds: there is a target certificate Λ with $B = \Lambda A$, or there is $h \in \ker A$ with $Bh \neq 0$.*

The proposition yields a compact report. A completion witness gives one economy satisfying the modules. A completion-failure witness gives a weighted contradiction among module values. A target certificate computes the counterfactual. A target-failure witness gives two compatible economies with different counterfactuals.

The polyhedral theorem supplies the global inequality extension. Smooth nonlinear modules require local charts or global convexification, and sampling uncertainty requires confidence regions for module values. Those extensions must preserve the distinction between rejection of a completion model and population target ambiguity.

6.3. Decision value

A module matters economically when it changes a feasible decision or the loss attached to that decision. Let $a \in \mathcal{A}$ be an action and $u(a,\theta)$ its payoff. For a completion set $\mathcal{C}$, define minimax regret

$$\mathcal{R}(\mathcal{C}) = \inf_{a\in\mathcal{A}} \sup_{\theta\in\mathcal{C}} \left\{ \sup_{a'\in\mathcal{A}} u(a',\theta) - u(a,\theta) \right\}. \tag{38}$$

The value of adding a module is the reduction in $\mathcal{R}$. This definition allows a wide target set to have little value when every completion recommends the same action, and a narrow target set to matter when it crosses a policy threshold. It connects acquisition to the decision problem while preserving the separate identification report. Comparison with Blackwell's ordering of experiments and econometric decision theory is essential for the general result (Blackwell 1951, 1953; Manski 2021).

The threshold problem gives a closed-form illustration. Let the scalar target lie in the sharp interval $[L, U]$, let $a \in \{0, 1\}$, and let the payoff be $u(a,\tau) = a(\tau - \kappa)$. Action 1 is adopted when the target exceeds the policy cost κ.

PROPOSITION 3 (Decision value of a target interval). *The minimax regret is zero when* $U \le \kappa$ *or* $L \ge \kappa$. *When* $L < \kappa < U$,

$$\mathcal{R}([L, U]) = \min\{U - \kappa,\ \kappa - L\}. \tag{39}$$

The minimax action is $a = 1$ *exactly when* $\kappa \le (L + U)/2$. *Hence a module can reduce target width while having zero decision value, and a module that moves one endpoint across* κ *can resolve the decision.*

This formula converts the sharp interval in Theorem 8 into a target-specific acquisition payoff. Richer action sets give a finite or convex regret program. The threshold case already records the distinction needed for empirical design: statistical information receives value through the decision boundary it moves.

7. A disclosure certificate for the stress test

Public stress-test disclosure has two related purposes. It informs market participants about bank resilience, and it permits outsiders to assess the supervisory calculation. These purposes have motivated a literature on the benefits and costs of disclosure (Goldstein and Sapra 2014; Goldstein and Leitner 2018). They also enter the Federal Reserve's proposal to publish comprehensive model documentation and invite public comment on material model changes (Federal Reserve Board 2025b). A target certificate asks a narrower operational question: which disclosed values establish a particular public conclusion?

7.1. A cumulative-income certificate

The Federal Reserve's 2025 results file reports bank-level cumulative amounts over the stress horizon (Federal Reserve Board 2025a). For bank i, collect pre-provision net revenue, other revenue, provisions, securities losses, trading and counterparty losses, and other losses in

$$x_i = (R_i, O_i, P_i, S_i, T_i, L_i)^\top. \tag{40}$$

Cumulative pre-tax net income is carried by the accounting certificate

$$\tau_i = c^\top x_i, \qquad c = (1, 1, -1, -1, -1, -1)^\top. \tag{41}$$

This certificate uses the aggregate revenue module. The file also reports net interest income, noninterest income, and noninterest expense, which provide a more detailed certificate for the first coordinate.

The public amounts are displayed to one decimal place in billions of dollars. Treat a displayed amount $\widetilde{x}_{ij}$ as the interval $[\widetilde{x}_{ij} - .05, \widetilde{x}_{ij} + .05]$. Linear programming then gives the sharp reporting interval

$$\tau_i \in [c^\top \widetilde{x}_i - .30, c^\top \widetilde{x}_i + .30]. \tag{42}$$

The six endpoints attain the bounds. The width records published precision.

The public file permits a direct comparison with the separately published pre-tax net-income values. Every reported value is compatible with the component interval. The largest distance between a reported value and the certificate center is 0.1 billion. The component certificate places cumulative pre-tax income strictly on one side of zero for 20 banks. Barclays US and Northern Trust have intervals that cross zero. A direct report of pre-tax income at the same decimal precision resolves the reporting interval further for those banks.

7.2. Why a path minimum needs path evidence

Capital adequacy is assessed using the minimum capital ratio reached over the projection horizon. The public-disclosure rule accordingly calls for beginning, ending, and minimum regulatory capital ratios. The target-date set below allows the reported minimum to be defined over projection-horizon dates while the initial actual ratio remains a checkpoint. Cumulative revenue and loss components leave that minimum dependent on the interior path.

PROPOSITION 4 (Cumulative disclosure and a path minimum). *Let r_0 and r_T be the beginning and ending values of a capital ratio over a horizon with at least one interior date, and let $\mathcal{H} \subseteq \{0, \dots, T\}$ be the dates entering the reported minimum. Suppose $\mathcal{H}$ contains an interior date and the evidence fixes r_0, r_T, and the cumulative change $r_T - r_0$. For every $\ell \leq \min\{r_0, r_T\}$, there is a compatible path whose minimum over $\mathcal{H}$ is ℓ. Hence these cumulative modules leave the reported path minimum unbounded below.*

The proposition separates model disclosure from target verification. Equations and cumulative inputs can reveal how a total is formed. A conclusion about the weakest point of the path requires a direct minimum report, intermediate capital states, or restrictions on period-by-period changes. The Federal Reserve's results file supplies the direct minimum CET1 ratio. Within the present framework, that field is a target module that carries the regulatory minimum; the cumulative accounting modules leave it open.

Period-by-period restrictions turn this open path into a finite completion problem. Let $0 = t_0 < t_1 < \cdots < t_q = T$ index disclosed checkpoints, and suppose every increment

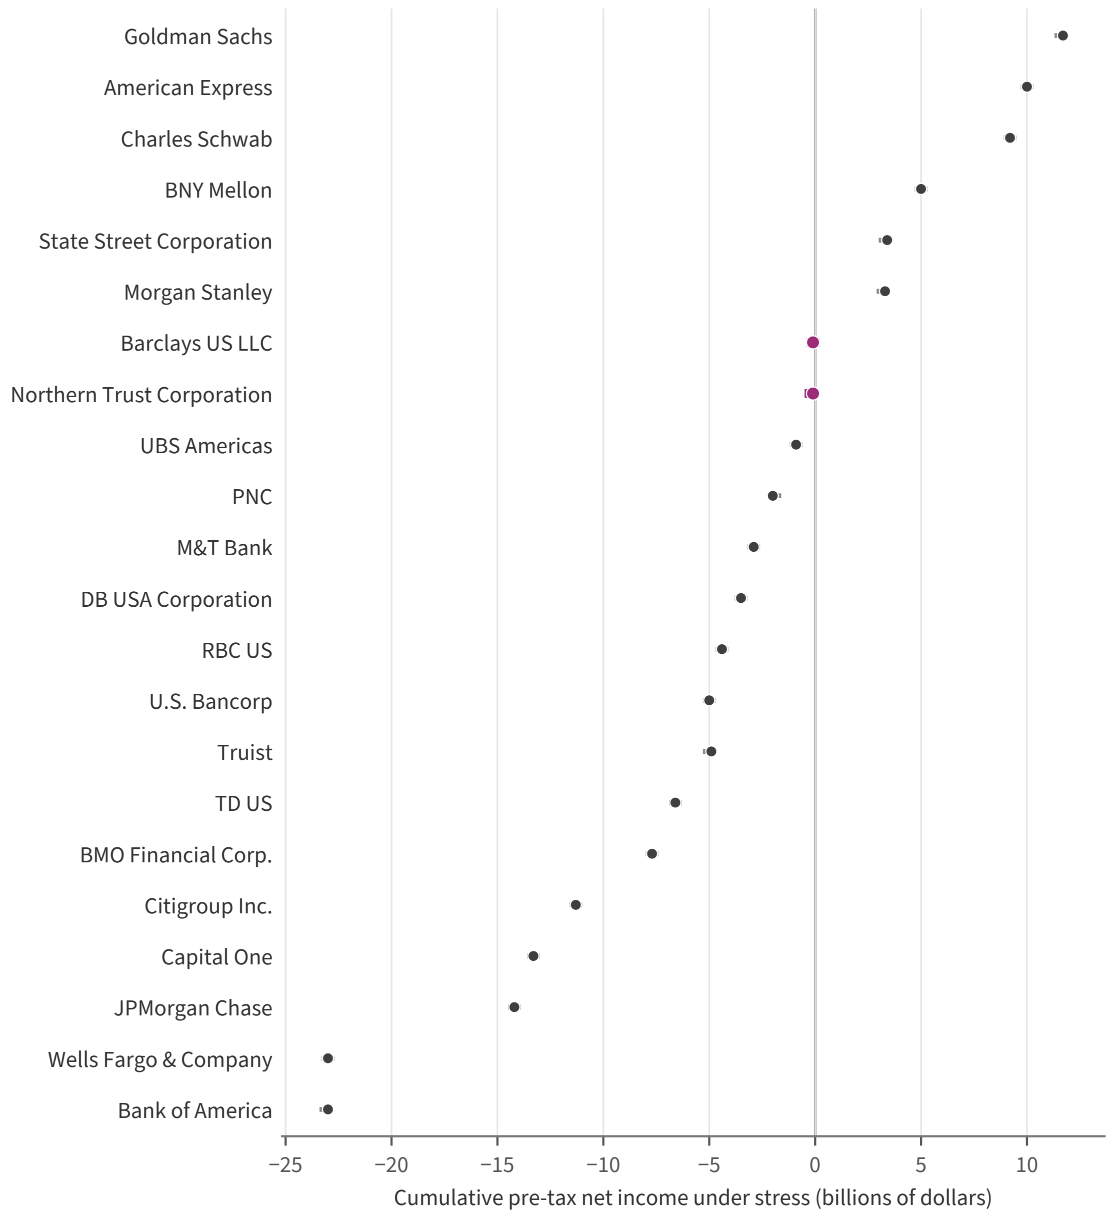

FIGURE 1. A public certificate for cumulative pre-tax net income

Note. Segments are the sharp intervals obtained from the six component disclosures in (41), treating each displayed amount as rounded to the nearest 0.1 billion. Magenta marks the two intervals that cross zero; points show the separately reported cumulative pre-tax net-income values. The sample contains the 22 bank-level rows in the 2025 Federal Reserve stress-test results file and excludes the aggregate row. Source: Board of Governors of the Federal Reserve System, 2025 stress-test results data.

$d_s = r_s - r_{s-1}$ lies in $[\ell_s, u_s]$. For $t_{j-1} \le t \le t_j$, define

$$\underline{r}_t = \max\left\{ r_{t_{j-1}} + \sum_{s=t_{j-1}+1}^{t} \ell_s,\ r_{t_j} - \sum_{s=t+1}^{t_j} u_s \right\}. \tag{43}$$

The first term follows the lowest admissible increments from the preceding checkpoint. The second retains enough room to reach the next checkpoint.

PROPOSITION 5 (Sharp path-minimum completion). *Suppose every checkpoint segment admits an increment sequence satisfying its endpoint change and the bounds $\ell_s \le d_s \le u_s$. Let $m_{\mathcal{H}}(r) = \min_{t\in\mathcal{H}} r_t$ for a nonempty set of target dates $\mathcal{H}$. The sharp lower endpoint of the compatible values of $m_{\mathcal{H}}(r)$ is*

$$\underline{m} = \min_{t\in\mathcal{H}} \underline{r}_t. \tag{44}$$

Partition the target dates across segments by $\mathcal{H}_j = \mathcal{H} \cap \{t_{j-1}+1, \ldots, t_j\}$, assigning date 0 to $\mathcal{H}_1$ when it belongs to $\mathcal{H}$. For each nonempty $\mathcal{H}_j$, define

$$\begin{aligned}
\overline{m}_j = \max_{d,m} \quad & m \\
\textit{subject to} \quad & r_{t_{j-1}} + \sum_{s=t_{j-1}+1}^{t} d_s \ge m \qquad (t \in \mathcal{H}_j), \\
& \ell_s \le d_s \le u_s \qquad\qquad\qquad (t_{j-1} < s \le t_j), \\
& \sum_{s=t_{j-1}+1}^{t_j} d_s = r_{t_j} - r_{t_{j-1}}.
\end{aligned} \tag{45}$$

The sharp upper endpoint is $\overline{m} = \min_{j:\mathcal{H}_j\neq\varnothing} \overline{m}_j$, and the sharp target set is $[\underline{m}, \overline{m}]$.

The result gives each checkpoint a target-specific disclosure value. For a regulatory threshold κ, the disclosed checkpoints and change bounds certify that every compatible path stays above the threshold exactly when $\underline{m} \ge \kappa$. If $\underline{m} < \kappa \le \overline{m}$, the public record admits paths on both sides of the threshold. Module costs can then be placed on candidate checkpoints and the least-cost certificate program can select a release whose lower envelope clears the stated threshold.

The distinction also changes how disclosure burden is evaluated. Detailed model documentation can support review of the mapping from scenarios and bank characteristics to projected components. A target certificate records the shorter set of disclosed values that establishes one stated conclusion. The two forms of transparency can therefore be assessed separately, with the certificate report naming the conclusion each disclosure carries.

8. Conclusion

The paper starts from local theoretical, measurement, and empirical restrictions and asks for the least complicated economy that gets a declared counterfactual right. Fragmented evidence can answer that counterfactual while leaving many complete economies. A target certificate records when this happens and supplies the formula that combines the relevant modules. In the linear benchmark, its support gives the exact least complicated economy, while the same equation gives target identification, minimum repair, efficient combination, and an attained approximation frontier. The module menu can generate arbitrary monotone patterns of complementarity, which makes the economic design of that menu part of the analysis.

The global report begins one step earlier. It checks whether equality and inequality modules admit a common completion. Conditional on completion, dual target certificates return the sharp target interval and the support of a least-cost carrying bundle. A joint confidence region for the module values preserves coverage after the bundle and certificate are selected. A decision problem then converts the remaining interval into regret.

The stress-test application shows the distinction in public data. Revenue and loss modules carry cumulative pre-tax income up to reporting precision. Beginning, ending, and cumulative-change modules leave an interior capital minimum undetermined. Period-specific change bounds and disclosed checkpoints yield a sharp interval for that minimum and a direct threshold certificate. A disclosure certificate therefore states the conclusion supported by released components and the path information that carries a different conclusion.

This sequence gives the least complicated economy an operational output: a verified economic resolution for a declared counterfactual, together with its completion witness, target certificate or ambiguity interval, cost record, sampling uncertainty, and decision consequence. Its empirical value is measured by the distinctions that can be suppressed, the records that can remain unlinked, and the next observation the certificate directs the analyst to collect.

References

Andrews, Donald W. K. and Gustavo Soares. 2010. "Inference for Parameters Defined by Moment Inequalities Using Generalized Moment Selection." *Econometrica* 78 (1): 119–157. `10.3982/ECTA7502`.

Armstrong, Timothy B. and Michal Kolesár. 2021. "Sensitivity Analysis Using Approximate Moment Condition Models." *Quantitative Economics* 12 (1): 77–108. `10.3982/QE1609`.

Bardhi, Arjada. 2024. "Attributes: Selective Learning and Influence." *Econometrica* 92 (2): 311–353. `10.3982/ECTA18355`.

Bardhi, Arjada and Steven Callander. 2026. "Learning in a Correlated World." *Annual Review of Economics* 18: 391–415. `10.1146/annurev-economics-051624-072515`.

Belloni, Alexandre, Victor Chernozhukov, Denis Chetverikov, and Ying Wei. 2018. "Uniformly Valid Post-Regularization Confidence Regions for Many Functional Parameters in Z-Estimation Framework." *Annals of Statistics* 46 (6B): 3643–3675. `10.1214/17-AOS1671`.

Beresteanu, Arie, Ilya Molchanov, and Francesca Molinari. 2011. "Sharp Identification Regions in Models With Convex Moment Predictions." *Econometrica* 79 (6): 1785–1821. `10.3982/ECTA8680`.

Blackwell, David. 1951. "Comparison of Experiments." In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pp. 93–102. University of California Press.

Blackwell, David. 1953. "Equivalent Comparisons of Experiments." *Annals of Mathematical Statistics* 24 (2): 265–272. `10.1214/aoms/1177729032`.

Bontemps, Christian, Thierry Magnac, and Eric Maurin. 2012. "Set Identified Linear Models." *Econometrica* 80 (3): 1129–1155. `10.3982/ECTA7637`.

Carneiro, Pedro, Sokbae Lee, and Daniel Wilhelm. 2020. "Optimal Data Collection for Randomized Control Trials." *The Econometrics Journal* 23 (1): 1–31. `10.1093/ectj/utz020`.

Chamberlain, Gary. 1987. "Asymptotic Efficiency in Estimation with Conditional Moment Restrictions." *Journal of Econometrics* 34 (3): 305–334. `10.1016/0304-4076(87)90015-7`.

Chernozhukov, Victor, Whitney K. Newey, and Rahul Singh. 2022. "Debiased Machine Learning of Global and Local Parameters Using Regularized Riesz Representers." *The Econometrics Journal* 25 (3): 576–601. `10.1093/ectj/utac002`.

Elfving, Gustav. 1952. "Optimum Allocation in Linear Regression Theory." *Annals of Mathematical Statistics* 23 (2): 255–262. `10.1214/aoms/1177729442`.

Fan, Yanqin, Robert Sherman, and Matthew Shum. 2014. "Identifying Treatment Effects under Data Combination." *Econometrica* 82 (2): 811–822. `10.3982/ECTA10571`.

Federal Reserve Board. 2025a. "2025 Federal Reserve Stress Test Results." , Board of Governors of the Federal Reserve System, Washington, DC.

Federal Reserve Board. 2025b. "Enhanced Transparency and Public Accountability of the Supervisory Stress Test." , Board of Governors of the Federal Reserve System, Washington, DC.

Goldstein, Itay and Yaron Leitner. 2018. "Stress Tests and Information Disclosure." *Journal of Economic Theory* 177: 34–69. `10.1016/j.jet.2018.05.013`.

Goldstein, Itay and Haresh Sapra. 2014. "Should Banks' Stress Test Results Be Disclosed? An Analysis of the Costs and Benefits." *Foundations and Trends in Finance* 8 (1): 1–54. `10.1561/0500000038`.

Gu, Jiaying, Thomas M. Russell, and Thomas Stringham. 2026. "Counterfactual Identification and Latent Space Enumeration in Discrete Outcome Models." *Review of Economic Studies* 93 (3): 1847–1888. `10.1093/restud/rdaf058`.

Hansen, Lars Peter. 1982. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica* 50 (4): 1029–1054. `10.2307/1912775`.

Kaji, Tetsuya. 2021. "Theory of Weak Identification in Semiparametric Models." *Econometrica* 89 (2): 733–763. `10.3982/ECTA17847`.

Karchmer, Mauricio and Avi Wigderson. 1993. "On Span Programs." In *Proceedings of the Eighth Annual Structure in Complexity Theory Conference*, pp. 102–111. IEEE. `10.1109/SCT.1993.336536`.
Karp, Richard M. 1972. "Reducibility Among Combinatorial Problems." In *Complexity of Computer Computations*, edited by Raymond E. Miller and James W. Thatcher, pp. 85–103. New York: Plenum Press. `10.1007/978-1-4684-2001-2_9`.
Komarova, Tatiana, Denis Nekipelov, and Evgeny Yakovlev. 2018. "Identification, Data Combination, and the Risk of Disclosure." *Quantitative Economics* 9 (1): 395–440. `10.3982/QE568`.
Manski, Charles F. 2021. "Econometrics for Decision Making: Building Foundations Sketched by Haavelmo and Wald." *Econometrica* 89 (6): 2827–2853. `10.3982/ECTA19279`.
Ridder, Geert and Robert Moffitt. 2007. "The Econometrics of Data Combination." In *Handbook of Econometrics*, vol. 6B, edited by James J. Heckman and Edward E. Leamer, pp. 5469–5547. Elsevier. `10.1016/S1573-4412(07)06075-8`.
Rockafellar, R. Tyrrell. 1970. *Convex Analysis*. Princeton, NJ: Princeton University Press.
Severini, Thomas A. and Gautam Tripathi. 2012. "Efficiency Bounds for Estimating Linear Functionals of Nonparametric Regression Models with Endogenous Regressors." *Journal of Econometrics* 170 (2): 491–498. `10.1016/j.jeconom.2012.05.018`.
Zhang, Haotian, Raid Ayoub, and Shreyas Sundaram. 2017. "Sensor Selection for Kalman Filtering of Linear Dynamical Systems: Complexity, Limitations and Greedy Algorithms." *Automatica* 78: 202–210. `10.1016/j.automatica.2016.12.025`.

Appendix A. Proofs

PROOF OF THEOREM 1. Choose θ_0 with $A\theta_0 = \mu$. Every element of the completion fiber has the form $\theta_0 + h$ for $h \in \ker A$. Thus $B\theta$ is constant on the fiber exactly when $Bh = 0$ for every $h \in \ker A$, which proves the equivalence of statements 1 and 2. In finite-dimensional Euclidean spaces, $\mathrm{row}(A) = (\ker A)^{\perp}$. Applying this identity row by row to B proves the equivalence of statements 2 and 3. Statement 3 holds exactly when every row of B is a linear combination of the rows of A, which is statement 4. Statement 4 implies statement 5 with $d(x) = \Lambda x$, and statement 5 implies statement 1. If the conditions fail, choose $h \in \ker A$ with $Bh \neq 0$ and set $\theta_1 = \theta_0 + h$. □

PROOF OF COROLLARY 1. The completion fiber is $\theta_0 + \ker A$. Applying the linear map B gives (7). The image is a singleton exactly when $B(\ker A) = \{0\}$, which is statement 2 of Theorem 1. □

PROOF OF THEOREM 2. Let $U = \mathrm{row}(A)$ and let $\pi : \mathbb{R}^p \to \mathbb{R}^p/U$ be the quotient map. The target is carried after adding C exactly when $\mathrm{row}(B) \subseteq U + \mathrm{row}(C)$, or equivalently when $\pi\{\mathrm{row}(B)\} \subseteq \pi\{\mathrm{row}(C)\}$. The dimension of $\pi\{\mathrm{row}(B)\}$ is

$$\dim\{U + \mathrm{row}(B)\} - \dim(U) = \mathrm{rank}\begin{bmatrix} A \\ B \end{bmatrix} - \mathrm{rank}(A) = \delta_B(A).$$

Since $\dim \pi\{\mathrm{row}(C)\} \leq \mathrm{rank}(C) \leq r$, every successful repair has at least $\delta_B(A)$ scalar rows. Choose rows of C whose quotient images form a basis of $\pi\{\mathrm{row}(B)\}$. This construction uses $\delta_B(A)$ rows and satisfies $\mathrm{row}(B) \subseteq U + \mathrm{row}(C)$. □

Proof of Theorem 3. If $B = \Lambda A_J$, discarding every block outside $\mathrm{supp}_J(\Lambda)$ gives $B = \Lambda_{\mathrm{supp}} A_{\mathrm{supp}}$, so the support carries the target. Conversely, if S carries the target, Theorem 1 supplies a matrix Λ_S satisfying $B = \Lambda_S A_S$. Extending Λ_S by zero blocks gives a full-menu certificate whose support is contained in S. These two maps preserve feasibility and weakly reduce cost because module costs are nonnegative, which proves equality of the minima and the stated description of the minimizing collections. If the full menu has no certificate, Theorem 1 gives $h \in \ker A_J$ with $Bh \neq 0$. Since $\ker A_J \subseteq \ker A_S$ for every $S \subseteq J$, the same direction witnesses failure for every subcollection. □

Proof of Proposition 1. Use the construction in (11). Each singleton collection has target defect one, equal to the empty collection's defect, while the pair has defect zero. A rule based on strictly positive one-module defect reductions therefore stops at the empty collection. □

Proof of Theorem 4. For statement 1, let $\mathcal{M}$ be the minimal elements of $\mathfrak{A}$. Introduce one common coordinate indexed by 0. For every $T \in \mathcal{M}$, choose a distinguished element $d(T) \in T$ and introduce private coordinates (T, i) for $i \in T \setminus \{d(T)\}$. The target row is $b^\top = e_0^\top$. Module i contains the row $e_{T,i}^\top$ whenever $i \in T \setminus \{d(T)\}$ and contains

$$e_0^\top - \sum_{i \in T \setminus \{d(T)\}} e_{T,i}^\top$$

whenever $i = d(T)$. If S contains T, summing the displayed row and the private rows for that T gives $e_0^\top$, so S carries the target. Conversely, consider a linear combination of rows available to S that equals $e_0^\top$. Write a_T for the coefficient on the displayed row for T. A missing private row forces $a_T = 0$ because its private coordinate occurs nowhere else, and a missing distinguished module makes the displayed row unavailable. Thus a_T can be nonzero only when $T \subseteq S$. The coefficient on e_0 requires $\sum_T a_T = 1$, so at least one minimal authorized set is contained in S. Hence $S \in \mathfrak{A}$.

For statement 2, take an instance of weighted set cover with universe $\{1, \ldots, p\}$ and sets F_j. Let the target row be $b^\top = \sum_{i=1}^p e_i^\top$ and let module j report the coordinate rows $e_i^\top$

for $i \in F_j$. Then $b^\top \in \mathrm{row}(A_S)$ exactly when $\bigcup_{j \in S} F_j$ covers the universe. Module costs equal the set-cover weights, so a least-cost carrying collection solves weighted set cover.

For statement 3, the construction in (11) violates diminishing returns: the gain from module 2 is zero at the empty collection and one after module 1 is present. To violate increasing returns, let either of two modules report the scalar target directly. The gain from module 2 is one at the empty collection and zero after module 1 is present. □

Proof of Theorem 5. Set $M = A^\top \Omega^{-1} A$. Positive definiteness of Ω gives $\mathrm{range}(M) = \mathrm{range}(A^\top)$. Since $\mathrm{row}(B) \subseteq \mathrm{row}(A)$, each column of $B^\top$ belongs to $\mathrm{range}(M)$ and $BM^\dagger M = B$. Hence

$$\Lambda_* A = BM^\dagger A^\top \Omega^{-1} A = BM^\dagger M = B.$$

For another exact certificate, write $\Lambda = \Lambda_* + D$. Then $DA = 0$. The cross term vanishes because

$$\Lambda_* \Omega D^\top = BM^\dagger A^\top D^\top = BM^\dagger (DA)^\top = 0.$$

Therefore

$$\Lambda \Omega \Lambda^\top = \Lambda_* \Omega \Lambda_*^\top + D \Omega D^\top \succeq \Lambda_* \Omega \Lambda_*^\top.$$

Finally, the Moore–Penrose identities and $BM^\dagger M = B$ give $\Lambda_* \Omega \Lambda_*^\top = BM^\dagger B^\top$. □

Proof of Theorem 6. Let $r = b - A^\top \lambda$. Write $x = G^{1/2} h$. The constraint becomes $\|x\|_2 \le R$ and $r^\top h = (G^{-1/2} r)^\top x$. Cauchy–Schwarz gives the right side of (18), and equality is attained at $x = R G^{-1/2} r / \|G^{-1/2} r\|_2$ when $r \neq 0$. Minimizing the squared bound over λ is weighted least squares. Its normal equation is (19), and its residual norm is the G^{-1} distance from b to $\mathrm{row}(A)$. □

Proof of Theorem 7. On the event $\mu_0 \in \mathcal{Q}_n$, an exact selected certificate satisfies

$$B\theta_0 = \widehat{\Lambda} A_{\widehat{S}} \theta_0 = \widehat{\Lambda}(\mu_0)_{\widehat{S}} \in \mathcal{I}_n(\widehat{S}, \widehat{\Lambda}).$$

This inclusion holds pointwise for every realized selection rule. When no certificate is returned, coverage is immediate from the convention $\mathcal{I}_n = \mathbb{R}^k$. Therefore the coverage probability is bounded below by $P\{\mu_0 \in \mathcal{Q}_n\}$, and (24) follows from (22). □

Proof of Proposition 2. The fundamental theorem of linear algebra gives

$$\mathrm{range}(A)^\perp = \ker(A^\top).$$

Thus $\mu \in \mathrm{range}(A)$ exactly when $z^\top \mu = 0$ for every $z \in \ker(A^\top)$. If μ lies outside the range, its orthogonal projection onto $\ker(A^\top)$ supplies a vector with nonzero inner product. Conditional on completion, the second alternative is Theorem 1. □

Proof of Theorem 8. For statement 1, Farkas' lemma applied to $A_S\theta = \mu_S$ and $G_S\theta \leq g_S$ gives exactly the alternative in (30). For statement 2, linear-programming duality gives

$$\sup_{\theta \in \mathcal{C}_S} d^\top \theta = U_S(d)$$

for every direction d with a finite primal optimum. Apply the identity to $d = b$ and $d = -b$.

Suppose the certificate pair in statement 3 exists. For every $\theta \in \mathcal{C}_S$,

$$b^\top \theta = (y^+)^\top \mu_S + (z^+)^\top G_S\theta \leq (y^+)^\top \mu_S + (z^+)^\top g_S$$

and, using the certificate for $-b$,

$$-b^\top \theta \leq (y^-)^\top \mu_S + (z^-)^\top g_S.$$

The two upper bounds sum to zero by (34), so both inequalities bind and (35) follows. Hence S carries the target.

Conversely, suppose S carries the target with common value c. The two primal programs that maximize $b^\top\theta$ and $-b^\top\theta$ over $\mathcal{C}_S$ have finite optimal values c and $-c$. Strong duality supplies optimal dual pairs (y^+, z^+) and (y^-, z^-) satisfying (32) and (33). Their objectives sum to zero, which gives (34). □

Proof of Corollary 2. Let a certificate pair be supported on S. The full system has a completion, so $\mathcal{C}_S$ is nonempty. Theorem 8 then shows that S carries the target. Conversely, every target-carrying S has a certificate pair supported on its equality and inequality blocks by the same theorem. Minimizing cost on either side therefore gives the same collections and value. □

Proof of Proposition 3. For $a = 0$, the largest regret over $[L, U]$ is $\max\{0, U - \kappa\}$. For $a = 1$, it is $\max\{0, \kappa - L\}$. Taking the smaller of these two quantities gives zero when the interval lies on one side of κ and gives (39) when the interval straddles it. In the latter case, action 1 minimizes regret exactly when $\kappa - L \leq U - \kappa$, which is equivalent to $\kappa \leq (L + U)/2$. □

Proof of Proposition 4. Fix $\ell \leq \min\{r_0, r_T\}$ and an interior date $t^* \in \mathcal{H}$. Set $r_{t^*} = \ell$ and connect the three fixed points $(0, r_0)$, (t^*, ℓ), and (T, r_T) while keeping every other target-date value weakly above ℓ. The resulting increments telescope to $r_T - r_0$, so the path satisfies every stated cumulative module and has minimum ℓ over $\mathcal{H}$. Since ℓ can be arbitrarily small, the compatible minima have no finite lower bound. □

Proof of Proposition 5. Consider a date t in segment j and write $D_j = r_{t_j} - r_{t_{j-1}}$. The feasible prefix sum through t lies in the interval

$$\left[\max\left\{\sum_{s=t_{j-1}+1}^{t} \ell_s,\ D_j - \sum_{s=t+1}^{t_j} u_s\right\},\ \min\left\{\sum_{s=t_{j-1}+1}^{t} u_s,\ D_j - \sum_{s=t+1}^{t_j} \ell_s\right\}\right].$$

Every point in this interval is attainable because sums of interval-restricted increments fill the corresponding interval, separately for the prefix and suffix. Adding $r_{t_{j-1}}$ gives (43) as the smallest feasible value of r_t.

Every compatible path therefore satisfies $m_{\mathcal{H}}(r) \geq \min_{t\in\mathcal{H}} \underline{r}_t$. Choose a target date attaining this finite minimum and a compatible segment path attaining its pointwise lower endpoint. Complete the other segments arbitrarily. The resulting global path has target-date minimum at most $\min_{t\in\mathcal{H}} \underline{r}_t$, which proves (44) and attainment.

Program (45) gives the largest attainable target-date minimum in segment j, so every global path has target-date minimum at most $\overline{m}_j$. Choose a segment optimizer whenever $\mathcal{H}_j$ is nonempty and any feasible path in the remaining segments. Their checkpoint values agree, and their concatenation attains $\min_{j:\mathcal{H}_j\neq\varnothing} \overline{m}_j$. Compactness gives both endpoints; continuity of $m_{\mathcal{H}}(r)$ on the connected path polytope fills the interval between them. □