The Counterfactual Reconstruction Theorem

When Local Economies Form One Global Economy

Chae-Yeon Xon

Yonsei University

September 2026

Abstract

Economic counterfactuals are often learned in pieces. Each experiment can yield a minimal state on its own policy domain, and overlaps can make the states locally compatible. A common economic history imposes a further requirement: the compatible pieces must admit one global lift. This paper represents local histories and their declared future responses by a morphism of sheaves. The counterfactual reconstruction theorem identifies the minimal global economy and a connecting map whose image is the complete obstruction to a global lift. Every valid linear specification restriction factors uniquely through this map. Its dimension is the gap between local and global state counts. A baseline archive can reveal part of the obstruction; its kernel determines the remaining bridge problem. A downstream counterfactual selects a canonical target bridge state. Records identify the target precisely when they carry this state, equivalently when they preserve every finite decision problem whose payoff depends on the target. The theorem gives a sharp identified set, a necessary and attainable record count, and the exact amplification of bounded record error. Across several targets, record requirements form a polymatroid whose intersections measure information reuse. The residual target spectrum gives exact frontiers for record counts, normalized measurement noise, and continuous precision budgets. Singular-value brackets and Gaussian projection inference retain coverage across rank changes. An Oregon Medicaid application reconstructs approval rates across separate reports and gives bridge counts zero, one, and four for three declared targets.

JEL Codes: C10, C13, C14

Keywords: counterfactuals, state representation, identification, model combination

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

1. Introduction

An empirical literature rarely observes one model of an economy. One experiment records how a household responds to prices and transfers. Another records labor supply under a different policy menu. Administrative data connect some states across the two settings. A structural model supplies transition restrictions on another domain. Each source can support a minimal state for the counterfactuals it sees. The collection still leaves a separate question: do these local states arise from one economic history?

Agreement on overlaps establishes local compatibility. Global generation asks for a history whose restriction produces every local state at once. Consider three contexts carrying coordinates (x, y), (y, z), and (z, x). Pairwise agreement leaves the three-dimensional set indexed by (x, y, z). A common history may impose $z = x + y$, leaving two dimensions. The missing coordinate is visible in every local atlas and absent from every globally generated economy.

The paper places this distinction in a response morphism $T : \mathcal{H} \to \mathcal{Y}$. The history sheaf $\mathcal{H}$ records locally available pasts. The response sheaf $\mathcal{Y}$ records the declared future responses. The image sheaf $\mathcal{Z} = \operatorname{im} T$ contains the locally minimal response states, and the kernel sheaf $\mathcal{K} = \ker T$ contains local histories with the same declared responses. Two operations now compete. Gluing the local minima gives $\Gamma(\mathsf{X}, \mathcal{Z})$. Gluing histories first and then minimizing gives $\operatorname{im} \Gamma(T)$.

The counterfactual reconstruction theorem identifies the difference through one exact sequence:

$$0 \longrightarrow \Gamma(\mathcal{K}) \longrightarrow \Gamma(\mathcal{H}) \longrightarrow \Gamma(\mathcal{Z}) \xrightarrow{\partial} H^1(\mathcal{K}) \longrightarrow H^1(\mathcal{H}). \tag{1}$$

The image of the connecting map ∂ is the reconstruction obstruction. A compatible collection of locally minimal states has a common global history exactly when its image under ∂ is zero. The global minimal economy is $\operatorname{im} \Gamma(T)$, and every sufficient global representation maps onto it. Every valid linear specification restriction factors uniquely through the same connecting map. Thus ∂ supplies the lift criterion, the minimal state, and the complete family of linear diagnostics.

The obstruction also organizes measurement design. A fully estimated local atlas makes its coordinates available as specification restrictions. A fragmented archive may reveal only a linear image of the obstruction, leaving a missing subspace N. For any linear counterfactual or welfare target G, the quotient $N/(N \cap \ker G)$ is the target bridge state. Every identifying record system maps onto it. Selected links identify G precisely when their restrictions span the row space of $G|_N$, and this condition preserves the value of observing G in every finite decision problem. The resulting affine image is the sharp identified set. With unrestricted measurements, $\operatorname{rank}(G|_N)$ records attain the bound. Across a portfolio of targets, these ranks form a representable polymatroid, so overlap among missing row spaces measures the exact reuse of bridge records. The norm of the canonical decoder gives the sharp amplification of bounded record error. Full obstruction recovery uses $\dim N$ records; with a finite technology and record-specific costs, it is a minimum-weight basis problem. Under a smaller unrestricted budget, the residual target covariance gives the exact linear decision frontier. The same spectrum gives a closed-form frontier under homoskedastic uncorrelated noise and allocates a continuous precision budget across record directions. For a finite atlas, an Euler–Poincaré calculation links the global state dimension to local states, overlaps, and higher-order assembly terms.

Several mathematical components have established uses elsewhere. Minimal realization connects behavioral rank to state dimension (Ho and Kalman 1966; Willems 1986, 2007). Sheaves encode compatible heterogeneous data as global sections (Robinson 2017), and sheaf cohomology records the failure of global lifting in a short exact sequence (Bredon 1997). Contextuality research also uses global-section obstructions for families of local observations (Abramsky and Brandenburger 2011). Recent work applies sheaf Laplacians and cohomological diagnostics to generative causal models (Wu, Xie, and Li 2026). Functional observability and goal-oriented and A-optimal design study measurements chosen for a declared target or posterior risk (Montanari et al. 2022; Zhang, Fernando, and Darouach 2024; Attia, Alexanderian, and Saibaba 2018; Alexanderian et al. 2016). Comparison of experiments orders information by its value across decision problems (Blackwell 1953). The theorem here derives the missing target from the difference between local response images and common-history generation. Its universal factorization classifies the resulting specification restrictions. The relative acquisition result then conditions measurement choice and its decision value on the information already present in the archive.

The finite example makes the connecting map explicit as $\partial(x, y, z) = z - x - y$. An application to the Oregon Health Insurance Experiment reads outreach, adjudication, and program reports as local views of one administrative history. A projection result turns connecting-map coordinates into specification statistics, and rank-boundary confidence sets retain coverage when record dimensions change. Smooth and global sections state the regularity conditions under which the same reconstruction survives beyond finite vector spaces.

2. Local responses and global behavior

Let X be a finite policy-context space or a space equipped with a finite cover. A context can index an experiment, an administrative regime, a model domain, or an observed policy menu. Let $\mathcal{H}$ and $\mathcal{Y}$ be sheaves of finite-dimensional real vector spaces on X. A section of

$\mathcal{H}$ over a context records histories available there. A section of $\mathcal{Y}$ records response profiles over the declared future probes.[1]

A response morphism

$$T : \mathcal{H} \longrightarrow \mathcal{Y} \tag{2}$$

maps histories into their future-response profiles and commutes with restriction across contexts. Define the kernel and image sheaves

$$\mathcal{K} = \ker T, \qquad \mathcal{Z} = \operatorname{im} T. \tag{3}$$

The stalk of $\mathcal{K}$ contains locally silent history directions. The stalk of $\mathcal{Z}$ is the target-specific local response state. Using the sheaf image is consequential: a response state has a local history lift around every context, while a single lift over all contexts remains an additional requirement.

Write $\Gamma(\mathcal{A}) = \Gamma(X, \mathcal{A})$ for global sections. Two spaces follow from the same local response morphism:

$$\mathcal{L} = \Gamma(\mathcal{Z}), \qquad \mathcal{M} = \operatorname{im}\left[\Gamma(\mathcal{H}) \xrightarrow{\Gamma(T)} \Gamma(\mathcal{Z})\right]. \tag{4}$$

The space $\mathcal{L}$ glues the locally minimal response states. The space $\mathcal{M}$ contains the response states produced by one global history. Their quotient

$$\mathfrak{O} = \mathcal{L}/\mathcal{M}, \qquad o = \dim \mathfrak{O}, \tag{5}$$

is the reconstruction obstruction space.

The corresponding global behavior joins a history $h \in \Gamma(\mathcal{H})$ to a linear future probe $f \in \mathcal{L}^*$:

$$B(h, f) = f(\Gamma(T)h)\,. \tag{6}$$

A state realization of this behavior is a vector space W and maps $q : \Gamma(\mathcal{H}) \to W$ and $g : W \to \mathcal{L}$ satisfying $\Gamma(T) = gq$. Its reachable part is $\operatorname{im} q$. A realization is minimal when the reachable state has the smallest dimension among such factorizations.

LEMMA 1 (Behavioral image). *The minimal global state dimension is*

$$d^\star = \operatorname{rank} \Gamma(T) = \dim \mathcal{M}. \tag{7}$$

[1] A history is the common latent record from which the declared response family is generated. Its empirical content changes with the contexts and future probes included in the response morphism.

The image $\mathcal{M}$ attains the bound. Every realization has a unique response-preserving surjection from its reachable state onto $\mathcal{M}$, and this map is an isomorphism exactly for a minimal realization.

The declared future family enters through $\mathcal{Y}$ and T. Enlarging that family replaces the image sheaf and may retain more history directions. The reconstruction problem therefore begins after the target has been declared.[2]

3. A counterfactual atlas

Let $\mathcal{N}$ be the nerve of a finite cover of X. Its vertices index policy contexts and its faces index observed overlaps. For a sheaf $\mathcal{A}$ and a p-face σ, write $\mathcal{A}_\sigma$ for the section space on the corresponding overlap. The degree-p Čech cochain space is

$$C^p(\mathcal{A}) = \bigoplus_{\sigma \in \mathcal{N}_p} \mathcal{A}_\sigma. \tag{8}$$

Restriction maps and an orientation of the nerve define coboundaries $\delta^p_{\mathcal{A}} : C^p(\mathcal{A}) \to C^{p+1}(\mathcal{A})$ with $\delta^{p+1}_{\mathcal{A}} \delta^p_{\mathcal{A}} = 0$.

Assume the cover computes the cohomology of $\mathcal{H}$, $\mathcal{K}$, and $\mathcal{Z}$, and all groups used below are finite dimensional. Then

$$\mathcal{L} = H^0(\mathcal{Z}) = \ker \delta^0_{\mathcal{Z}}. \tag{9}$$

A vector in $\mathcal{L}$ assigns a target-specific minimal response state to every local context and makes those responses agree on every observed overlap.

The response morphism gives the short exact sequence of sheaves

$$0 \longrightarrow \mathcal{K} \longrightarrow \mathcal{H} \longrightarrow \mathcal{Z} \longrightarrow 0. \tag{10}$$

Exactness here is local. A section of $\mathcal{Z}$ has local history lifts. The existence of one global lift is determined after applying the global-section functor.

For later use, define

$$\chi(\mathcal{Z}) = \sum_{p=0}^{q} (-1)^p \dim C^p(\mathcal{Z}), \qquad \eta(\mathcal{Z}) = \sum_{p=1}^{q} (-1)^{p+1} \dim H^p(\mathcal{Z}). \tag{11}$$

[2] The image construction follows the minimal-realization lineage in Ho and Kalman (1966); Willems (1986, 2007). It also echoes exact aggregation, where restrictions determine when heterogeneous primitives admit a common reduced representation (Gorman 1968).

Euler–Poincaré gives

$$\dim \mathcal{L} = \chi(\mathcal{Z}) + \eta(\mathcal{Z}). \tag{12}$$

The first term uses local state and overlap dimensions. The second records the higher-order assembly terms of the response sheaf.

The notation now makes the order of operations explicit:

$$\mathrm{Glue}(\mathrm{Min}\, T) = \Gamma(\mathcal{Z}) = \mathcal{L}, \qquad \mathrm{Min}(\mathrm{Glue}\, T) = \mathrm{im}\,\Gamma(T) = \mathcal{M}. \tag{13}$$

The first expression minimizes responses inside each context and then glues them. The second glues histories and then takes the response image of the assembled economy.

DEFINITION 1 (Reconstruction diagnostic). *A linear map $A : \mathcal{L} \to \mathbb{R}^m$ is a valid reconstruction diagnostic when $\mathcal{M} \subseteq \ker A$. It is complete when*

$$\ker A = \mathcal{M}. \tag{14}$$

A valid diagnostic records restrictions satisfied by every response state generated by a common global history. A complete diagnostic separates those states from every remaining compatible local response.

Diagnostics and measurements are distinct objects. Once a compatible local state $z \in \mathcal{L}$ has been estimated, a coordinate of the connecting map is a computable restriction. Data acquisition enters when the baseline archive reveals only a linear image $P\pi(z)$ of the obstruction class $\pi(z) \in \mathfrak{O}$. A bridge technology is a collection of additional scalar measurements $a_j \in \mathfrak{O}^*$, with declared costs $c_j > 0$. Linked administrative units, paired experimental arms, and common simulator seeds can supply such measurements. The theorem states when a given technology completes the baseline archive and how many missing obstruction directions remain.[3]

4. The counterfactual reconstruction theorem

The local response morphism determines three objects: the glued local minimum $\mathcal{L}$, the globally generated minimum $\mathcal{M}$, and the obstruction space $\mathfrak{O} = \mathcal{L}/\mathcal{M}$. The theorem identifies all three through the connecting map of (10). For a linear record $R : \mathfrak{O} \to E$, a

[3]Functional observability and goal-oriented design begin from a fixed state or target and select informative sensors (Montanari et al. 2022; Attia, Alexanderian, and Saibaba 2018). Here the response images and baseline archive first determine the missing domain on which bridge measurements act.

finitely supported law μ on $\mathfrak{O}$, a finite action set A, and a bounded loss $\ell : W \times \mathsf{A} \to \mathbb{R}$, write

$$\mathfrak{B}(R; \mu, \ell) = \inf_{\delta : R(\operatorname{supp}\mu) \to \mathsf{A}} \mathbb{E}_{\mu}[\ell(GU, \delta(RU))]\,. \tag{15}$$

This is the smallest risk attainable from the record when the payoff-relevant state is GU.

THEOREM 1 (Counterfactual reconstruction). *Consider a response morphism $T : \mathcal{H} \to \mathcal{Y}$ on X, with $\mathcal{K} = \ker T$ and $\mathcal{Z} = \operatorname{im} T$. Suppose the cohomology groups below are finite dimensional and the chosen finite cover computes them. Then:*

1. *There is a natural exact sequence*

$$0 \longrightarrow \Gamma(\mathcal{K}) \longrightarrow \Gamma(\mathcal{H}) \xrightarrow{\Gamma(T)} \mathcal{L} \xrightarrow{\partial} H^1(\mathcal{K}) \longrightarrow H^1(\mathcal{H}). \tag{16}$$

   The connecting map induces a canonical isomorphism

$$\mathfrak{O} = \mathcal{L}/\mathcal{M} \cong \operatorname{im}\partial = \ker\Big[H^1(\mathcal{K}) \to H^1(\mathcal{H})\Big]. \tag{17}$$

2. *A compatible local response $z \in \mathcal{L}$ is generated by one global history exactly when $\partial z = 0$. Hence ∂z is a complete certificate of global-history failure.*
3. *The space $\mathcal{M} = \operatorname{im}\Gamma(T)$ is the minimal global state for the behavior* (6). *Every sufficient global state has a unique response-preserving surjection onto $\mathcal{M}$. Minimal global states are therefore unique up to a linear isomorphism.*
4. *The following statements are equivalent:*
   - *(a) local minimization and global gluing commute,* $\operatorname{Glue}(\operatorname{Min} T) = \operatorname{Min}(\operatorname{Glue} T)$*;*
   - *(b) every compatible collection of locally minimal responses has a global history lift;*
   - *(c)* $\mathfrak{O} = 0$*;*
   - *(d)* $\partial = 0$*;*
   - *(e) the map $H^1(\mathcal{K}) \to H^1(\mathcal{H})$ is injective.*
5. *The minimal global state dimension satisfies the reconstruction identity*

$$d^{\star} = \dim \mathcal{M} = \chi(\mathcal{Z}) + \eta(\mathcal{Z}) - o, \qquad o = \dim \mathfrak{O}. \tag{18}$$

6. *Let V be any finite-dimensional vector space. Composition with the quotient map $\pi : \mathcal{L} \to \mathfrak{O}$ and the isomorphism in* (17) *gives natural bijections*

$$\{A : \mathcal{L} \to V : \mathcal{M} \subseteq \ker A\} \cong \operatorname{Hom}(\mathfrak{O}, V) \cong \operatorname{Hom}(\operatorname{im}\partial, V). \tag{19}$$

Hence every valid linear reconstruction diagnostic factors uniquely through ∂. It is complete exactly when the induced map on $\operatorname{im}\partial$ *is injective. In particular,*

$$\min\{m : \textit{a complete m-coordinate diagnostic exists}\} = o. \tag{20}$$

7. *Let $P : \mathfrak{O} \to D$ collect the obstruction information already carried by a baseline archive, let $N = \ker P$, and let $G : \mathfrak{O} \to W$ be a declared linear counterfactual or welfare target. Consider a finite bridge technology $\mathcal{A} = \{a_1, \ldots, a_J\} \subset \mathfrak{O}^*$. For $S \subseteq \{1, \ldots, J\}$, write $A_S u = (a_j(u))_{j \in S}$ and $R_S = (P, A_S)$. The following statements are equivalent:*
   - *(a) G is identified from the selected archive: $R_S u = R_S u'$ implies $Gu = Gu'$;*
   - *(b)* $\ker R_S \subseteq \ker G$*;*
   - *(c) $G = C_S R_S$ for a linear decoder C_S;*
   - *(d)*

$$\operatorname{im}\big((G|_N)^*\big) \subseteq \operatorname{span}\{a_j|_N : j \in S\} \subseteq N^*. \tag{21}$$

   - *(e) for every finitely supported μ, finite A, and bounded ℓ,*

$$\mathfrak{B}(R_S; \mu, \ell) = \mathfrak{B}(G; \mu, \ell). \tag{22}$$

For an observed record $R_S u$, the sharp identified set is the affine space

$$Gu + G(\ker R_S). \tag{23}$$

Under unrestricted scalar measurements, the minimum number of additional records that identifies G is

$$r_G(P) = \operatorname{rank}(G|_N). \tag{24}$$

This count is the dimension of the canonical target bridge state

$$\mathfrak{Q}_{G,P} := N/(N \cap \ker G) \cong G(N). \tag{25}$$

Every identifying augmentation $A : N \to E$ induces a unique response-preserving surjection $\operatorname{im} A \to \mathfrak{Q}_{G,P}$*. Hence $\mathfrak{Q}_{G,P}$ is the minimal state carried by any archive augmentation that preserves the declared target. For full obstruction recovery, $G = I_{\mathfrak{O}}$, this count is $q = \dim N$ and condition* (21) *becomes $\operatorname{span}\{a_j|_N : j \in S\} = N^*$. When the dictionary spans N^*, a minimum-cost exact design under positive record costs is a minimum-weight basis of the restricted vector matroid. The increasing-cost greedy algorithm attains it.*

8. *Equip $\mathfrak{O}$, the record space E_S of R_S, and W with inner products. Suppose the selected archive is observed with an error e satisfying $\|e\| \le \varepsilon$, while u ranges over $\mathfrak{O}$. The sharp minimax target error is*

$$\inf_{\delta:E_S\to W} \sup_{\substack{u\in\mathfrak{O} \\ \|e\|\le\varepsilon}} \|Gu - \delta(R_S u + e)\| = \begin{cases} \varepsilon\|GR_S^{\dagger}\|_{\mathrm{op}}, & \ker R_S \subseteq \ker G, \\ +\infty, & \ker R_S \nsubseteq \ker G, \end{cases} \tag{26}$$

where $R_S^{\dagger}$ is the Moore–Penrose inverse. In the finite case, the decoder $\delta^{\star}(y) = GR_S^{\dagger}y$ attains the bound. Thus $\|GR_S^{\dagger}\|_{\mathrm{op}}$ is the exact stability price of the selected archive.

9. *Equip W with an inner product and let U be a centered square-integrable random obstruction. Write $\widetilde{V}_{G,P}$ for the residual from the best linear prediction of GU using PU, and let $\lambda_1(G,P) \ge \cdots \ge \lambda_{\dim W}(G,P) \ge 0$ be its covariance eigenvalues. If at most m unrestricted scalar bridge records may be added to the baseline archive, then*

$$\mathcal{R}_m^{\star}(G,P) := \inf_{\substack{B:\mathfrak{O}\to\mathbb{R}^m \\ C_0:D\to W,\ C_1:\mathbb{R}^m\to W}} \mathbb{E}\|GU - C_0PU - C_1BU\|^2 = \sum_{j>m} \lambda_j(G,P). \tag{27}$$

Records equal to the first m principal coordinates of $\widetilde{V}_{G,P}$ attain the bound. If the baseline archive already identifies G, all residual eigenvalues vanish. If it reveals no obstruction coordinate and $G = I_{\mathfrak{O}}$, the exact count is 0, which gives the unconditional diagnostic dimension in (20).

10. *Let $\widetilde{U}_P$ be the residual from the best linear prediction of U using PU, and let $\Sigma_P = \mathrm{Cov}(\widetilde{U}_P)$. Suppose the added record is*

$$Y_B = BU + \varepsilon, \qquad \mathbb{E}\varepsilon = 0, \qquad \mathrm{Cov}(\varepsilon) = \sigma^2 I_m, \qquad \mathrm{Cov}(\varepsilon, U) = 0, \tag{28}$$

and impose the signal normalization $B\Sigma_P B^* = I_m$. For $m \le \operatorname{rank}\Sigma_P$, the minimum best-linear risk is

$$\mathcal{R}_{m,\sigma}^{\star}(G,P) := \inf_{\substack{B\Sigma_P B^*=I_m \\ C_0:D\to W,\ C_1:\mathbb{R}^m\to W}} \mathbb{E}\|GU - C_0PU - C_1Y_B\|^2 \tag{29}$$

$$= \sum_j \lambda_j(G,P) - \frac{1}{1+\sigma^2}\sum_{j=1}^{m} \lambda_j(G,P) \tag{30}$$

$$= \sum_{j>m} \lambda_j(G,P) + \frac{\sigma^2}{1+\sigma^2}\sum_{j=1}^{m} \lambda_j(G,P), \tag{31}$$

where the eigenvalue sequence is padded by zeros. A normalized record aligned with the first m eigen-directions attains the bound. Thus observation noise discounts the value of each selected direction by the common factor $(1+\sigma^2)^{-1}$*, and* (27) *is the case* $\sigma^2 = 0$.

The theorem gives the obstruction a common role across theory, testing, and acquisition. It is the cokernel of global response generation, the image of the connecting map, the exact state-dimension loss, the universal domain of every valid linear diagnostic, and the missing state in the acquisition problem. The relative formulation matters empirically. A fully estimated local atlas makes the obstruction diagnostic computable. A fragmented archive can leave a subspace N that linked observations must recover.[4]

The connecting map is constructive on a finite cover. Given $z \in \mathcal{L}$, choose a local history lift on every context. Their differences on overlaps lie in $\mathcal{K}$ and form a one-cocycle. Changing the local lifts changes that cocycle by a coboundary. Its cohomology class is ∂z. The class vanishes precisely when the local lifts can be adjusted to agree as one global history.[5]

COROLLARY 1 (Acyclic-history form). *If* $H^1(\mathcal{H}) = 0$*, then*

$$\mathfrak{O} \cong H^1(\mathcal{K}). \tag{32}$$

Global reconstruction holds exactly when $H^1(\mathcal{K}) = 0$*. If the baseline archive reveals no obstruction coordinate, the number of additional scalar measurements required for exact recovery is* $\dim H^1(\mathcal{K})$.

COROLLARY 2 (Value of bridge information). *Suppose a planner chooses a linear estimate* $\widehat{V}$ *of* GU *after observing the baseline archive and at most* m *additional scalar records, with loss* $\|GU - \widehat{V}\|^2$*. With noiseless records, the smallest expected loss is* $\mathcal{R}^\star_m(G,P)$*, the value of the first* m *records is* $\sum_{j=1}^m \lambda_j(G,P)$*, and the marginal value of record* m *is* $\lambda_m(G,P)$*. Under* (28) *and the signal normalization, these two values become*

$$\frac{1}{1+\sigma^2}\sum_{j=1}^{m}\lambda_j(G,P) \qquad \textit{and} \qquad \frac{\lambda_m(G,P)}{1+\sigma^2}. \tag{33}$$

COROLLARY 3 (Precision allocation). *Let* $q = \operatorname{rank}\Sigma_P$*, pad* $\{\lambda_j(G,P)\}$ *to length* q*, and let a planner choose residual-signal-orthonormal directions* $b_1,\ldots,b_q$ *together with precisions* $\tau_j \geq 0$.

[4]The count $r_G(P)$ allows unrestricted noiseless scalar records. A constrained measurement dictionary can require more coordinates or leave the target partially identified; condition (21) gives the corresponding criterion.

[5]The decision-sufficiency equivalence uses finite action sets because a two-state, two-action loss already separates any target variation left inside a record fiber. Extensions to general action spaces add the corresponding measurability and attainment conditions.

The added records are

$$Z_j = \sqrt{\tau_j}\, b_j U + \varepsilon_j, \qquad b_i \Sigma_P b_j^* = \mathbf{1}\{i = j\}, \qquad \mathrm{Cov}(\varepsilon) = I_q, \qquad \mathrm{Cov}(\varepsilon, U) = 0. \tag{34}$$

Under the budget $\sum_j \tau_j \le b$, the minimum best-linear risk is

$$\mathcal{R}^\star(b; G, P) = \min_{\substack{\tau_j \ge 0 \\ \sum_j \tau_j \le b}} \sum_{j=1}^{q} \frac{\lambda_j(G,P)}{1+\tau_j} = \sum_{j=1}^{q} \frac{\lambda_j(G,P)}{1+\tau_j^\star(b)}. \tag{35}$$

For $b > 0$ and $\lambda_1(G,P) > 0$, the allocation is

$$\tau_j^\star(b) = \left(\sqrt{\frac{\lambda_j(G,P)}{\nu(b)}} - 1\right)_+, \qquad \sum_{j=1}^{q}\left(\sqrt{\frac{\lambda_j(G,P)}{\nu(b)}} - 1\right)_+ = b. \tag{36}$$

The records follow the residual target eigen-directions, an eigen-direction receives positive precision when $\lambda_j(G,P) > \nu(b)$, and $\nu(b)$ is the marginal value of the precision budget.

COROLLARY 4 (Archive refinement). *Let $P_1 : \mathfrak{O} \to D_1$ and $P_2 : \mathfrak{O} \to D_2$ be two baseline archives for the same target G, and suppose $P_1 = HP_2$ for a linear map H. Then, for every $u \in \mathfrak{O}$ and every record budget m,*

$$Gu + G(\ker P_2) \subseteq Gu + G(\ker P_1), \qquad r_G(P_2) \le r_G(P_1), \qquad \mathcal{R}_m^\star(G, P_2) \le \mathcal{R}_m^\star(G, P_1). \tag{37}$$

For every additional record map $A : \mathfrak{O} \to E$, finitely supported μ, finite action set A, and bounded loss ℓ,

$$\mathfrak{B}\big((P_2, A); \mu, \ell\big) \le \mathfrak{B}\big((P_1, A); \mu, \ell\big). \tag{38}$$

COROLLARY 5 (Target ordering). *Suppose two declared future families induce response morphisms with global minimal states $\mathcal{M}_1$ and $\mathcal{M}_2$, and the first family is obtained by restricting the second. Then there is a canonical surjection $\mathcal{M}_2 \to \mathcal{M}_1$, so $\dim \mathcal{M}_1 \le \dim \mathcal{M}_2$.*

COROLLARY 6 (A portfolio of counterfactual targets). *Fix a baseline archive P and let $G_i : \mathfrak{O} \to W_i$, $i \in I$, be a finite family of counterfactual or welfare targets. For $S \subseteq I$, let $G_S = (G_i)_{i \in S}$ and define $\rho_P(S) = r_{G_S}(P)$. Then*

$$\rho_P(S) = \dim \sum_{i \in S} \mathrm{im}\big((G_i|_N)^*\big), \qquad N = \ker P. \tag{39}$$

The set function ρ_P is normalized, increasing, and submodular. In particular,

$$\rho_P(S) + \rho_P(T) \ge \rho_P(S \cup T) + \rho_P(S \cap T). \tag{40}$$

If $i \notin S$ and $U_S = \sum_{j \in S} \operatorname{im}((G_j|_N)^*)$, the new record dimension contributed by target i is

$$\rho_P(S \cup \{i\}) - \rho_P(S) = \dim \operatorname{im}((G_i|_N)^*) - \dim\big(\operatorname{im}((G_i|_N)^*) \cap U_S\big)\,. \tag{41}$$

Thus the same bridge coordinate can serve several downstream decisions, and the overlap of their missing row spaces gives the exact reuse.

The target-ordering result follows from response restriction. The portfolio result applies the same quotient construction to a stacked target and shows how a common archive serves several future questions. The descent obstruction can move in either direction when the target family changes because both the image sheaf and its kernel sheaf change. The acquisition problem is therefore posed after the future family and the baseline archive have been declared.

5. Three local economies

5.1. A finite cycle

Three policy contexts report the locally minimal response states

$$X_1 = (x, y), \qquad X_2 = (y, z), \qquad X_3 = (z, x). \tag{42}$$

Stack them as $v = (x_1, y_1, y_2, z_2, z_3, x_3)^\top \in \mathbb{R}^6$. Pairwise agreement is $Dv = 0$, where

$$D = \begin{pmatrix} 0 & 1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & -1 & 0 \\ -1 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}. \tag{43}$$

The matrix has rank three. Hence the glued local minimum $\mathcal{L} = \ker D$ has dimension three and is indexed by (x, y, z). Each local two-coordinate projection is free, so pairwise agreement preserves every local state.

A common global history has coordinates $h = (u, v) \in \mathbb{R}^2$ and produces

$$Qh = \begin{pmatrix} u & v & v & u+v & u+v & u \end{pmatrix}^\top. \tag{44}$$

Thus $DQ = 0$, $\operatorname{rank} Q = 2$, and the globally generated states satisfy

$$z = x + y. \tag{45}$$

Every local context remains two dimensional because (u, v), $(v, u+v)$, and $(u+v, u)$ each give invertible coordinates for the common history.

In the cochain representation of the short exact sequence, the connecting map is

$$\partial : \mathcal{L} \to \mathbb{R}, \qquad \partial(x, y, z) = z - x - y. \tag{46}$$

Its kernel is $\mathcal{M} = \operatorname{im} Q$, its image is one dimensional, and

$$\mathfrak{O} \cong \mathbb{R}, \qquad o = 1, \qquad d^{\star} = \dim \mathcal{L} - o = 2. \tag{47}$$

The scalar in (46) is a complete reconstruction diagnostic. If the three local coordinates are jointly estimated, it is computed directly and the baseline map on $\mathfrak{O}$ is the identity. The diagnostic then requires no new observation. If the baseline archive contains only separate context records and carries no unit-level obstruction information, $P = 0$ and $N = \mathfrak{O}$. One linked scalar measurement proportional to $z - x - y$ completes the archive; every scalar measurement orthogonal to this direction leaves the obstruction unresolved.

The relative formulation also separates available and useful links. Suppose an institution can collect candidate scalar records $a_j(u) = \alpha_j u$ at costs c_j. For a downstream welfare gap $G(u) = \gamma u$, the baseline identifies the gap when $\gamma = 0$. When $\gamma \neq 0$, one record with $\alpha_j \neq 0$ is necessary and sufficient, and the minimum-cost design selects the cheapest such record. With a centered population obstruction of variance λ, a quadratic decision loses $\gamma^2\lambda$ before linkage and zero after that record. Thus the state gap, the test restriction, and the value of linkage all use the same one-dimensional class while remaining different empirical operations.

The example separates pairwise compatibility from generation by one history. Every chart is locally minimal, every overlap agrees, and the cycle carries one response direction unsupported by the common economy. The connecting class records that direction directly.

Changing the declared future family changes the image sheaf before the theorem is applied. If the future uses only (x, y), the target-specific local image discards z and the obstruction vanishes. The comparison shows why local minimality and bridge complexity are both defined relative to the declared responses.

5.2. An Oregon Medicaid application

The Oregon Health Insurance Experiment offered lottery-selected low-income adults an opportunity to apply for Medicaid (Finkelstein et al. 2012). Its public-use records link the invitation, application, and adjudication stages (National Bureau of Economic Research 2026). Let n_g, k_g, and r_g denote invited people, submitted applications, and approvals in

TABLE 1. Three administrative reports and the reconstruction defect

Reception group	Application	Approval given application	Approval after invitation	Log defect 95% interval
English, phone	62.5	49.8	31.0	-0.4 [-7.0, 6.3]
English, no phone	50.5	48.0	27.0	11.0 [-7.6, 29.6]
Other language, phone	63.5	40.1	21.1	-18.9 [-45.6, 7.8]
Other language, no phone	62.4	49.0	27.0	-12.3 [-95.0, 70.4]

Notes. Rates are percentages. A stable household hash assigns outreach, adjudication, and program reporting to three disjoint extracts within the eight lottery draws. The log defect is $\log d_g - \log a_g - \log b_g$. Intervals use a household bootstrap within draw and extract and one simultaneous max-t critical value across the four reception groups.

reception group g. The three stage probabilities are

$$\begin{aligned} a_g &= \Pr(\text{application} \mid \text{invitation}, g), \\ b_g &= \Pr(\text{approval} \mid \text{application}, g), \\ d_g &= \Pr(\text{approval} \mid \text{invitation}, g). \end{aligned} \tag{48}$$

A common administrative history gives $d_g = a_g b_g$. In log coordinates, $(x_g, y_g, z_g) = (\log a_g, \log b_g, \log d_g)$, its reconstruction obstruction is

$$\delta_g = z_g - x_g - y_g. \tag{49}$$

The outreach report supplies a_g, the adjudication report supplies b_g, and the program report supplies d_g. Equation (49) asks whether these locally valid reports can be read as extracts from one linked production system.

The application uses 29,799 people in 24,901 households across eight lottery draws. Reception groups combine the language requested for mailed material and phone availability at lottery signup. A stable household hash assigns the linked source records to three disjoint extracts within each draw. Each extract supplies one of the reports in (48). This split makes the reconstruction restriction an empirical comparison across records drawn from the same institutional population. Inference resamples households within draw and extract. One max-t critical value covers the four group defects jointly.[6]

The aggregate direct approval rate is 29.59 percent. The product of the separately estimated application and adjudication rates is 29.75 percent, a gap of −0.16 percentage point with a bootstrap interval of $[-1.58, 1.26]$. The joint statistic for $(\delta_1, \dots, \delta_4)$ is 5.34 on four degrees of freedom, with $p = 0.254$. Every simultaneous group interval contains zero.

[6]The split is an analysis design applied to linked public-use records. It evaluates whether separate administrative reports reconstruct the same approval process. The randomized causal interpretation in this application is reserved for the lottery effect on enrollment.

TABLE 2. Targets select different bridge records

Target	Estimate	$r_G(P)$	Bridge state
Lottery effect on enrollment	25.6 [24.9, 26.3]	0	Retained experiment
Approval after invitation, population	29.2 [28.6, 29.8]	1	Weighted application score
Approval after invitation, four groups	23.3–30.4	4 {4}	Four application scores

Notes. The evaluated archive retains the lottery–enrollment experiment and group conditional approval reports. Applicant totals and invitation–application links remain upstream. The first estimate is the lottery-selection coefficient from a linear probability model with lottery-draw and household-size fixed effects; its interval clusters by household. The other intervals resample households within lottery draw. Braces report the simultaneous 95% confidence set for the four-group record count. The first estimate is in percentage points; the remaining estimates are percentages.

Repeating the hash assignment over 100 prespecified seeds places the fifth and ninety-fifth percentiles of the aggregate gap at −1.10 and 1.26 percentage points.

The linked source also permits a target-specific archive exercise. Consider an evaluated release that carries the lottery–enrollment experiment and four conditional approval rates and keeps applicant totals and invitation–application links upstream. The minimum bridge count is zero for the lottery effect on enrollment because its experiment remains in the archive. Population approval after invitation uses one weighted application-score direction. The four group approval probabilities use four directions. Household-clustered and household-bootstrap intervals give the record-count confidence set $\{4\}$ for the group portfolio.

The target therefore chooses the bridge state. The same reporting handoff yields record counts zero, one, and four for three downstream questions. This is the portfolio rank in Corollary 6 expressed in administrative reporting units. The exercise evaluates a declared release on linked public records and records the information carried by that release.

FIGURE 1. Reconstructing approval across administrative reports

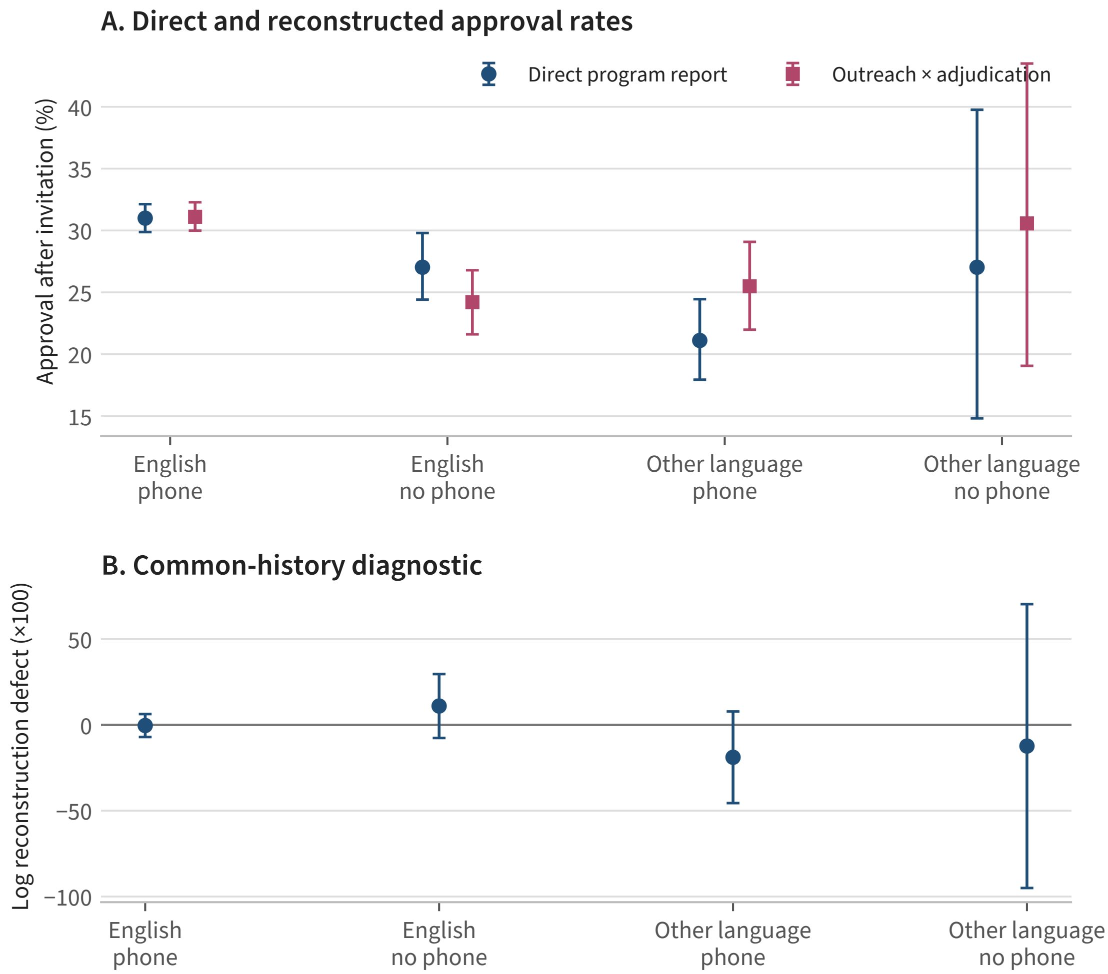

Notes. Panel A compares the program extract's approval-after-invitation rate with the product of application and conditional-approval rates from the other two extracts. Bars are pointwise household-bootstrap 95% intervals. Panel B reports $100\delta_g$ with simultaneous 95% max-t intervals. A stable hash assigns households to disjoint extracts within lottery draw.

6. Testing reconstruction and choosing links

The connecting map turns global reconstruction into a set of linear restrictions. Choose coordinates on $\operatorname{im}\partial$ and write the resulting complete diagnostic as $A^\star : \mathcal{L} \to \mathbb{R}^o$. By Theorem 1,

$$\mathcal{M} = \ker A^\star, \qquad \operatorname{rank} A^\star = o. \tag{50}$$

Let $\widehat{z} \in \mathcal{L}$ estimate a compatible collection of local response states. Suppose

$$\sqrt{n}(\widehat{z} - z_0) \Rightarrow N(0, \Sigma), \tag{51}$$

where Σ is positive definite on $\mathcal{L}$. Define $\Omega = A^\star \Sigma A^{\star\top}$.

PROPOSITION 1 (Reconstruction test). *If $z_0 \in \mathcal{M}$ and Ω is nonsingular, then*

$$n(A^\star\widehat{z})^\top \widehat{\Omega}^{-1}(A^\star\widehat{z}) \Rightarrow \chi^2_o \tag{52}$$

for every consistent estimator $\widehat{\Omega}$ of Ω. The degrees of freedom equal the dimension of a complete reconstruction diagnostic.

The statistic evaluates the same object that appears in the population theorem. Its coordinates can be changed by any nonsingular transformation without changing the test. In the three-context economy, $A^\star\widehat{z} = \widehat{z} - \widehat{x} - \widehat{y}$, so the single reconstruction obstruction yields a one-degree-of-freedom statistic. This test uses an estimated compatible atlas. When the archive reveals only $P\pi(z)$, the missing subspace $N = \ker P$ governs additional linkage, and the dictionary condition in (21) determines whether the available links can complete it.[7]

Raw local estimates may also violate within-context or overlap restrictions. Let V contain the stacked raw estimates, let $L \subseteq V$ impose local and overlap compatibility, and embed $\mathcal{M} \subseteq L$. For a positive definite weight matrix W, the nested projection identity is

$$\widehat{v} = P^W_{\mathcal{M}}\widehat{v} + (P^W_L - P^W_{\mathcal{M}})\widehat{v} + (I - P^W_L)\widehat{v}. \tag{53}$$

The last term records local or overlap disagreement. The middle term records the globally unresolved response class. Under a Gaussian limit with $W = \Sigma^{-1}$ and fixed subspaces, the two squared norms are independent chi-square variables with degrees of freedom $\dim V - \dim L$ and o.

[7] For a fixed diagnostic rank, the statistic is the usual Wald quadratic form; the moment-based asymptotic framework follows Hansen (1982). The rank-boundary result below treats the case in which the diagnostic dimension itself changes.

Rank selection adds a separate layer. When the rank of the response morphism, archive map, or target restriction changes, the obstruction and record dimensions change with it. Rank testing is irregular at such boundaries (Chen and Fang 2019), and local approaches to identification failure lead to weakly regular targets (Kaji 2021). The next result separates the discontinuous exact rank from two objects that retain uniform coverage.

PROPOSITION 2 (Inference at a rank boundary). *Let A be a $p \times q$ matrix and $m = \min\{p, q\}$.*

1. *In the Gaussian rank experiment $\widehat{A}_n = A + n^{-1/2}Z$, where $\operatorname{vec} Z \sim N(0, I)$, take a rank-$r < m$ matrix A_0 and unit vectors $u \in \ker A_0^*$ and $v \in \ker A_0$. For $h > 0$, let $A_{n,h} = A_0 + hn^{-1/2}uv^*$. Every rank estimator $\widehat{r}_n$ satisfies*

   (54) $$\Pr_{A_0}(\widehat{r}_n \neq r) + \Pr_{A_{n,h}}(\widehat{r}_n \neq r + 1) \geq 2\{1 - \Phi(h/2)\}.$$

   Hence exact rank has no uniformly consistent estimator on neighborhoods containing $n^{-1/2}$ rank changes.

2. *Suppose an estimator and a random radius satisfy*

   (55) $$\inf_{A \in \mathcal{A}_n} \Pr_A\Big(\|\widehat{A} - A\|_{\mathrm{op}} \leq c_n\Big) \geq 1 - \alpha - o(1).$$

   For $\tau \geq 0$, define the effective rank $r_\tau(A) = \#\{j : \sigma_j(A) > \tau\}$ and

   (56) $$\underline{r}_{n,\tau} = \#\{j : \sigma_j(\widehat{A}) > \tau + c_n\},$$

   (57) $$\overline{r}_{n,\tau} = \#\{j : \sigma_j(\widehat{A}) \geq \max\{\tau - c_n, 0\}\}.$$

   Then, simultaneously for every $\tau \geq 0$,

   (58) $$\inf_{A \in \mathcal{A}_n} \Pr_A\Big(\underline{r}_{n,\tau} \leq r_\tau(A) \leq \overline{r}_{n,\tau} \textit{ for every } \tau \geq 0\Big) \geq 1 - \alpha - o(1).$$

3. *Let θ collect the finite-dimensional primitives that determine an archive fiber and its sharp target set $\mathcal{I}_G(\theta)$, and let C_n be a confidence region satisfying $\inf_{\theta \in \Theta_n} \Pr_\theta(\theta \in C_n) \geq 1 - \alpha - o(1)$. The projection set*

   (59) $$\mathcal{C}_{G,n} = \overline{\bigcup_{\vartheta \in C_n} \mathcal{I}_G(\vartheta)}$$

   covers the true sharp target set uniformly over Θ_n. If Θ_{rec} denotes the primitives that admit global reconstruction, the test that rejects when $C_n \cap \Theta_{\mathrm{rec}} = \varnothing$ has uniform asymptotic size at most α.

Part 1 locates the discontinuity. At $\tau = 0$, the upper endpoint in (56) may include all m directions, reflecting the information available near a zero singular value. Positive τ yields an honest statement about directions whose response exceeds a declared economic resolution. The projection set carries every rank admitted by the primitive confidence region into the final counterfactual. It therefore supplies the inference layer for the identified set in (23) while allowing the archive and target ranks to vary.

COROLLARY 7 (Gaussian critical values and projection). *Suppose first that* $\widehat{A} = A + \sigma n^{-1/2} Z$, *where* $\operatorname{vec} Z \sim N(0, I_{pq})$ *and* σ *is known. Let* $q_{p,q}(1-\alpha)$ *be the* $(1-\alpha)$ *quantile of* $\|Z\|_{\mathrm{op}}$ *and set*

$$c_{n,\alpha} = \frac{\sigma q_{p,q}(1-\alpha)}{\sqrt{n}}. \tag{60}$$

The brackets in (56) *then cover* $r_\tau(A)$ *simultaneously for every* $\tau \geq 0$ *with probability at least* $1-\alpha$ *in finite samples. The quantile depends only on* (p, q, α) *and can be simulated before observing* $\widehat{A}$.

Suppose next that $\widehat{\theta} \sim N_d(\theta, V/n)$ *with known positive definite* V. *With* $\chi^2_{d,1-\alpha}$ *denoting the corresponding chi-square quantile, define*

$$C_{n,\alpha} = \left\{ \vartheta : n(\widehat{\theta} - \vartheta)^\top V^{-1} (\widehat{\theta} - \vartheta) \leq \chi^2_{d,1-\alpha} \right\}. \tag{61}$$

Using $C_{n,\alpha}$ *in* (59) *gives finite-sample coverage at least* $1-\alpha$ *for the sharp target set and a reconstruction test with size at most* α. *If* V *is estimated and the studentized quadratic form has a uniform chi-square limit, the same conclusions hold asymptotically.*

This implementation leaves the discontinuous rank inside the confidence region. The reported output can include the effective-rank bracket, the union of target sets over the ellipsoid, and the reconstruction decision from the same primitive region. Monte Carlo evaluation in the replication files uses independent draws for the operator-norm quantile and the coverage calculation.

7. Regular extensions

Smooth local reconstruction. Let $\mathcal{L}$ be a manifold of compatible locally minimal response states near z, and let $Q : H \to \mathcal{L}$ map a common history into its assembled response. Suppose Q has constant rank near h with $Q(h) = z$. Let the derivatives of the local restriction maps form a finite tangent complex $\mathcal{C}_h$ whose zeroth cohomology is $T_z\mathcal{L}$. Define the tangent obstruction dimension

$$o_h = \dim T_z\mathcal{L} - \operatorname{rank} DQ_h. \tag{62}$$

PROPOSITION 3 (Smooth reconstruction at a regular history). *Under the stated constant-rank and tangent-complex conditions, the smallest local state dimension through which the response germ Q factors is*

$$d_h = \operatorname{rank} DQ_h = \chi(\mathcal{C}_h) + \eta(\mathcal{C}_h) - o_h. \tag{63}$$

If $Q(H)$ is locally the zero set of a smooth constant-rank representative of the connecting map, then its derivative has kernel $\operatorname{im} DQ_h$ *and rank o_h. The derivative therefore gives a locally minimal reconstruction diagnostic.*

The constant-rank conditions supply coordinates in which the response uses d_h state coordinates and the obstruction uses o_h transverse coordinates. At a rank change, one chart ceases to describe the state dimension. Nonlinear realization theory studies related regularity and uniqueness questions (Sussmann 1976; Gauthier and Bornard 1982).

Global regular reconstruction. A local rank calculation supplies a state count around one history. A global state requires the response equivalence classes to form one smooth quotient. Let R be a connected history manifold and let $T : R \to Y$ collect the declared future responses. Define

$$E_T = \{(r, r') \in R \times R : T(r) = T(r')\}. \tag{64}$$

A regular global reconstruction is a Hausdorff smooth manifold Z, a surjective submersion $\pi : R \to Z$, and an injective immersion $g : Z \to Y$ satisfying $T = g\pi$.

PROPOSITION 4 (Global regular reconstruction). *Suppose T has constant rank. A regular global reconstruction exists exactly when E_T is a closed embedded submanifold of $R \times R$ and the first projection $E_T \to R$ is a submersion. When it exists, the quotient*

$$Z_T = R/E_T \tag{65}$$

is a Hausdorff smooth manifold, supplies a minimal smooth realization of T, and is unique up to diffeomorphism among regular global reconstructions. Its dimension is

$$\dim Z_T = \operatorname{rank} DT = \dim R - \dim \ker(DT|_{TR}). \tag{66}$$

The condition on E_T is the quotient-manifold criterion of Godement (2017). Singular fiber dimensions can break the submersion condition. Nonclosed equivalence relations can break Hausdorff separation. The theorem locates both boundaries in the response equivalence relation.[8]

Stochastic behavior. When local responses are distributions or conditional laws, characteristic functions, score operators, or decision losses can form the response sheaf. Equality of means gives one image sheaf and equality of laws gives another. The descent sequence then records which locally coherent distributional responses admit one global stochastic history.

8. Conclusion

Fragmented evidence can supply a minimal response state in every policy context and a compatible state on every overlap. A common economic history imposes a global lifting requirement. The counterfactual reconstruction theorem places that requirement in the connecting map ∂. Its kernel is the globally generated economy. Its image is the reconstruction obstruction. Its rank is the exact state-dimension gap and the dimension of a complete linear diagnostic.

The theorem therefore turns the reconstruction problem into one object. It determines whether the local responses form one economy, constructs the minimal economy in the feasible cases, measures the missing information in the remaining cases, and classifies every valid linear diagnostic. After the baseline archive is declared, its unobserved obstruction subspace determines the acquisition task. Each downstream counterfactual selects a canonical quotient inside that subspace, which gives its sharp identified set, its exact record count, the records that preserve every finite decision problem based on the target, and the exact amplification of bounded record error. A portfolio of counterfactuals yields a polymatroid of record requirements; its intersections measure which acquired coordinates can serve several decisions. Full reconstruction reduces to a spanning design. The residual target spectrum prices smaller budgets under quadratic loss, shows how normalized measurement noise discounts each acquired direction, and allocates continuous measurement precision by its shadow value. Singular-value brackets and Gaussian primitive projection sets provide inference across rank changes. The Oregon application shows how the obstruction becomes a comparison among reports produced at successive administrative stages and how the declared target changes the bridge count from zero to one to four. The finite atlas gives a computable cochain representation. The smooth and global results describe the regular settings in which the same state survives beyond linear coordinates.

[8] Closedness gives the separation property needed for a Hausdorff quotient, while the embedded-submanifold and submersion conditions give compatible smooth charts on the equivalence classes and their quotient.

References

Abramsky, Samson and Adam Brandenburger. 2011. "The Sheaf-Theoretic Structure of Non-Locality and Contextuality." *New Journal of Physics* 13 (11): 113036. 10.1088/1367-2630/13/11/113036.

Alexanderian, Alen, Noemi Petra, Georg Stadler, and Omar Ghattas. 2016. "A Fast and Scalable Method for A-Optimal Design of Experiments for Infinite-Dimensional Bayesian Nonlinear Inverse Problems." *SIAM Journal on Scientific Computing* 38 (1): A243–A272. 10.1137/140992564.

Attia, Ahmed, Alen Alexanderian, and Arvind K. Saibaba. 2018. "Goal-Oriented Optimal Design of Experiments for Large-Scale Bayesian Linear Inverse Problems." *Inverse Problems* 34 (9): 095009. 10.1088/1361-6420/aad210.

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

Bredon, Glen E. 1997. *Sheaf Theory*, vol. 170 of *Graduate Texts in Mathematics*. 2 ed. New York: Springer. 10.1007/978-1-4612-0647-7.

Chen, Qihui and Zheng Fang. 2019. "Improved Inference on the Rank of a Matrix." *Quantitative Economics* 10 (4): 1787–1824. 10.3982/QE1139.

Eckart, Carl and Gale Young. 1936. "The Approximation of One Matrix by Another of Lower Rank." *Psychometrika* 1 (3): 211–218. 10.1007/BF02288367.

Finkelstein, Amy, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P. Newhouse, Heidi Allen, Katherine Baicker, and Oregon Health Study Group. 2012. "The Oregon Health Insurance Experiment: Evidence from the First Year." *The Quarterly Journal of Economics* 127 (3): 1057–1106. 10.1093/qje/qjs020.

Gauthier, J. P. and G. Bornard. 1982. "Existence and Uniqueness of Minimal Realizations in the C^∞ Case." *Systems & Control Letters* 1 (6): 395–398. 10.1016/S0167-6911(82)80053-4.

Godement, Roger. 2017. *Introduction to the Theory of Lie Groups*. Universitext. Cham: Springer. 10.1007/978-3-319-54375-8.

Gorman, W. M. 1968. "The Structure of Utility Functions." *The Review of Economic Studies* 35 (4): 367–390. 10.2307/2296766.

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

Ho, B. L. and R. E. Kalman. 1966. "Effective Construction of Linear State-Variable Models from Input/Output Functions." *Regelungstechnik* 14: 545–548. 10.1524/auto.1966.14.112.545.

Kaji, Tetsuya. 2021. "Theory of Weak Identification in Semiparametric Models." *Econometrica* 89 (2): 733–763. 10.3982/ECTA16413.

Montanari, Arthur N., Chao Duan, Luis A. Aguirre, and Adilson E. Motter. 2022. "Functional Observability and Target State Estimation in Large-Scale Networks." *Proceedings of the National Academy of Sciences* 119 (1): e2113750119. 10.1073/pnas.2113750119.

National Bureau of Economic Research. 2026. "The Oregon Health Insurance Experiment: Public Use Data." 10.60592/rx0z-sp89.

Robinson, Michael. 2017. "Sheaves Are the Canonical Data Structure for Sensor Integration." *Information Fusion* 36: 208–224. 10.1016/j.inffus.2016.12.002.

Sussmann, Hector J. 1976. “Existence and Uniqueness of Minimal Realizations of Nonlinear Systems.” *Mathematical Systems Theory* 10 (1): 263–284. `10.1007/BF01683278`.
Willems, Jan C. 1986. “From Time Series to Linear System—Part II: Exact Modelling.” *Automatica* 22 (6): 675–694. `10.1016/0005-1098(86)90005-1`.
Willems, Jan C. 2007. “The Behavioral Approach to Open and Interconnected Systems.” *IEEE Control Systems Magazine* 27 (6): 46–99. `10.1109/MCS.2007.906923`.
Wu, Rui, Hong Xie, and Yongjun Li. 2026. “Cohomological Obstructions to Global Counterfactuals: A Sheaf-Theoretic Foundation for Generative Causal Models.” `https://arxiv.org/abs/2603.17384`.
Zhang, Yuan, Tyrone Fernando, and Mohamed Darouach. 2024. “Functional Observability, Structural Functional Observability, and Optimal Sensor Placement.” *IEEE Transactions on Automatic Control*. `10.1109/TAC.2024.3462558`.

Appendix A. Proofs

Proof of Lemma 1. Let $\tau = \Gamma(T)$. Every factorization $\tau = gq$ satisfies $\operatorname{rank}\tau \leq \dim \operatorname{im} q$. The factorization through $\mathcal{M} = \operatorname{im}\tau$ attains equality. For a general factorization, define $\pi : \operatorname{im} q \to \mathcal{M}$ by $\pi(qh) = \tau h$. If $qh = qh'$, then $\tau h = gqh = gqh' = \tau h'$, so π is well defined. It is linear, surjective, response preserving, and fixed on every reachable element, which gives uniqueness. It is an isomorphism exactly when the reachable state has dimension $\operatorname{rank}\tau$. □

Proof of Theorem 1. The kernel and sheaf image of T give the short exact sequence (10). Applying the derived global-section functors gives the long exact cohomology sequence. Its first terms are (16). Exactness at $\mathcal{L}$ and $H^1(\mathcal{K})$ gives

$$\ker\partial = \operatorname{im}\Gamma(T) = \mathcal{M}, \qquad \operatorname{im}\partial = \ker\Big[H^1(\mathcal{K}) \to H^1(\mathcal{H})\Big]. \tag{A1}$$

The first isomorphism theorem applied to ∂ yields (17). It also gives $z \in \mathcal{M}$ exactly when $\partial z = 0$.

Lemma 1 identifies $\mathcal{M}$ as the minimal global state and supplies its universal property. The equality $\mathcal{L} = \mathcal{M}$ is equivalent to $\mathfrak{O} = 0$. The obstruction is zero exactly when $\operatorname{im}\partial = 0$, which is equivalent to $\partial = 0$ and, by exactness, to injectivity of $H^1(\mathcal{K}) \to H^1(\mathcal{H})$. These equivalences prove part 4.

Rank–nullity for ∂ gives $\dim\mathcal{M} = \dim\mathcal{L} - o$. Euler–Poincaré and (11) give $\dim\mathcal{L} = \chi(\mathcal{Z}) + \eta(\mathcal{Z})$, proving (18).

Let $\pi : \mathcal{L} \to \mathfrak{O}$ be the quotient map. A linear map $A : \mathcal{L} \to V$ vanishes on $\mathcal{M}$ exactly when the universal property of the quotient supplies a unique map $\bar{A} : \mathfrak{O} \to V$ with $A = \bar{A}\pi$. Exactness makes the induced map $\bar{\partial} : \mathfrak{O} \to \operatorname{im}\partial$ an isomorphism. Substituting $\pi = \bar{\partial}^{-1}\partial$ gives the two natural bijections in (19). A valid diagnostic is complete exactly when $\bar{A}$ is injective. Hence a complete $A : \mathcal{L} \to \mathbb{R}^m$ has $m \geq o$. Choosing a coordinate isomorphism $J : \operatorname{im} \partial \to \mathbb{R}^o$ gives $A^\star = J\partial$ with kernel $\mathcal{M}$, proving (20).

Fix a baseline map $P : \mathfrak{O} \to D$, a selected set S, and write $R_S = (P, A_S)$. The implication from identification to the kernel inclusion follows by applying identification to u and $u + k$ for $k \in \ker R_S$. Conversely, the kernel inclusion makes the rule $C_S(R_S u) = Gu$ well defined on $\operatorname{im} R_S$; extend it linearly to the codomain of R_S. This proves the equivalence of the first three statements. Since $\ker R_S = \ker(A_S|_N)$, finite-dimensional duality gives

$$\ker(A_S|_N) \subseteq \ker(G|_N) \quad \Longleftrightarrow \quad \operatorname{im}\big((G|_N)^*\big) \subseteq \operatorname{im}\big((A_S|_N)^*\big)\,. \tag{A2}$$

The last image is $\operatorname{span}\{a_j|_N : j \in S\}$, proving (21). The fiber of R_S through u is $u + \ker R_S$, whose image under G is (23).

If $G = C_S R_S$, a decision rule based on GU can be implemented from $R_S U$ by composition with C_S. Observing GU already reveals the payoff-relevant state, so the two observations give the same minimum risk for every μ, A, and ℓ. Conversely, suppose identification fails. There are $u_0, u_1 \in \mathfrak{O}$ with $R_S u_0 = R_S u_1$ and $Gu_0 \neq Gu_1$. Place probability one half on each point, take two actions, and assign loss zero when action i is chosen at target Gu_i and loss one for the other action. Observing G gives risk zero. The selected archive gives the same record at both points, so its minimum risk is one half. Thus (22) is equivalent to the first four conditions.

Every identifying design has $\operatorname{rank}(A_S|_N) \geq \operatorname{rank}(G|_N)$, so it uses at least $r_G(P)$ scalar records. Choose a basis of $\operatorname{im}((G|_N)^*) \subseteq N^*$ and extend its elements from N to $\mathfrak{O}$. The common kernel of these $r_G(P)$ measurements on N is $\ker(G|_N)$, so they identify G and prove (24). The first isomorphism theorem gives (25). For any identifying augmentation $A : N \to E$, define $\varphi_A : \operatorname{im} A \to \mathcal{Q}_{G,P}$ by $\varphi_A(An) = [n]$. The kernel inclusion makes this map well defined, and every class has a preimage. Its action on each reachable record fixes it uniquely. The map $[n] \mapsto Gn$ embeds $\mathcal{Q}_{G,P}$ into W, so φ_A preserves the target response. This proves the minimal-state claim. For $G = I_{\mathfrak{O}}$, identification requires the selected restrictions to span N^*. A spanning family contains a basis of $q = \dim N$ elements, while fewer than q functionals have a nonzero common kernel. With positive costs, every exact design contains a basis whose cost is weakly smaller. The independent subsets of $\{a_j|_N\}$ form a vector matroid, and the greedy basis theorem shows that sorting by increasing cost and adding a functional whenever it raises the span produces a minimum-weight basis.

For the deterministic error claim, suppose first that $\ker R_S \nsubseteq \ker G$. Choose $k \in \ker R_S$ with $Gk \neq 0$. The states tk and $-tk$ yield the same record at zero error and targets separated by $2t\|Gk\|$, so the minimax error diverges with t. Now suppose $\ker R_S \subseteq \ker G$. The map $C_0(R_S u) = Gu$ is well defined on $\operatorname{im} R_S$, and its zero extension to $(\operatorname{im} R_S)^\perp$ is $C_0 = GR_S^\dagger$. The decoder $C_0 y$ has error $\|C_0 e\| \leq \varepsilon \|C_0\|_{\mathrm{op}}$. For the reverse inequality, choose a unit $v \in \operatorname{im} R_S$ approaching a maximizing right singular vector of C_0. At the common observation zero, the two admissible state–error pairs $u_\pm = \pm\varepsilon R_S^\dagger v$ and $e_\pm = \mp\varepsilon v$ have targets $\pm\varepsilon C_0 v$. Every decoder incurs error at least $\varepsilon\|C_0 v\|$ at one of these pairs. Taking the supremum over v proves (26).

For the noiseless stochastic frontier, let $\widehat{U}_P$ be the orthogonal projection of the random vector U in L^2 onto the finite-dimensional space of linear functions of PU, and set $\widetilde{U}_P = U - \widehat{U}_P$. Linearity of orthogonal projection gives $\widetilde{V}_{G,P} = G\widetilde{U}_P$. The random vector $P\widetilde{U}_P$ is both linear in PU and orthogonal to every linear function of PU, so it is zero almost surely. For any additional encoder B, the component $B\widehat{U}_P$ is already a linear function of the baseline record. The incremental record is therefore $B\widetilde{U}_P$. Every encoder–decoder pair gives a rank-at-most-m linear approximation $C_1 B\widetilde{U}_P$ to $G\widetilde{U}_P$, and every rank-at-most-m linear approximation of this form admits such a factorization. The principal-component approximation theorem (Eckart and Young 1936) gives the covariance eigenvalue tail as a lower bound. It is attained by taking the records to be the first m principal coordinates of $G\widetilde{U}_P$, each of which is a linear functional of U after a baseline-measurable term is removed. This proves (27).

For the noisy-record claim, set $X = \widetilde{U}_P$ and $\Sigma_P = \mathrm{Cov}(X)$. Conditional linear projection on the baseline removes $B\widehat{U}_P$ from Y_B, leaving $BX + \varepsilon$. The signal normalization and the orthogonality of ε give

$$\text{(A3)} \qquad \mathrm{Cov}(BX+\varepsilon) = (1+\sigma^2)I_m, \qquad \mathrm{Cov}(GX, BX+\varepsilon) = G\Sigma_P B^*.$$

The best linear decoder therefore has risk

$$\text{(A4)} \qquad \mathrm{tr}(G\Sigma_P G^*) - \frac{1}{1+\sigma^2}\,\mathrm{tr}(G\Sigma_P B^* B\Sigma_P G^*).$$

Choose coordinates on $\mathfrak{D}$ and write $C = B\Sigma_P^{1/2}$. The normalization becomes $CC^* = I_m$, while the second trace in (A4) is

$$\text{(A5)} \qquad \mathrm{tr}\Big(C\Sigma_P^{1/2}G^*G\Sigma_P^{1/2}C^*\Big).$$

Ky Fan's maximum principle makes its largest value the sum of the first m eigenvalues of $\Sigma_P^{1/2}G^*G\Sigma_P^{1/2}$. Its positive eigenvalues equal those of $G\Sigma_P G^* = \mathrm{Cov}(\widetilde{V}_{G,P})$. Rows along the leading eigenvectors attain the maximum; directions in range Σ_P complete the normalized record when zero eigenvalues are selected. Substitution into (A4) proves (29). □

Proof of Corollary 1. If $H^1(\mathcal{H}) = 0$, then the map from $H^1(\mathcal{K})$ in (16) has zero codomain. Equation (17) therefore gives $\mathfrak{D} \cong H^1(\mathcal{K})$. The remaining claims follow from parts 4 and 7 of Theorem 1. □

Proof of Corollary 2. The decision rule is a linear decoder of the baseline and additional records, so its minimum expected loss is the optimization in (27). With zero added records the loss is the sum of all eigenvalues of the residual target covariance. Subtracting the loss at budget m gives $\sum_{j=1}^{m} \lambda_j(G, P)$, and the increment from $m - 1$ to m is $\lambda_m(G, P)$. Under normalized noisy records, subtracting (29) from the zero-record loss gives the first expression in (33). Taking the increment from $m - 1$ to m gives the second. □

Proof of Corollary 3. Residualize on the baseline archive as in the proof of Theorem 1. For fixed precisions, the best-linear risk reduction contributed by the normalized directions is a weighted sum of Rayleigh quotients with weights $\tau_j/(1 + \tau_j)$. The weighted Ky Fan inequality pairs the largest weight with the largest eigenvalue of the residual target covariance. Relabeling the precisions therefore gives the objective in (35). The objective is convex in τ and decreases along every coordinate with positive λ_j. The Karush–Kuhn–Tucker conditions for an active coordinate give

$$\frac{\lambda_j(G, P)}{(1 + \tau_j)^2} = \nu. \tag{A6}$$

Complementary slackness gives (36); its left side is continuous and strictly decreasing in ν on the relevant interval, so the budget determines $\nu(b)$. The envelope theorem identifies $\nu(b)$ as the reduction in optimized risk from one additional unit of precision. □

Proof of Corollary 4. The factorization $P_1 = HP_2$ gives $\ker P_2 \subseteq \ker P_1$. Applying G gives the identified-set inclusion, and restricting G to the two nested kernels gives $\operatorname{rank}(G|_{\ker P_2}) \le \operatorname{rank}(G|_{\ker P_1})$. Every linear decoder that uses $P_1 U$ and m added records can be implemented from $P_2 U$ and the same records by composing its baseline coefficient with H. The feasible decision rules under P_1 are therefore contained in those under P_2, which gives the risk inequality. For an arbitrary added record A, every rule $\delta(P_1 U, AU)$ can be implemented as $\delta(HP_2 U, AU)$. Taking the infimum over rules gives (38). □

Proof of Corollary 5. Let P restrict the larger response profile to the smaller future family. The global response operators satisfy $\Gamma(T_1) = P\Gamma(T_2)$. Hence $P(\mathcal{M}_2) = \mathcal{M}_1$. The restriction of P to $\mathcal{M}_2$ is surjective, and the dimension inequality follows. □

Proof of Corollary 6. The adjoint of the stacked restriction has image

$$\operatorname{im}((G_S|_N)^*) = \sum_{i \in S} \operatorname{im}((G_i|_N)^*). \tag{A7}$$

The record-count formula in Theorem 1 therefore gives (39). Normalization and monotonicity follow from the empty sum and inclusion of row spaces. Write $U_S = \operatorname{im}((G_S|_N)^*)$.

The dimension identity

$$\dim U_S + \dim U_T = \dim(U_S + U_T) + \dim(U_S \cap U_T) \tag{A8}$$

and the inclusion $U_{S\cap T} \subseteq U_S \cap U_T$ give (40). Finally, $U_{S\cup\{i\}} = U_S + \operatorname{im}((G_i|_N)^*)$. Applying the same dimension identity gives (41). □

THREE-CONTEXT ECONOMY. The three rows of D have disjoint pivot pairs, so $\operatorname{rank} D = 3$ and $\dim \mathcal{L} = \dim\ker D = 3$. Direct multiplication using (44) gives $DQ = 0$. The first two rows of Q form the identity, so $\operatorname{rank} Q = 2$. In the coordinates (x, y, z) on $\mathcal{L}$, define $\ell(x, y, z) = z - x - y$. Then $\ker \ell = \operatorname{im} Q$ and $\operatorname{rank} \ell = 1$. The short exact sequence of cochain complexes

$$0 \longrightarrow (0 \to \mathbb{R}) \longrightarrow (\mathbb{R}^3 \xrightarrow{\ell} \mathbb{R}) \longrightarrow (\mathbb{R}^3 \to 0) \longrightarrow 0 \tag{A9}$$

has connecting map ℓ. Thus $\partial = \ell$, $\mathfrak{O} \cong \mathbb{R}$, and $d^\star = 2$. The three local coordinate maps have matrices I_2, $\left(\begin{smallmatrix}0&1\\1&1\end{smallmatrix}\right)$, and $\left(\begin{smallmatrix}1&1\\1&0\end{smallmatrix}\right)$, each with nonzero determinant. □

PROOF OF PROPOSITION 1. Under the null, $A^\star z_0 = 0$. The continuous mapping theorem gives $\sqrt{n}A^\star\widehat{z} \Rightarrow N(0, \Omega)$. Studentization by a consistent $\widehat{\Omega}$ yields (52). The matrix $A^\star$ has rank o by Theorem 1, which gives the degrees of freedom. For the nested projection statement, whiten the Gaussian limit by $\Sigma^{-1/2}$. The ranges of $I - P_L^W$ and $P_L^W - P_{\mathcal{M}}^W$ are orthogonal, so their Gaussian components are independent and their squared norms have the stated chi-square dimensions. □

PROOF OF PROPOSITION 2. Under A_0 and $A_{n,h}$, the vectorized Gaussian observation has equal covariance and standardized means separated by the vector $h\operatorname{vec}(uv^*)$, whose Euclidean norm is h. The sum of errors of any test between two simple hypotheses is bounded below by one minus their total variation distance. Equal-covariance Gaussian measures separated by norm h have total variation distance $2\Phi(h/2) - 1$. A rank estimator induces such a test because A_0 and $A_{n,h}$ have ranks r and $r + 1$. This gives (54).

On the event in (55), Weyl's inequality gives $|\sigma_j(\widehat{A}) - \sigma_j(A)| \le c_n$ for every singular value. If $\sigma_j(\widehat{A}) > \tau + c_n$, then $\sigma_j(A) > \tau$. If $\sigma_j(A) > \tau$, then $\sigma_j(\widehat{A}) > \tau - c_n$; the weak inequality and the zero truncation in the upper count also cover $\tau \le c_n$. The two implications hold for every j and τ on the same event, proving (58).

If the true primitive θ belongs to C_n, then $\mathcal{I}_G(\theta)$ is one of the sets in the union defining (59). Hence the projected set contains the true sharp target set whenever the primitive confidence region covers. Under global reconstruction, rejection through $C_n \cap \Theta_{\mathrm{rec}} = \varnothing$ implies that the true primitive lies outside C_n. The two uniform claims follow from the coverage assumption on C_n. □

PROOF OF COROLLARY 7. The random variable $\sqrt{n}\|\widehat{A}-A\|_{\mathrm{op}}/\sigma$ has the same distribution as $\|Z\|_{\mathrm{op}}$. Hence the event in (55) with radius (60) has probability $1-\alpha$. Part 2 of Proposition 2 gives the simultaneous bracket on this event. For the primitive vector,

$$n(\widehat{\theta}-\theta)^{\top}V^{-1}(\widehat{\theta}-\theta) \sim \chi^2_d, \tag{A10}$$

so $C_{n,\alpha}$ covers θ with probability $1-\alpha$. Part 3 of the proposition gives coverage of the sharp target set and the size bound. Replacing V by a uniformly consistent estimator and applying the assumed uniform chi-square approximation gives the asymptotic statements. □

PROOF OF PROPOSITION 3. The tangent complex gives $\dim T_z\mathcal{L} = \chi(\mathcal{C}_h) + \eta(\mathcal{C}_h)$. Rank–nullity and the definition of o_h yield

$$\operatorname{rank} DQ_h = \dim T_z\mathcal{L} - o_h = \chi(\mathcal{C}_h) + \eta(\mathcal{C}_h) - o_h. \tag{A11}$$

The constant-rank theorem supplies local coordinates in which Q uses $\operatorname{rank} DQ_h$ response coordinates, so a local realization of that dimension exists. Every smooth factorization through a state manifold Z satisfies $\operatorname{rank} DQ_h \le \dim Z$, proving minimality. If $Q(H)$ is locally the zero set of a constant-rank connecting map, the constant-rank level-set theorem identifies its tangent space with the kernel of the derivative. That tangent space is $\operatorname{im} DQ_h$, and rank–nullity gives derivative rank o_h. □

PROOF OF PROPOSITION 4. Suppose first that E_T is closed and embedded and that its first projection is a submersion. The quotient-manifold criterion gives a unique Hausdorff smooth structure on $Z_T = R/E_T$ for which the quotient map $\pi : R \to Z_T$ is a surjective submersion. Since T is constant on the equivalence classes, it induces a unique smooth map $g : Z_T \to Y$ with $T = g\pi$. Its set map is injective by the definition of E_T. Constant rank gives $\ker DT_r = T_r\pi^{-1}(\pi(r)) = \ker D\pi_r$. The identity $DT_r = Dg_{\pi(r)}D\pi_r$ and surjectivity of $D\pi_r$ make $Dg_{\pi(r)}$ injective, so g is an injective immersion.

Every smooth factorization of T through a manifold Z' satisfies $\operatorname{rank} DT \le \dim Z'$. The quotient construction has $\dim Z_T = \operatorname{rank} DT$, proving minimality and (66). If (Z', π', g') is regular, injectivity of g' makes the fibers of π' equal the classes in E_T. The quotient maps induce mutually inverse smooth maps between Z_T and Z', giving the diffeomorphism.

Conversely, suppose a regular global reconstruction exists. Injectivity of g gives $E_T = R\times_Z R$, the kernel pair of the submersion π. This fiber product is an embedded submanifold, its first projection is a submersion, and it is closed because Z is Hausdorff. □