The Polar Express

Dynamic Aggregation under Policy Holonomy

Chae-Yeon Xon
Yonsei University

September 2026

Abstract

A policy cycle can restore an aggregate report while changing the response to a later intervention. This paper studies that failure of dynamic aggregation. A feasible report path is lifted through a specified policy technology; the terminal change in a common future probe is target holonomy. Opposite loop orientations identify its local curvature, and the resulting response map describes which hidden directions require additional measurements. A phase-conditional Gorman representation gives local closure under an additive differential-frame condition. A 72-state transition economy illustrates the mechanism: a closed output–consumption loop changes a subsequent consumption response by up to 0.00714 percentage point, or 6.04 percent of its ordinary value. Public delivery records provide a separate predictive-closure exercise. Randomized loops and common terminal probes are the observation design for identifying the holonomy contrast.

JEL Codes: C14, C32, C61, C93, D12, E47, L91

Keywords: dynamic aggregation, holonomy, policy experiments, inverse design, Gorman aggregation, heterogeneous agents, delivery operations

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

1. Introduction

A policy cycle can return output and consumption to their initial levels while changing how households respond to a later intervention. During the cycle, saving and income transitions change the distribution of assets across households. Two terminal distributions can therefore share the same aggregate report and imply different responses to a common policy probe.[1] The economic question is which additional records make the report adequate for that later response. The comparison holds the future intervention and its measurement horizon fixed, so the policy cycle changes the state confronted by the probe. Delivery networks provide a second setting: restoring throughput and backlog can leave parcel ages and courier loads altered.

The paper studies the future-response defect left by a closed aggregate path, conditional on the policy technology and the future probe. Let $q : \mathcal{S} \to \mathcal{X}$ map the economic state into a report and let $\tau(s)$ denote the conditional mean response to a standardized future probe.[2] A feasible closed report path γ based at $q(s)$ is lifted through the policy technology and implementation rule. If the lifted path ends at $\mathrm{Hol}_\gamma(s)$, define the target defect

$$h_\tau(\gamma, s) = \tau\big(\mathrm{Hol}_\gamma(s)\big) - \tau(s). \tag{1}$$

A report is response-sufficient for the chosen policy loops and probe when these defects vanish. The object is therefore relative to four primitives: the report, the feasible policy technology, the implementation rule that lifts report motion into state motion, and the future target.

The local result uses small closed loops. Traversing the same rectangle in opposite directions isolates the second-order state displacement generated by the policy lift. If v and w are attainable report directions and $\Omega_s(v, w)$ is the vertical curvature of the lift, then

$$\frac{\tau(\mathrm{Hol}_{\gamma_\varepsilon^+}(s)) - \tau(\mathrm{Hol}_{\gamma_\varepsilon^-}(s))}{2\varepsilon^2} \longrightarrow D\tau_s \Omega_s(v, w). \tag{2}$$

Mirrored loop arms remove the cubic term.[3] The resulting target-curvature map has two uses. Its row space describes the response defects that local combinations of feasible loops can offset. Its pullback describes latent distinctions that a report must retain for the specified probe. Each use is conditional on the policy lift and the admissible loops. This

[1] Closure refers to the reported aggregate path. The transported micro state can finish at a different point of its report fiber.

[2] The probe is held fixed across terminal states, including its horizon, assignment rule, and outcome normalization.

[3] The design compares opposite loop orientations with matched arm lengths and a common initial report.

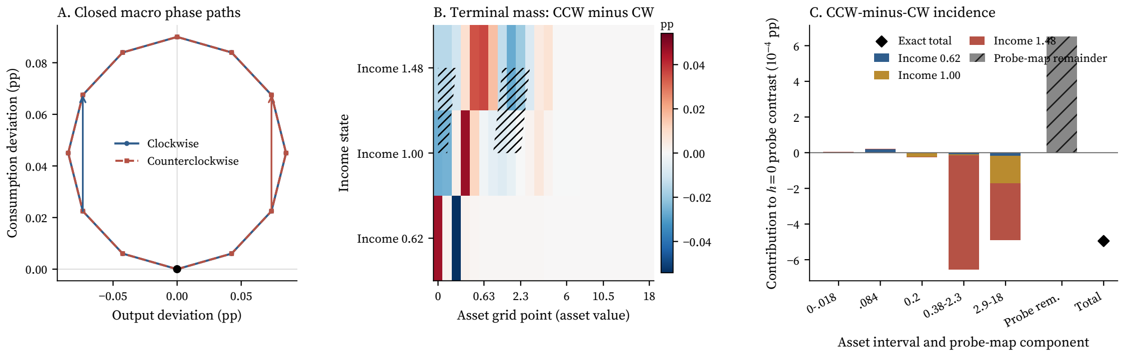

FIGURE 1. A closed report path, transported state, and target defect

Panel A plots the clockwise and counterclockwise output–consumption paths in the 72-state transition economy. Both begin and end at the black report point. Panel B gives the percentage-point difference between the two terminal asset–income distributions; the color scale is symmetric around zero, and hatching marks negative mass differences. Panel C decomposes the horizon-0 consumption-probe contrast into asset–income contributions and the probe-map remainder. The underlying transition paths, terminal distributions, and probe weights generate every plotted coordinate.

gives the economic comparison both a measurement requirement and an experimental design.

Figure 1 shows the mechanism in the transition economy used below. Clockwise and counterclockwise policy paths trace the same closed output–consumption loop. Their terminal asset–income distributions differ. A common future consumption probe converts that hidden distributional difference into a response contrast.

The paper then asks what an aggregate must retain so that these defects disappear. The dynamic aggregation theorem gives a necessary and sufficient first-order condition: the measurement differentials must span every target covector transported around the declared policy loops. The dimension of that span is a lower bound for every differentiable report, and additive moments attain it exactly when the span admits an additive differential frame. Enlarging the target, shock, or loop library can only raise this minimum dimension. The resulting quotient is target- and path-specific; a representation adequate for one family of future probes can omit directions needed by another.

Two exercises clarify the empirical content. In a 72-state transition economy, the full state and policy lift are observed. A closed output–consumption loop changes a subsequent consumption response by as much as 0.00714 percentage point, or 6.04 percent of its ordinary value.[4] In public LaDe delivery records, six additive parcel, courier, and delivery-area moments reduce next-day backlog RMSE by 6.13 percent after conditioning on dashboard history and accepted flow.[5] Policy-holonomy identification would addition-

[4]The percentage uses the ordinary horizon-zero response of 0.11826 percentage point at the same initial state as its denominator.

[5]The delivery records come from observed operations and supply the predictive-closure comparison.

ally require randomized report loops and common terminal probes. The paper states the corresponding four-arm experiment and gives its bias–variance tradeoff.

The analysis connects classical aggregation to dynamic state reduction. Exact aggregation results derive additive micro-unit subaggregates under conditions that make an economy-level relation valid across admissible micro allocations (Gorman 1968). Firm heterogeneity can also turn temporary component shocks into persistent and volatile aggregate dynamics (Abadir and Talmain 2002). State-space pruning asks which state variables are needed for a useful approximation, while sufficient-statistic approaches ask which measured objects carry a particular response. Dynamic allocation of attention across correlated sources provides a target-indexed information benchmark (Liang et al. 2022). The present problem adds policy history to the measurement requirement: the micro state is transported along feasible intervention paths before the future target is evaluated. Experimentation in endogenous organizations provides a complementary economic setting in which policy history changes the organization that later decisions confront.

The geometry is related to control-theoretic accessibility and connections on state bundles (Hermann and Krener 1977; Kobayashi and Nomizu 1963; Kelly and Murray 1995). Recent work on low-dimensional response structure and state compression provides complementary ways to summarize high-dimensional dynamics (Montanari et al. 2022; Wu et al. 2023). The empirical design also relates to policy learning and experimentation when interventions alter the state on which later decisions are based (Perdomo et al. 2020; Johari et al. 2022; Bajari et al. 2023; Holtz et al. 2025; Hu and Shum 2012).[6] Rothenberg's rank perspective provides a familiar identification analogue for the curvature-rank calculations (Rothenberg 1971).

2. Policy transport and target holonomy

Begin with a policy cycle whose observable path closes and whose latent composition changes. The general model then specifies the instruments, implementation rule, and future probe needed to interpret such a cycle. Curvature measures the leading response consequence of reversing its order.

2.1. A two-instrument policy cycle

A reported state can return while a hidden coordinate changes. Interpret the two reported coordinates as policy-controlled operating summaries and the third as composition relevant to a later response. Let $\mathbb{S} = \mathbb{R}^3$, with state $s = (x, y, h)$ and report $q(s) = (x, y)$. The

[6] Trajectoid constructions give a geometric example in which a closed base path changes a transported state (Sobolev et al. 2023).

horizontal vector fields associated with the two phase coordinates are

$$H_1 = \partial_x, \qquad H_2 = \partial_y + x\partial_h. \tag{3}$$

A move of length a along H_1 changes x. A move of length b along H_2 changes y and adds xb to h. Compose the moves $+aH_1, +bH_2, -aH_1, -bH_2$. In this example, Hol_γ denotes the endpoint map generated by those moves, and τ is the response to the same future probe at either endpoint. The general policy technology and endpoint map are defined below.

PROPOSITION 1 (Policy rectangle). *The rectangle returns (x, y) to its initial value and maps h to $h + ab$. Reversing its orientation maps h to $h - ab$. For a linear probe $\tau(s) = h\ell \in \mathbb{R}^r$,*

$$\tau(\mathrm{Hol}_{\gamma^+_{a,b}}(s)) - \tau(\mathrm{Hol}_{\gamma^-_{a,b}}(s)) = 2ab\ell. \tag{4}$$

An additive correction $c = -ab$ closes the scalar target, and two opposite rectangles close it by composition.

The Lie bracket $[H_1, H_2] = \partial_h$ is vertical, and ab is the signed phase area.[7] The product ab records how the order of two otherwise offsetting policy changes alters the hidden coordinate. The following model states which report motions are feasible, how they are implemented, and when this calculation extends to a target response.

2.2. Economic state, report, and attainable phase motion

Let $\mathcal{S}$ be an n-dimensional smooth state manifold and let $\mathcal{X}$ be an m-dimensional smooth phase manifold equipped with a Riemannian metric. The map $q : \mathcal{S} \to \mathcal{X}$ reports the phase point. Its vertical space at s is

$$\mathcal{V}_s = \ker Dq_s = T_s q^{-1}(q(s)). \tag{5}$$

The controlled state law is

$$\dot{s} = f_s + G_s u, \qquad d_s = Dq_s f_s, \qquad A_s = Dq_s G_s : \mathbb{R}^k \to T_{q(s)}\mathcal{X}, \tag{6}$$

where f is the autonomous drift, $u \in \mathbb{R}^k$ is the policy deviation, and R_s is the positive-definite implementation-cost matrix. A prime on a map into a phase tangent denotes the adjoint under the Euclidean policy inner product and the phase metric. The attainable relative phase velocities form $\operatorname{im} A_s$.

ASSUMPTION 1 (Constant-rank policy technology). *The maps q, f, G, and R are C^3, q is*

[7] For connections and geometric phases in controlled mechanical systems, see Kobayashi and Nomizu (1963); Kelly and Murray (1995); Hermann and Krener (1977) relates Lie-bracket spans to nonlinear accessibility and observability.

a submersion, R_s is symmetric positive definite, and A_s has constant rank $\rho \le m$. There are a smooth rank-ρ subbundle $\mathcal{E} \subseteq T\mathcal{X}$ and a vector field d on $\mathcal{X}$ such that

$$\operatorname{im} A_s = \mathcal{E}_{q(s)}, \qquad Dq_s f_s = d_{q(s)}. \tag{7}$$

The projectability conditions require states with the same report to have the same set of attainable phase velocities. This restriction can be assessed by comparing policy-to-phase maps within an observational fiber. A requested phase velocity $v \in T_x\mathcal{X}$ at $x = q(s)$ is feasible when $e = v - d_x \in \mathcal{E}_x$. Among the policies that deliver this velocity, select

$$u^E(s, v) = \operatorname*{arg\,min}_{u \in \mathbb{R}^k} \frac{1}{2} u' R_s u \quad \text{subject to} \quad A_s u = v - d_x. \tag{8}$$

Write $C_s = A_s R_s^{-1} A_s'$ and let $C_s^\dagger$ be its Moore–Penrose inverse. Define the relative state lift $L_s^E : \mathcal{E}_x \to T_s\mathcal{S}$ by

$$u^E(s, v) = R_s^{-1} A_s' C_s^\dagger (v - d_x), \qquad L_s^E e = G_s R_s^{-1} A_s' C_s^\dagger e. \tag{9}$$

PROPOSITION 2 (Partial policy connection). *Under assumption 1, equation* (9) *is the unique solution of* (8). *For every $e \in \mathcal{E}_x$, it satisfies $Dq_s L_s^E e = e$ and determines the decomposition*

$$Dq_s^{-1}(\mathcal{E}_x) = \mathcal{V}_s \oplus \mathcal{H}_s^E, \qquad \mathcal{H}_s^E = \operatorname{im} L_s^E. \tag{10}$$

If $\mathcal{E}$ is involutive, $\mathcal{H}^E$ is an Ehresmann connection on the restriction of the reported-state bundle to each integral leaf of $\mathcal{E}$. If $\rho = m$, equation (10) *becomes $T_s\mathcal{S} = \mathcal{V}_s \oplus \mathcal{H}_s^E$ and $C_s^\dagger = C_s^{-1}$.*

The instrument set, equilibrium restrictions, drift compensation, and R_s jointly determine the partial connection. Changing one of these objects changes L_s^E and can alter the state displacement accumulated along a feasible phase loop. Where rank changes, the state space separates into regular strata. Within a stratum, the smallest positive singular value of $A_s R_s^{-1/2}$ records the local conditioning of implementation.

When $d_x \in \mathcal{E}_x$, decompose the autonomous drift as

$$f_s = L_s^E d_x + b_s, \qquad b_s \in \mathcal{V}_s. \tag{11}$$

The field b_s is the vertical residual drift relative to the selected minimum-cost lift. On a domain where $b_s = 0$, equation (12) reduces to $\dot{s}_t = L_{s_t}^E \dot{\gamma}_t$. Spatial rectangles are then flow commutators of the relative horizontal lifts, which is the setting used for the local curvature results below. A controller can produce the same reduction by offsetting b_s and recentering the policy. With a general b_s, finite-time fiber transport also depends on loop duration; the path-space and polar-transport objects retain this dependence.

Forward-looking implementation. The local lift has a sequence-space counterpart when expectations and terminal conditions determine the report path jointly. Its path-space construction carries a feasibility residual. A common terminal comparison then evaluates an admissible path.

2.3. Accessible paths and target holonomy

A C^1 phase path $\gamma : [0,1] \to \mathcal{X}$ is policy-feasible when $\dot\gamma_t - d_{\gamma(t)} \in \mathcal{E}_{\gamma(t)}$. For an initial state s satisfying $q(s) = \gamma(0)$, its lift solves

$$\dot s_t = f_{s_t} + L^E_{s_t}\big(\dot\gamma_t - d_{\gamma(t)}\big), \qquad s_0 = s. \tag{12}$$

Equations (7) and (10) give $Dq_{s_t}\dot s_t = \dot\gamma_t$, hence $q(s_t) = \gamma(t)$ throughout the lift. The analysis uses feasible paths for which this solution stays in the constant-rank domain. If $\gamma(1) = \gamma(0) = x$, let $\mathcal{D}_\gamma \subseteq q^{-1}(x)$ be the open set of initial states whose lifted solutions exist through time one and remain in the constant-rank domain. The endpoint defines a local accessible fiber map

$$\mathrm{Hol}_\gamma : \mathcal{D}_\gamma \to q^{-1}(x), \qquad \mathrm{Hol}_\gamma(s) = s_1. \tag{13}$$

The drift and local sections of $\mathcal{E}$ generate control orbits containing the feasible phase paths. Closed paths within an orbit generate a local semigroup of partially defined fiber maps. If the negative of every feasible lifted field is also feasible, reversed paths provide local inverses, and the maps form a local pseudogroup on their common domains. This formulation allows $\rho < m$ and accommodates a noninvolutive $\mathcal{E}$, for which iterated brackets can enlarge the reachable phase set.

Let a standardized future probe policy be fixed. The target $\tau : \mathcal{S} \to \mathbb{R}^r$ is the conditional mean of its recorded response, which can contain several response horizons. The target holonomy defect is equation (1).

DEFINITION 1 (Loop sufficiency). *Let $\mathcal{O}^{\mathrm{Hol},E}_s$ be the set reached from s by lifts of feasible contractible loops based at $q(s)$. The phase report is sufficient for τ on the accessible loop orbit at s when τ is constant on $\mathcal{O}^{\mathrm{Hol},E}_s$.*

PROPOSITION 3 (Holonomy-orbit criterion). *Suppose $\mathcal{O}_s^{\mathrm{Hol},E}$ is a connected immersed submanifold. The phase report is sufficient for τ on the accessible loop orbit at s exactly when*

$$D\tau_z\xi = 0 \quad \text{for every } z \in \mathcal{O}_s^{\mathrm{Hol},E} \text{ and } \xi \in T_z\mathcal{O}_s^{\mathrm{Hol},E}. \tag{14}$$

If the accessible holonomy orbit equals the connected component of the fiber, this condition is equivalent to constancy of τ on that fiber component. When parallel transports of feasible curvature and commutator values span the accessible holonomy-orbit tangent space, equation (14) *is equivalent to annihilation of those transported values.*

Environment comparisons. The environment is held fixed for the loop contrast. Varying it yields a comparative-static derivative and requires an additional joint-inference design.

2.4. Target curvature and phase sufficiency

Curvature on an accessible leaf. Suppose first that $\mathcal{E}$ is involutive near x, and let $\mathcal{L}_x$ be its ρ-dimensional integral leaf. On $Dq_s^{-1}(\mathcal{E}_x)$, let $P_s^{\mathcal{V},E} = I - L_s^E Dq_s$ be the vertical projection from (10). For $v, w \in T_x\mathcal{L}_x = \mathcal{E}_x$, extend them to commuting coordinate vector fields on the leaf and let v^H, w^H denote their relative horizontal lifts. Define

$$\Omega_s^E(v,w) = P_s^{\mathcal{V},E}[v^H, w^H]_s \in \mathcal{V}_s, \qquad \kappa_\tau^E(s; v, w) = D\tau_s\Omega_s^E(v,w) \in \mathbb{R}^r. \tag{15}$$

The first object is the curvature of the partial policy connection under the stated sign convention. The second retains the component read by the future probe. Subsequent formulas suppress the superscript E.

Let the two coordinate rectangles have ordered sides

$$\gamma_\varepsilon^+ : (+\varepsilon v, +\varepsilon w, -\varepsilon v, -\varepsilon w), \qquad \gamma_\varepsilon^- : (+\varepsilon w, +\varepsilon v, -\varepsilon w, -\varepsilon v). \tag{16}$$

Both loops start and finish at x.

THEOREM 1 (Direction reversal and target curvature). *Suppose the accessible-leaf connection and τ are C^4, $d \in \mathcal{E}$, and the vertical residual drift in* (11) *vanishes on the flow neighborhood. In a state chart around s,*

$$\mathrm{Hol}_{\gamma_\varepsilon^+}(s) = s + \varepsilon^2\Omega_s(v,w) + O(\varepsilon^3), \tag{17}$$

$$\mathrm{Hol}_{\gamma_\varepsilon^-}(s) = s - \varepsilon^2\Omega_s(v,w) + O(\varepsilon^3), \tag{18}$$

and

$$\frac{\tau(\mathrm{Hol}_{\gamma_\varepsilon^+}(s)) - \tau(\mathrm{Hol}_{\gamma_\varepsilon^-}(s))}{2\varepsilon^2} = \kappa_\tau(s; v, w) + O(\varepsilon). \tag{19}$$

Let $\Delta(\varepsilon)$ denote the numerator in (19)*, and construct mirrored loops by replacing ε with $-\varepsilon$ in every side. Then*

$$\frac{\Delta(\varepsilon) + \Delta(-\varepsilon)}{4\varepsilon^2} = \kappa_\tau(s; v, w) + O(\varepsilon^2). \tag{20}$$

The two-loop contrast identifies the sign associated with oriented area. The third-order term changes sign under $\varepsilon \mapsto -\varepsilon$, so the mirror average removes it. In Proposition 1, the remainder is zero and equation (20) holds for every a, b in the stated domain.

COROLLARY 1 (Curvature diagnostic). *If $\kappa_\tau(s; v, w) \neq 0$, the phase report is locally deficient for the probe target along the policy connection. If the transported curvature values span the holonomy-orbit tangent space and $D\tau$ annihilates every one of them, the phase report is loop-sufficient on that orbit.*

The diagnostic is specific to the chosen probe response. A distribution may move after a loop in a direction contained in $\ker D\tau_s$, yielding zero target curvature for that probe. A different probe can load on the same movement and produce a nonzero contrast.

Noninvolutive access and Lie-word closure. When $\mathcal{E}$ is noninvolutive, feasible phase paths lie on the control orbit generated by its local sections. Let $e_1, \ldots, e_\rho$ be a local frame, let $H_i = L^E e_i$ be the relative state lifts, and let $\mathrm{Lie}_{\leq \ell}(H)$ denote bracket words of length at most ℓ. The target-relevant vertical word space at s is

$$\mathcal{W}_{\ell,\tau}(s) = \mathrm{span}\left\{D\tau_s Z_s : Z \in \mathrm{Lie}_{\leq \ell}(H),\ Dq_s Z_s = 0\right\} \subseteq \mathbb{R}^r. \tag{21}$$

PROPOSITION 4 (Accessible Lie-word bridge). *Consider a nonempty finite reversible library of commutator controls with leading vertical words $Z_1, \ldots, Z_J$ of lengths $\ell_1, \ldots, \ell_J$, and write $\ell_{\max} = \max_j \ell_j$. Suppose the lifted fields and target have bounded derivatives through order $\ell_{\max} + 1$ on a compact flow neighborhood. Suppose each projected commutator endpoint admits a feasible phase-closing segment whose state contribution has the next asymptotic order. Parameterize the jth signed word by its leading coefficient a_j. Then its target displacement is*

$$h_{\tau,j}(a_j, s) = a_j D\tau_s Z_j(s) + O\left(|a_j|^{1+1/\ell_j}\right). \tag{22}$$

The baseline defects removable at leading order form

$$\mathrm{span}\{D\tau_s Z_j(s) : j = 1, \ldots, J\}. \tag{23}$$

Full local target reachability requires this span to have dimension r.

The curvature bridge is the length-two case: Z_j is a vertical commutator and a_j is oriented area. Higher words provide additional closure coordinates when instantaneous phase rank is limited and policy sequencing expands the accessible Lie algebra. The corresponding control-system example is a rolling constraint. Its admissible velocities have lower dimension than the configuration space, while repeated commutators still generate displacement (Bloch et al. 1996).

3. Identification and inverse design

The oriented-loop contrast connects the empirical and design parts of the analysis. Random assignment identifies how state transport changes the target. The estimated derivative can then be used as a coordinate system for constructing a loop that offsets an observed target defect.

A causal loop comparison needs assignments that vary the order of the report motions and then expose the terminal states to the same probe. The construction below states those assignments and their scale. The transition-economy exercise supplies all model states and policies; the LaDe exercise supplies predictive evidence under observed operations. Their distinct information sets are retained when interpreting the experimental result.

3.1. Four randomized arms

Index loop orientation by $o \in \{+, -\}$ and mirror sign by $\sigma \in \{+, -\}$. Arm (o, σ) receives the corresponding loop with signed side scale $\sigma\varepsilon$, followed by the common probe. Let $Y^p \in \mathbb{R}^r$ be the recorded probe response and let

$$m_{o,\sigma}(\varepsilon) = \mathbb{E}[Y^p \mid o, \sigma], \qquad \Delta_\sigma(\varepsilon) = m_{+,\sigma}(\varepsilon) - m_{-,\sigma}(\varepsilon). \tag{24}$$

Random assignment identifies these four conditional means. Because the initial-state distribution is shared across arms, their differences isolate the assigned paths. Conditioning and covariate adjustment can be incorporated in the arm means when they are prespecified.

Each assigned path is centered at the unit's initial report $q(S)$ and uses the same local direction fields across arms. Potential-outcome consistency for the probe means

$$\mathbb{E}[Y^p(o, \sigma) \mid S = s] = \tau\Big(\mathrm{Hol}_{\gamma^o_{\sigma\varepsilon,s}}(s)\Big), \tag{25}$$

where $\gamma^o_{\sigma\varepsilon,s}$ is the corresponding report loop based at $q(s)$. Thus the randomized arm contrast averages terminal-state probe responses under a common initial-state law.

Define the population and sample curvature contrasts

$$\kappa_\varepsilon = \frac{\Delta_+(\varepsilon) + \Delta_-(\varepsilon)}{4\varepsilon^2}, \qquad \widehat{\kappa}_\varepsilon = \frac{\overline{Y}_{+,+} - \overline{Y}_{-,+} + \overline{Y}_{+,-} - \overline{Y}_{-,-}}{4\varepsilon^2}. \tag{26}$$

THEOREM 2 (Shrinking-loop identification and inference). *Let $\varepsilon = \varepsilon_n > 0$ satisfy $\varepsilon_n \to 0$. Suppose independent experimental units are randomized across the four arms with $n_{o,\sigma}/n \to \pi_{o,\sigma} \in (0,1)$, equation (25) holds, the arm outcomes have uniformly bounded $2+\delta$ moments along the shrinking sequence, their covariance matrices converge to $\Sigma_{o,\sigma}$, and the smooth expansion in theorem 1 holds uniformly over the initial-state support. Then*

$$\kappa_\varepsilon = \mathbb{E}[\kappa_\tau(S; v, w)] + O(\varepsilon^2). \tag{27}$$

The area-normalized sampling limit is

$$\sqrt{n}\varepsilon^2(\widehat{\kappa}_\varepsilon - \kappa_\varepsilon) \implies \mathcal{N}(0, V), \qquad V = \frac{1}{16}\sum_{o,\sigma}\frac{\Sigma_{o,\sigma}}{\pi_{o,\sigma}}. \tag{28}$$

Consistency follows when $n\varepsilon^4 \to \infty$. Centering the limit at $\mathbb{E}[\kappa_\tau(S; v, w)]$ additionally follows when $\sqrt{n}\varepsilon^4 \to 0$. For $\varepsilon = n^{-a}$, both conditions hold when $1/8 < a < 1/4$.

In a market experiment, the independent units of assignment are policy clusters indexed by c, with $A_c = (o, \sigma)$ denoting the assigned path. The potential cluster outcome $\overline{Y}_c^p(a)$ averages terminal probe outcomes within cluster c under path a and allows users and sellers in that cluster to interact. Partial interference requires this outcome to depend on the cluster's assignment and within-cluster exposure, while assignments to other clusters leave it unchanged.

THEOREM 3 (Cluster-assigned path identification). *Suppose policy clusters are independently and identically sampled and paths are randomized with probabilities $\pi_a \in (0,1)$. Suppose partial interference holds at the chosen cluster partition and the arm-specific cluster means have uniformly bounded $2+\delta$ moments along the sampling sequence. Then each arm mean identifies $m_a = \mathbb{E}[\overline{Y}_c^p(a)]$, the clockwise–counterclockwise contrast identifies $m_{+,+} - m_{-,+}$, and the four-arm contrast in (26) identifies κ_ε. If arm a contains J_a clusters and M_a terminal observations per cluster, with*

$$\mathrm{Var}\big(\overline{Y}_c^p(a)\big) = \Sigma_{B,a} + \frac{\Sigma_{W,a}}{M_a}, \tag{29}$$

then, conditional on the realized arm count J_a or under fixed-count randomization,

$$\mathrm{Var}(\overline{Y}_a) = \frac{\Sigma_{B,a}}{J_a} + \frac{\Sigma_{W,a}}{J_a M_a}. \tag{30}$$

Along a sequence with $J = \sum_a J_a \to \infty$, $J_a/J \to \pi_a$, and convergent arm-specific cluster-mean covariance matrices, the cluster central limit theorem gives the finite-loop and area-normalized limits after substituting (30) *into their arm contrasts. Cluster-robust arm covariances remain valid when within-cluster dependence is summarized directly by* $\mathrm{Var}(\overline{Y}_c^p(a))$.

Additional assignment designs. Randomized environments, cluster-assigned paths, and matched-set comparisons require their own observations and covariance laws. The four-arm contrast is the reference experiment for the local curvature calculation.

3.2. Power frontier and assignment level

The scaling in Theorem 2 also determines which curvature values the experiment can detect. Consider the scalar Gaussian submodel in which arm $a = (o, \sigma)$ supplies n_a independent observations with variance $\sigma_a^2 > 0$ and mean

$$m_a(\varepsilon; \kappa) = \mu_a + s_a \varepsilon^2 \kappa, \qquad (s_{+,+}, s_{-,+}, s_{+,-}, s_{-,-}) = (1, -1, 1, -1). \tag{31}$$

The nuisance means μ_a are held fixed across the two simple curvature hypotheses. The same calculation applies when their admissible nuisance subspace is orthogonal to the contrast vector $(s_a)_a$.

THEOREM 4 (Local-holonomy power frontier). *For two curvature values κ_0 and κ_1, the Kullback–Leibler divergence in the submodel* (31) *is*

$$\mathrm{KL}(P_{\kappa_0}, P_{\kappa_1}) = \frac{\varepsilon^4(\kappa_1 - \kappa_0)^2}{2} \sum_a \frac{n_a}{\sigma_a^2}. \tag{32}$$

Every level-α test of κ_0 against κ_1 has power bounded above by

$$\alpha + \frac{|\kappa_1 - \kappa_0|\varepsilon^2}{2} \left(\sum_a \frac{n_a}{\sigma_a^2} \right)^{1/2}. \tag{33}$$

Consequently, $n\varepsilon^4(\kappa_1 - \kappa_0)^2 \to 0$ with stable allocation and variances makes the two sequences asymptotically indistinguishable. For $V > 0$, $\varepsilon > 0$, $\kappa \neq 0$, $\alpha \in (0, 1)$, and $\beta \in (0, 1/2)$, the Wald statistic in (28) *has absolute local noncentrality $\sqrt{n}\varepsilon^2|\kappa|/\sqrt{V}$. A sufficient Gaussian planning rule for two-sided size α and power at least $1 - \beta$ takes the smallest integer weakly above the continuous value*

$$n_{\alpha,\beta}^{\mathrm{cont}} = \frac{(z_{1-\alpha/2} + z_{1-\beta})^2 V}{\varepsilon^4 \kappa^2}. \tag{34}$$

The randomization unit is the economy, jurisdiction, market, or other policy cluster that receives the path. Additional household draws from one terminal distribution reduce measurement error for that realized path. Independent policy assignments come from additional randomized clusters. For a macroeconomic implementation, cluster arm means and cluster-level covariance matrices therefore replace household-level quantities in (28)–(34), with the continuous sample-size value rounded upward. The finite-loop transport estimator pools transition information and uses the unscaled loop contrast.

3.3. Curvature-rank uncertainty

Experiments on several coordinate planes or at several base points supply the rows of a target-curvature matrix K. Let $\widehat{K}$ be the stacked estimator and suppose a simultaneous confidence construction yields

$$\Pr\{\|\widehat{K} - K\|_{\mathrm{op}} \le \rho_n\} \ge 1 - \alpha. \tag{35}$$

PROPOSITION 5 (Rank and bridge certificate). *Suppose $L \ge r$. On the event in* (35)*, every singular value satisfies $|\widehat{\sigma}_j - \sigma_j| \le \rho_n$. Hence $\widehat{\sigma}_r > \rho_n$ certifies full target rank and gives $\sigma_r(K) \ge \widehat{\sigma}_r - \rho_n$. The confidence set*

$$\mathcal{R}_n = \{\mathrm{rank}(K') : \|K' - \widehat{K}\|_{\mathrm{op}} \le \rho_n\} \tag{36}$$

contains the target-curvature rank with probability at least $1 - \alpha$.

Near a rank change, $\mathcal{R}_n$ can contain several integers, and the area bridge becomes poorly conditioned. Reporting the set makes that uncertainty visible in the design result and its residual bound.

3.4. Inverse loop design

Area coordinates and accessible phase dimension. Fix $s \in q^{-1}(x)$ on an involutive accessible leaf and choose an orthonormal basis $E_1, \dots, E_L$ of $\Lambda^2 \mathcal{E}_x$ under the metric induced from the phase manifold, where $L = \binom{\rho}{2}$. Stack the feasible target curvatures as rows:

$$K_s = \begin{bmatrix} \kappa_\tau(s; E_1)' \\ \vdots \\ \kappa_\tau(s; E_L)' \end{bmatrix} \in \mathbb{R}^{L \times r}. \tag{37}$$

An area vector $a \in \mathbb{R}^L$ produces the second-order target displacement aK_s. Given a baseline defect $b \in \mathbb{R}^r$, the area closure equation is

$$b + aK_s = 0. \tag{38}$$

PROPOSITION 6 (Area reachability, codimension, and accessible phase dimension). *Assume $r \geq 1$. Equation (38) has a solution exactly when $b \in \mathrm{row}(K_s)$. On this row space, $a^\star = -bK_s^\dagger$ is the minimum-Euclidean-norm solution. If $\mathrm{rank}(K_s) = r$, every local target defect is reachable to second order, the closure areas form*

$$a^\star + \ker_{\mathrm{L}}(K_s), \qquad \ker_{\mathrm{L}}(K_s) := \{a \in \mathbb{R}^L : aK_s = 0\}, \tag{39}$$

and this affine set has dimension $L - r$ and codimension r. Full second-order reachability implies

$$r \leq \binom{\rho}{2} \leq \binom{m}{2}, \qquad \rho \geq \rho^\star(r) := \left\lceil \frac{1 + \sqrt{1 + 8r}}{2} \right\rceil. \tag{40}$$

A p-coordinate loop library with area map $a = \eta P$, $P \in \mathbb{R}^{p \times L}$, has effective curvature PK_s and requires $p \geq r$ for full target reachability. When K_s has rank r and $p \geq r$, the matrices P for which PK_s has rank r form an open dense subset of $\mathbb{R}^{p \times L}$.

The bound in (40) counts oriented two-dimensional faces of the attainable phase distribution. A displayed phase plane contributes a face only when policy can move the report in both directions. One small loop moves three target-response horizons along a single curvature vector. To reach an arbitrary three-vector, the design needs at least three independent attainable faces, first available at $\rho = 3$, or a library that combines curvature and Lie-word vectors from several base points.

Transversality of the curvature rank. Let Θ be a finite-dimensional smooth manifold, let $\theta \in \Theta$ index policy technologies, implementation metrics, or phase reports, and let the C^1 map $\mathcal{K} : \Theta \to \mathbb{R}^{L \times r}$ send a design to its target-curvature matrix. Write $\Sigma_d = \{K : \mathrm{rank}(K) = d\}$.

PROPOSITION 7 (Rank strata of policy curvature). *Suppose $L \geq r$ and $\mathcal{K}$ is transverse to every matrix-rank stratum Σ_d. For each $d < r$, the set $\mathcal{K}^{-1}(\Sigma_d)$ is either empty or a submanifold of codimension*

$$(L - d)(r - d). \tag{41}$$

Full-rank designs form an open dense subset of Θ. If $L < r$, every target-curvature matrix has rank at most L.

Equation (41) states the regularity condition used by the phase-dimension result. It also shows which design perturbations can move a singular curvature matrix away from its rank stratum. Restrictions on G, q, or R may confine $\mathcal{K}(\Theta)$ to a lower-rank family, in which case transversality fails in a way that can be traced to those economic restrictions.

The oriented-area bridge. Implement a_j by an elementary rectangle in the plane E_j, with orientation $\operatorname{sign}(a_j)$ and side lengths proportional to $\sqrt{|a_j|}$. Concatenating the L rectangles returns the phase point to x. Let the resulting target displacement at the terminal state of the baseline cycle be

$$F_s(a) = aK_s + \mathcal{R}_s(a), \qquad \|\mathcal{R}_s(a)\| \le C\|a\|^{3/2}, \qquad C \ge 0 \tag{42}$$

on a neighborhood of zero.

THEOREM 5 (Curvature bridge and residual sequence). *Suppose K_s has rank r and* (42) *holds. For a baseline defect b, let $a^\star = -bK_s^\dagger$ and suppose $a^\star$ lies in the remainder neighborhood. The bridged defect satisfies*

$$\|b + F_s(a^\star)\| \le C\|K_s^\dagger\|^{3/2}\|b\|^{3/2}. \tag{43}$$

Consider repeated updates at the successive terminal states, suppose every proposed area update lies in its local remainder neighborhood, and suppose the remainder constant is bounded by C and $\|K_s^\dagger\| \le \beta$ along the sequence. With e_k denoting target-defect size and $M = C\beta^{3/2}$,

$$e_{k+1} \le Me_k^{3/2}. \tag{44}$$

If $M = 0$, then $e_k = 0$ for every $k \ge 1$. If $M > 0$ and $M^2 e_0 < 1$, then

$$e_k \le M^{-2}(M^2 e_0)^{(3/2)^k}. \tag{45}$$

The calculation uses the target-curvature matrix, its pseudoinverse, and elementary phase loops. The full state transition enters evaluations of target defects and curvature; matrix inversion takes place in dimensions L and r. Equation (45) supplies a stopping certificate for the repeated bridge.

PROPOSITION 8 (Estimated bridge residual). *Suppose $L \ge r$, $\rho, \delta \ge 0$, $\|\widehat{K} - K\|_{\mathrm{op}} \le \rho$, $\|\widehat{b} - b\| \le \delta$, and $\widehat{\sigma}_r > \rho$. Set $\widehat{a} = -\widehat{b}\widehat{K}^\dagger$. Then K and $\widehat{K}$ have target rank r, and*

$$\|b + \widehat{a}K\| \le \delta + \frac{\rho}{\widehat{\sigma}_r}\|\widehat{b}\|. \tag{46}$$

If $\widehat{a}$ lies in the remainder neighborhood, adding the nonlinear term in (42) *gives the implementable bound*

$$\|b + F_s(\widehat{a})\| \le \delta + \frac{\rho}{\widehat{\sigma}_r}\|\widehat{b}\| + C\|\widehat{a}\|^{3/2}. \tag{47}$$

4. Dynamic Gorman aggregation

The loop library determines which internal coordinates must accompany the report. Transported target covectors give the local dimension of these coordinates. A finite affine quotient then propagates them across sequences of loops.

4.1. Target-sensitive measurements

Fix a phase loop and linearize its observed history and future probe response around a baseline state. Write $\mathbb{K} = \mathbb{R}$, write perturbations as row vectors $z \in \mathbb{K}^N$, and let

$$z \longmapsto zO \in \mathbb{K}^d, \qquad z \longmapsto zD \in \mathbb{K}^r. \tag{48}$$

The tangent space to the observational fiber is $\{z : zO = 0\}$. The probe response is identified on this fiber when $zO = 0$ implies $zD = 0$.

Let the rows of N_O span $\{z : zO = 0\}$ and set

$$A = N_O D, \qquad d^\star = \operatorname{rank}(A). \tag{49}$$

Choose a full-row-rank matrix $B \in \mathbb{K}^{d^\star \times r}$ whose rows form a basis of $\operatorname{row}(A)$ and a generalized inverse $B^\dagger$ satisfying $zB^\dagger B = z$ on $\operatorname{row}(B)$. Define the target-sensitive measurement matrix

$$M^\star = DB^\dagger \in \mathbb{K}^{N \times d^\star}. \tag{50}$$

When $d^\star = 0$, the measurement and decoder are empty, and the formulas continue to hold with zero-dimensional matrices.

THEOREM 6 (Minimum measurement dimension on an observational fiber). *For every z satisfying $zO = 0$,*

$$zD = (zM^\star)B. \tag{51}$$

If a d_M-coordinate measurement matrix M admits a linear decoder C on the fiber, so that $(N_O M)C = N_O D$, then $d_M \ge d^\star$. Thus $d^\star$ is the minimum dimension among linearly decodable *measurements, and* (50) *attains it.*[8]

For a finite distribution μ, each coordinate of $\mu M^\star$ is a weighted micro moment. If two terminal distributions share the phase report and these $d^\star$ moments, equation (51) assigns them the same linearized probe response. Sample averages of the weights estimate the moments. The resulting empirical object has dimension $d^\star$, even when the distribution has many support points.

The randomized loop experiment uses probe outcomes to identify target curvature through (26). The measurement theorem addresses a related practical problem: a full probe experiment may be unavailable at every policy iteration. In that case, target-sensitive moments can predict and monitor the defect. A confidence region for (O, D) produces a set of candidate measurement ranks, and the operator ball in (35) carries this uncertainty into the area bridge.

4.2. A phase-conditional Gorman form

Gorman's polar form writes household expenditure as

$$e_i(p, u_i) = a_i(p) + b(p)u_i, \tag{52}$$

where the price index b is common across households (Gorman 1953, 1961). Shephard's lemma gives the aggregate Hicksian demand

$$H(p, \boldsymbol{u}) = \nabla_p A(p) + U\nabla_p b(p), \qquad A(p) = \sum_i a_i(p), \qquad U = \sum_i u_i. \tag{53}$$

Thus a vector of demands depends on the utility profile through one additive coordinate and a common price-dependent slope.

For the policy problem, let $M : \mathcal{S} \to \mathbb{R}^g$ collect target-sensitive additive moments and let $B : \mathcal{X} \to \mathbb{R}^{g\times r}$. A phase-conditional polar form is

$$\tau(s) = \alpha(q(s)) + M(s)B(q(s)). \tag{54}$$

For a finite distribution, $M(\mu) = \mu W$ for a matrix of microstate weights W. Equation (53) is the case $q = p$, $M = U$, $\alpha = \nabla_p A$, and $B = \nabla_p b$.

PROPOSITION 9 (Polar factorization of target holonomy). *Suppose* (54) *holds on a regularity domain. For every feasible phase loop γ based at $x = q(s)$ with $s \in \mathcal{D}_\gamma$,*

$$h_\tau(\gamma, s) = \left[M(\mathrm{Hol}_\gamma(s)) - M(s)\right]B(x). \tag{55}$$

[8] This finite-dimensional statement uses linear decoding on one observational fiber. Theorem 7 gives the differentiable first-order analogue for transported target covectors, and Proposition 10 states the local nonlinear quotient conditions.

For $v, w \in \mathcal{E}_x$ on an involutive accessible leaf, define $\kappa_M(s; v, w) = DM_s\Omega_s(v, w)$. Then

$$\kappa_\tau(s; v, w) = \kappa_M(s; v, w)B(x), \qquad K_s = K_{M,s}B(x), \tag{56}$$

and $\operatorname{rank}(K_s) \leq \min\{g, \operatorname{rank} K_{M,s}, \operatorname{rank} B(x)\}$.

Every target defect generated inside the phase fiber lies in the row space of $B(x)$. With $g = 1$, an r-coordinate target has a common polar direction. The single area coordinate of a two-dimensional phase diagram can move the target vector along that direction, and exact vector closure requires the baseline defect to lie on the same row. The factorization also reveals two ways that phase sufficiency can arise. Curvature may leave every moment in M fixed, or the common slope may annihilate the induced moment displacement.

4.3. Local dynamic Gorman closure

Fix $s \in \mathcal{S}$, write $x = q(s)$, and let $V_s = \ker Dq_s$ be the tangent space to the observational fiber. Let $\mathcal{L}_x$ contain the empty word and the feasible holonomy words based at x whose maps are defined on a neighborhood of s. For $\gamma \in \mathcal{L}_x$, the vertical differential

$$T_{\gamma,s} = D\,\mathrm{Hol}_\gamma(s)|_{V_s} : V_s \longrightarrow V_{\mathrm{Hol}_\gamma(s)} \tag{57}$$

transports a hidden-state perturbation around the reported loop. Pull the endpoint target sensitivities back to the initial fiber and define

$$\mathcal{W}^\star_{\tau,s} = \operatorname{span}\left\{T^*_{\gamma,s}\left(D\tau^a_{\mathrm{Hol}_\gamma(s)}|_{V_{\mathrm{Hol}_\gamma(s)}}\right) : \gamma \in \mathcal{L}_x,\ a = 1, \dots, r\right\} \subseteq V^*_s. \tag{58}$$

By the chain rule, each generator in this span equals $D(\tau^a \circ \mathrm{Hol}_\gamma)_s|_{V_s}$. The loop library, target family, and policy-induced connection jointly determine $\mathcal{W}^\star_{\tau,s}$.

For a measurement $M = (M_1, \dots, M_g)$, define first-order loop sufficiency at s by

$$\ker(DM_s|_{V_s}) \subseteq \bigcap_{\gamma \in \mathcal{L}_x} \ker\left(D(\tau \circ \mathrm{Hol}_\gamma)_s|_{V_s}\right). \tag{59}$$

When the state is a distribution μ, a Gorman measurement has coordinates $M_j(\mu) = \int \psi_j(z)\, d\mu(z)$, so each coordinate aggregates a microstate weight by addition.

THEOREM 7 (Dynamic aggregation characterization). *A C^1 measurement $M : \mathcal{S} \to \mathbb{R}^g$ satisfies* (59) *at s if and only if*

$$\mathcal{W}^\star_{\tau,s} \subseteq \operatorname{span}\left\{dM_{1,s}|_{V_s}, \dots, dM_{g,s}|_{V_s}\right\}. \tag{60}$$

Consequently $\operatorname{rank}(DM_s|_{V_s}) \geq d_G^{\mathrm{loc}}(s)$ *for every first-order loop-sufficient measurement, where*

$$d_G^{\mathrm{loc}}(s) = \dim \mathcal{W}^\star_{\tau,s}. \tag{61}$$

If $\mathcal{W}^\star_{\tau,s}$ *admits a basis consisting of differentials of additive moment functionals, those* $d_G^{\mathrm{loc}}(s)$ *functionals attain the bound among additive measurements. Thus additive attainability is equivalent to the existence of an additive differential frame for the transported target span.*

For an environment library $\mathcal{I}_0$, write $\Psi^a_{\eta,\gamma} = \tau^{\eta,a} \circ \mathrm{Hol}^\eta_\gamma$, let $T_{\eta,\gamma,s} = D\,\mathrm{Hol}^\eta_\gamma(s)|_{V_s}$, and define the shock-level response span

$$\mathcal{W}^\star_{\tau,\mathcal{I}_0,s} = \operatorname{span}\left\{ T^*_{\eta,\gamma,s}\left(D\tau^{\eta,a}_{\mathrm{Hol}^\eta_\gamma(s)}|_{V_{\mathrm{Hol}^\eta_\gamma(s)}} \right) : \eta \in \mathcal{I}_0,\ \gamma \in \mathcal{L}^\eta_x,\ a = 1, \ldots, r \right\}. \tag{62}$$

Fix $\eta_0 \in \mathcal{I}_0$ in the interior of the environment domain and a loop library feasible on its neighborhood. The shock-derivative span and the span required to preserve both the prespecified shock family and the first shock derivative are

$$\mathcal{J}^\star_{\tau,\eta_0,s} = \operatorname{span}\left\{ \partial_\eta D_s \Psi^a_{\eta,\gamma}(s)\Big|_{\eta=\eta_0} : \gamma \in \mathcal{L}_x,\ a = 1, \ldots, r \right\}, \qquad \mathcal{W}^{\star,1}_{\tau,\mathcal{I}_0,\eta_0,s} = \mathcal{W}^\star_{\tau,\mathcal{I}_0,s} + \mathcal{J}^\star_{\tau,\eta_0,s}. \tag{63}$$

Here level sufficiency preserves first-order state variation of every $\Psi_{\eta,\gamma}$ at the prespecified values of η. Shock-derivative robustness additionally preserves first-order state variation of $\partial_\eta \Psi_{\eta,\gamma}$ at η_0.

COROLLARY 2 (Minimum shock-library aggregation dimension). *A measurement is first-order sufficient for every prespecified environment, target, and loop in* (62) *exactly when its fiber differentials span* $\mathcal{W}^\star_{\tau,\mathcal{I}_0,s}$. *The minimum differentiable measurement rank is*

$$d_{G,0}^{\mathrm{loc}}(s; \mathcal{I}_0) = \dim \mathcal{W}^\star_{\tau,\mathcal{I}_0,s}, \tag{64}$$

with attainment by additive moments when the joint span admits an additive differential frame. For $\mathcal{I}_1 \subseteq \mathcal{I}_2$,

$$d_{G,0}^{\mathrm{loc}}(s; \mathcal{I}_1) \leq d_{G,0}^{\mathrm{loc}}(s; \mathcal{I}_2). \tag{65}$$

The minimum rank for level sufficiency over $\mathcal{I}_0$ *together with shock-derivative robustness at* η_0 *is*

$$d_{G,1}^{\mathrm{loc}}(s; \mathcal{I}_0, \eta_0) = \dim \mathcal{W}^{\star,1}_{\tau,\mathcal{I}_0,\eta_0,s} \geq d_{G,0}^{\mathrm{loc}}(s; \mathcal{I}_0), \tag{66}$$

and equality holds exactly when $\mathcal{J}^\star_{\tau,\eta_0,s} \subseteq \mathcal{W}^\star_{\tau,\mathcal{I}_0,s}$. *Additive moments attain the joint bound when the enlarged span admits an additive differential frame.*

The theorem and corollary translate the loop library into a local measurement rule.

Each basis covector assigns weights to microstate perturbations, and its pullback incorporates every policy history and environment included in the library. Enlarging the target, shock, or loop library can increase the dimension needed for level sufficiency; the shock jet may add further coordinates. A common C^1 decoder over an open environment neighborhood implies jet closure by differentiation. Closure verified only on a finite shock grid covers levels at those grid points, leaving the derivative between them to be studied separately. If the baseline s°_η moves with the environment, the aggregate comparative static also includes the observed tangent $D(q, M)_{s^\circ_\eta} \partial_\eta s^\circ_\eta$. Level and jet closure recover its target effect from this tangent. These statements concern first-order variation at s. The next result integrates a constant-rank family into a local quotient, followed by an exact finite affine construction.

4.4. The dynamic target quotient

Fix a phase point x and a relative-open subset U of its fiber that is invariant under a finite library of feasible holonomy maps $g_1, \dots, g_J$. Let $\mathcal{L}_x$ contain the empty word and every finite word in this library. For $\gamma \in \mathcal{L}_x$, write g_γ for the corresponding composition and $\Psi_\gamma = \tau \circ g_\gamma$. Two states have the same dynamic target type when

$$s \equiv_{\tau,\mathcal{L}_x} s' \quad \Longleftrightarrow \quad \Psi_\gamma(s) = \Psi_\gamma(s') \quad \text{for every } \gamma \in \mathcal{L}_x. \tag{67}$$

The relation is invariant under each generator: $s \equiv_{\tau,\mathcal{L}_x} s'$ implies $g_j(s) \equiv_{\tau,\mathcal{L}_x} g_j(s')$. A C^1 representation $M : U \to \mathbb{R}^g$ is dynamically target-sufficient when every Ψ_γ has a C^1 decoder through M. It is transport-closed when each $M \circ g_j$ has a C^1 update through M.

On the fiber, the transported target span in (58) becomes the response codistribution

$$\mathcal{W}^\star_{\tau,s} = \operatorname{span}\left\{D\Psi^a_\gamma(s) : \gamma \in \mathcal{L}_x,\ a = 1, \dots, r\right\} \subseteq T^*_s U, \tag{68}$$

where Ψ^a_γ is target coordinate a after word γ.

PROPOSITION 10 (Dynamic target quotient under policy holonomy). *Suppose $U_0 \subseteq U$ is a generator-invariant coordinate neighborhood and $\mathcal{W}^\star_\tau$ is a smooth codistribution of constant rank d on U_0. Suppose d scalar word-response coordinates have differentials that form a frame for $\mathcal{W}^\star_\tau$ on U_0. Let $F : U_0 \to \mathbb{R}^d$ be their joint map and suppose every nonempty level set $F^{-1}(y)$ in U_0 is connected. Then there is a C^1 submersion $M^\star : U_0 \to \mathbb{R}^d$ whose components are these response coordinates and whose fibers are the local equivalence classes in* (67). *The representation $M^\star$ is dynamically target-sufficient and transport-closed. Every dynamically target-sufficient C^1 representation M satisfies*

$$\operatorname{rank} DM_s \geq d \qquad \textit{for every } s \in U_0. \tag{69}$$

Thus $d = \operatorname{rank} \mathcal{W}^\star_\tau$ is the minimum local differentiable state dimension that propagates the target through every word in the loop library.

The quotient records state propagation separately from the condition for cycle closure. Its induced update describes how a closed report path moves the target-relevant aggregate state. A loop closes the target at s when the decoder has the same value at $M^\star(s)$ and at the updated quotient state. For a small oriented loop, curvature is the local derivative of this update.

Finite-loop extensions. Polar-quotient iteration and finite-loop transport inference add global or finite-cycle requirements to the local measurement criterion.

5. Evidence and implementation

The transition economy asks whether restoring the report can leave a different response to a fixed future probe. Because the economic state and policy technology are specified, it computes the transported distributions and the resulting response defect. Express records ask whether additional internal measurements help predict a declared target after conditioning on report history and observed flows. They also supply outcome-dispersion estimates for a prospective assigned-loop experiment. The model exercise establishes the mechanism within its specified equilibrium; the observational exercise evaluates the measurement requirement on available records.

5.1. Transition equilibrium

Distribution, prices, and path-space lift. The transition economy has 24 asset nodes and three income states. CRRA utility and logit-smoothed asset choice produce a 72-coordinate distribution μ_t. On the reported transition domain, $K(\mu_t) > 0$, $L(\mu_t) > 0$, $Y(\mu_t) > 0$, and $C_t > 0$. Aggregate assets and labor determine production capital and factor prices:

$$\begin{aligned} &K(\mu_t) = \zeta\mu_t a, \qquad && L(\mu_t) = \mu_t e, \qquad && Y(\mu_t) = ZK(\mu_t)^\alpha L(\mu_t)^{1-\alpha}, \\ &w(\mu_t) = (1-\alpha)\frac{Y(\mu_t)}{L(\mu_t)}, && r(\mu_t) = \alpha\frac{Y(\mu_t)}{K(\mu_t)} - \delta. && \end{aligned} \tag{70}$$

The calibration uses $\alpha = 0.36$, $\delta = 0.025$, $\zeta = 18.446$, and $Z = 0.57496$, which reproduce a wage of one and a net return of one percent at the stationary distribution. For every policy path, household value functions are solved backward along the endogenous wage and return path, distributions are propagated forward, and market prices are updated to a fixed point. A 12-quarter zero-policy tail joins the loop to the stationary continuation value. This is a finite-horizon perfect-foresight transition equilibrium in the sequence-space tradition (Reiter 2009; Auclert et al. 2021).

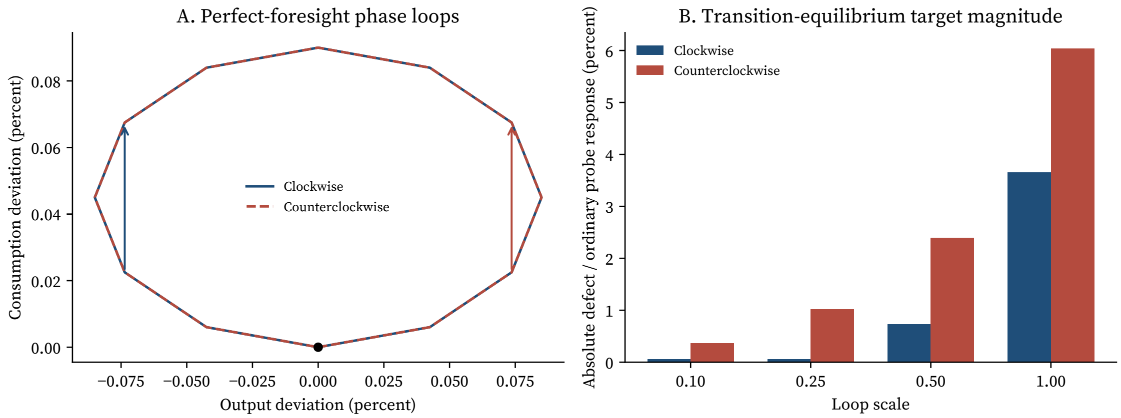

FIGURE 2. Policy holonomy in transition equilibrium

Panel A gives clockwise and counterclockwise 12-quarter output–consumption paths at scale 0.1. Both paths return to the stationary report point. Panel B reports the absolute horizon-0 consumption-probe defect divided by the ordinary horizon-0 probe response. Household decisions, factor prices, and distributions satisfy the perfect-foresight transition equilibrium along each path.

The policy vector contains a return wedge, a wealth-tilted transfer, and an income-tilted transfer in every quarter. The report is

$$q_t = 100\left(\log \frac{Y(\mu_{t+1})}{\overline{Y}},\ \log \frac{C_t}{\overline{C}}\right). \tag{71}$$

The path map $\Phi_{12} : \mathbb{R}^{36} \to \mathbb{R}^{24}$ includes the backward household solution, forward distribution, factor prices, and report sequence. The path-space lift in proposition A3 applies a minimum-norm Newton correction to all 36 policy coordinates. Its computed Jacobian has rank 24 along every reported loop. The smallest singular value across the four scales and two orientations is 0.0386, leaving a 12-dimensional local set of policy paths for each requested report path.

Closed output–consumption paths and future probes. For normalized path time $\theta \in [0, 1]$, orientation $o \in \{-1, 1\}$, and scale s, the requested report loop is

$$\gamma_{o,s}(\theta) = (os(0.85)\sin(2\pi\theta),\ s(0.45)[1 - \cos(2\pi\theta)])\,. \tag{72}$$

One traversal lasts 12 quarters. The standardized probe begins with a return-policy coordinate of −1.5 and decays at rate 0.68 over eight quarters. Its ordinary consumption effects at horizons 0, 3, and 7 are (−0.1183, −0.3934, −0.4414) percentage point. Figure 2 plots the scale-0.1 paths and the horizon-0 defects over four scales.

The counterclockwise defect rises from 0.00043 percentage point at scale 0.1 to 0.00714 at scale 1, corresponding to 0.36 and 6.04 percent of the ordinary response. The counter-

TABLE 1. Transition-equilibrium consumption-probe defects

Scale	CW defect	CCW defect	CW ratio (%)	CCW ratio (%)
0.10	0.00006	-0.00043	0.05	0.36
0.25	-0.00007	-0.00121	0.06	1.02
0.50	-0.00086	-0.00283	0.73	2.39
1.00	-0.00432	-0.00714	3.65	6.04

CW and CCW denote clockwise and counterclockwise. Defects are percentage-point changes in the horizon-0 response to the common probe, relative to the stationary initial distribution. Ratios divide absolute defects by the ordinary response of 0.11826 percentage point.

clockwise terminal-distribution distance rises from 0.00334 to 0.03961 in ℓ_1. The maximum report-path error across the table is 9.53×10^{-6} percentage point. At scale 0.1, the maximum policy-coordinate norm is 0.824 clockwise and 0.628 counterclockwise. This scale sequence records finite path transport in the transition equilibrium. A fixed-price calculation supplies shrinking-loop curvature and bridge diagnostics; the current-value general-equilibrium calculation supplies a second finite-loop polar-transport diagnostic.

5.2. Express operations

Operational reports and additive internal state. An Express network contains parcels of different ages and locations, couriers carrying different loads, and delivery areas with different task concentrations. Headquarters may summarize this state by reporting throughput, mean lead time, total backlog, and active capacity. Each operational region is treated as a high-dimensional organization, and the empirical question is whether this report is sufficient to predict its subsequent operating state. The public LaDe release records package acceptance and delivery events from operations in five Chinese cities (Wu et al. 2025).

For region r and day t, define the reported state

$$q_{rt} = \left(\log(1 + D_{rt}),\ \bar{\ell}_{rt},\ \log(1 + B_{rt}),\ \log(1 + C_{rt})\right), \tag{73}$$

where D_{rt} is delivered parcels, $\bar{\ell}_{rt}$ is mean lead time among deliveries with lead time at most seven days, B_{rt} is end-of-day backlog, and C_{rt} is the number of active couriers. The reported lead-time coordinate is set to zero on days without a measured delivery; the throughput coordinate records that event.

The Gorman coordinates are sums over parcels, couriers, or delivery areas. Let a_{it} be the age of backlogged parcel i, let L_{crt} and F_{crt} be courier c's backlog and delivered tasks, let N_{art} be delivered tasks in area of interest a, and let ℓ_{it} be parcel lead time. The six coordinates are

$$M_{rt} = \left(\sum_{i \in B_{rt}} a_{it},\ \sum_{i \in B_{rt}} a_{it}^2,\ \sum_c L_{crt}^2,\ \sum_c F_{crt}^2,\ \sum_a N_{art}^2,\ \sum_{i \in D_{rt}} \ell_{it}^2 \right). \tag{74}$$

The first two coordinates retain the backlog-age distribution. The next three describe concentration across couriers and delivery areas conditional on reported totals, and the last records delivery-time dispersion. Each coordinate has the additive form in Theorem 7: it sums a micro-unit weight over parcels, couriers, or areas. The target family contains future log throughput, mean lead time, and log backlog at horizons of one, three, and seven days. Backlog obeys the accounting relation $B_{r,t+h} = B_{rt} + \sum_{j=1}^{h} A_{r,t+j} - \sum_{j=1}^{h} D_{r,t+j}$. Conditional on a common future arrival path, the target therefore measures cumulative service capacity as a function of the initial internal state. Let H_{rt} collect the current report, its observed history, calendar information, and accepted-parcel flow. The following identity expresses the Gorman closure condition as a predictive-risk estimand.

PROPOSITION 11 (Predictive closure risk). *For a square-integrable scalar target Y, let $\mu_H = \mathbb{E}[Y \mid H]$ and $\mu_{HM} = \mathbb{E}[Y \mid H, M]$. The oracle squared-risk gain from carrying the additive state with the report is*

$$\inf_{g \in L^2(H)} \mathbb{E}[(Y - g)^2] - \inf_{f \in L^2(H,M)} \mathbb{E}[(Y - f)^2] = \mathbb{E}[(\mu_{HM} - \mu_H)^2]. \tag{75}$$

The report history is conditionally mean-sufficient for Y exactly when the right side is zero.

The sample analysis compares fixed learners on dates excluded from outcome fitting. The squared-loss contrast estimates a learner- and split-specific version of equation (75). If both learners consistently estimate their conditional means, this contrast converges to the oracle quantity. Absolute loss is also reported using the same fixed predictions.

Construction and predictive-closure design. After event validation, the sample contains 4,514,658 package records, 4,881 couriers, 172 operational regions, and 30,534 region-days. The analysis uses the 166 regions observed for at least 60 days. Online Appendix H describes the event accounting and variable construction.

For each target $\tau_{r,t+h}$, the analysis compares three specifications. The first uses the contemporaneous report. The history–flow benchmark adds accepted parcels and exact-calendar lags of reports and flows. The polar model further adds $M_{rt} B_h(q_{rt})$ with report-dependent common slopes. Within each region, the first 80 percent of dates with an observed target form the estimation sample, and the remaining dates form the evaluation sample. Sequential validation windows select penalties. Evaluation uses squared and absolute prediction loss with two-way cluster covariance by operational region and calendar date.

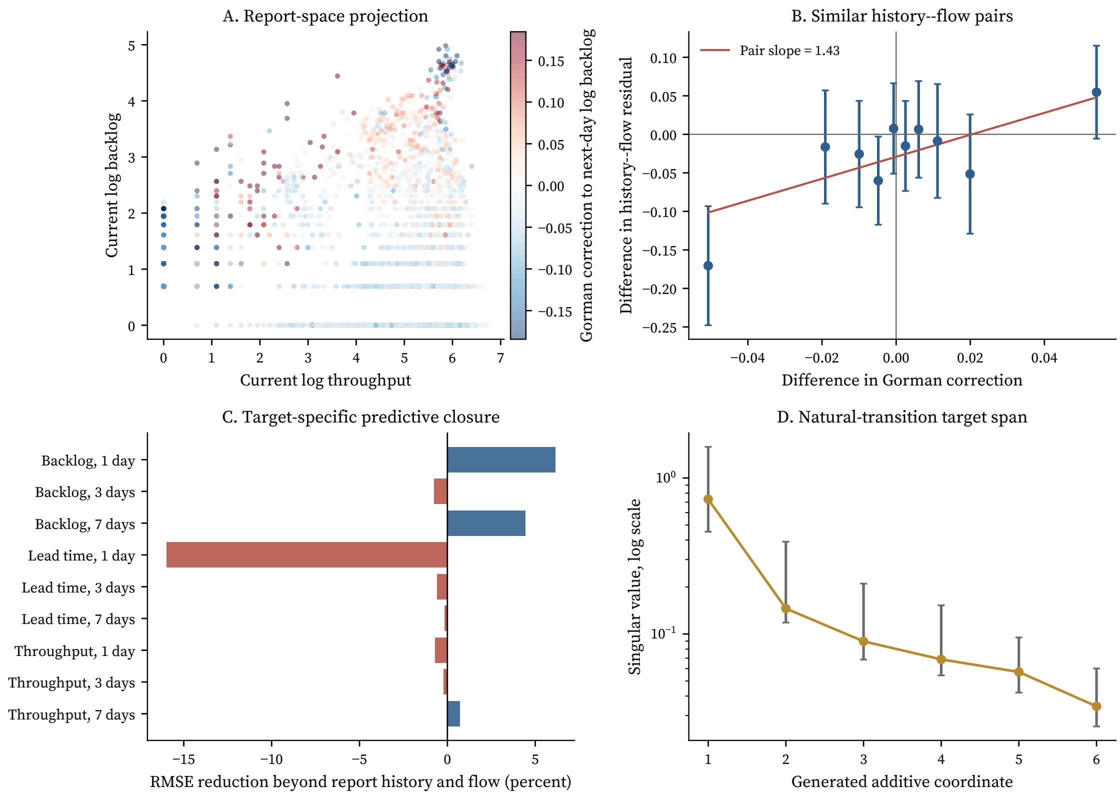

FIGURE 3. Reported operating state and additive internal coordinates

Panel A projects the four-coordinate current report onto log throughput and log backlog for 5,000 held-out region-days; color is the training-estimated Gorman correction to next-day log backlog conditional on report history and accepted flow. Panel B groups the closest half of city-date history–flow-matched pairs into deciles of their Gorman-correction difference; points give mean differences in the benchmark target residual and bars give pointwise 95-percent intervals. Panel C reports out-of-time RMSE reduction from adding the six moments to the history–flow benchmark; positive values indicate lower RMSE. Panel D gives singular values of the additive target span generated by the observed one-day transition; bars are operational-region cluster-bootstrap 95-percent intervals conditional on report history and accepted flow. The LaDe records contain observed operating transitions.

The second design forms pairs matched on report and flow history. Within each city-date cell and target sample, a greedy assignment creates disjoint pairs that minimize Euclidean distance in a standardized 25-coordinate vector. The vector contains the current report, current accepted flow, and four exact-calendar lags of each. For every pair, the observed future-outcome difference is compared with the differences predicted by the history–flow model and the Gorman correction. This is the operational analogue of a fiber comparison: the two regional units have nearby reported histories, the benchmark absorbs region and calendar terms, and their additive internal coordinates may differ.[9]

For next-day log backlog, exact-calendar report history and accepted flow reduce

[9]The matched-pair exercise is predictive and conditions on observed operations. A causal holonomy comparison additionally requires assigned paths and a common terminal probe.

TABLE 2. Predictive closure of the Express report

Target	Horizon	History–flow RMSE	History–flow + moments RMSE	Change (%)
Backlog	1	0.3945	0.3703	6.13
Backlog	3	0.4810	0.4846	-0.74
Backlog	7	0.7111	0.6795	4.44
Lead time	1	10.4928	12.1678	-15.96
Lead time	3	14.9590	15.0482	-0.60
Lead time	7	16.1213	16.1444	-0.14
Throughput	1	0.6074	0.6115	-0.69
Throughput	3	1.0041	1.0063	-0.22
Throughput	7	1.3959	1.3859	0.72

The benchmark contains the current report, current accepted flow, and exact-calendar lags at one, two, three, and seven days. The moment model adds equation (74) and current-report–moment interactions. Change is the percentage reduction in RMSE. Backlog and throughput targets are in logs; lead-time targets are in hours.

RMSE from 0.4228 to 0.3945, and the six additive moments reduce it to 0.3703. The final RMSE increment is 6.13 percent. Squared risk falls 11.89 percent, with a region-and-date two-way-clustered loss-gain t statistic of 2.94 and $p = 0.0038$; the Holm-adjusted p value within the nine-target squared-loss family is 0.0307. Absolute risk falls 8.83 percent, with $t = 6.17$ and an eighteen-comparison Holm-adjusted p value of 1.25×10^{-7}. The complete nine-target comparison in Table 2 shows gains for next-day and seven-day backlog and losses for several throughput and lead-time targets. Additive internal records improve prediction for particular outcomes and horizons; block and horizon decompositions show where the gains arise.

Figure 4 reports checks across time, workload support, location, and future inputs. The next-day gain remains positive in rolling-origin windows, under workload restrictions, in all five cities, and in leave-one-region calculations. In a conditional-alignment exercise, the observed next-day and seven-day backlog gains exceed all 999 within-cell permutation values under both losses. Providing both learners with realized cumulative future arrivals leaves absolute-risk reductions of 8.79, 5.16, and 6.27 percent at one, three, and seven days. Comparisons with nonlinear learners indicate how much of the public-data result is specific to the phase-conditional affine Gorman form. A schema assessment of the five pinned LaDe files finds parcel, courier, location, acceptance, delivery, and date fields, with zero randomized orientation–mirror path cells and zero terminal-probe assignment fields. The LaDe support therefore identifies the reported predictive contrasts and descriptive transition span. Policy holonomy uses the prospective assigned-path and common-probe design described below.

The observed one-day transition provides a descriptive transport operator for the six moments. Iterating the fitted operator on the nine target loadings gives singular

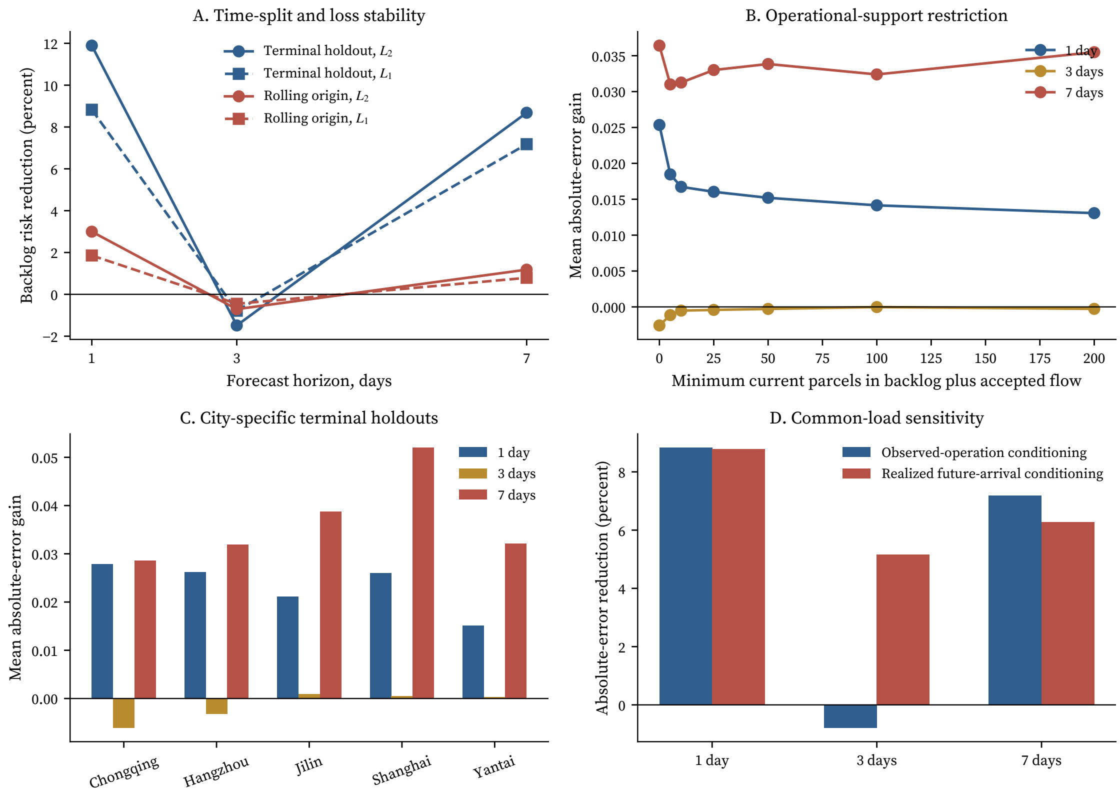

FIGURE 4. Predictive-risk stability across time, support, location, and future load

Panel A gives percentage backlog-risk reductions on the terminal holdout and four disjoint rolling-origin holdouts under squared and absolute loss. Panel B restricts the terminal holdout by minimum current workload and plots absolute-loss gains. Panel C gives city-specific absolute-loss gains. Panel D compares the observed-operation conditioning set with a specification that gives both learners realized cumulative future arrivals through the target horizon. Positive values indicate lower loss after adding the six Gorman coordinates.

values (0.732, 0.146, 0.090, 0.069, 0.057, 0.034). The 99-percent energy rank is four. Across 199 operational-region bootstrap draws, it is three in 28 draws, four in 126, and five in 45. This calculation estimates the target span generated by natural transitions, conditional on report history and accepted flow. Because the LaDe files contain historical operations, the result concerns predictive aggregation under observed transitions. A policy-induced connection and causal holonomy require assigned paths and a common terminal probe.

For prospective assignment, next-day residual dispersion implies a four-arm minimum detectable finite effect of 0.0754 log point, or 0.224 residual standard deviation, when 164 graph-separated regions supply 41 clusters to each arm at two-sided five-percent size and 80-percent power. The eight-cell path–shock design has a 0.1564 log-point finite-difference threshold with 160 balanced clusters. Retention, compliance, dependence, and multiplicity sensitivities accompany these calculations.

An assigned Express experiment could use routing priority, courier incentives, capacity

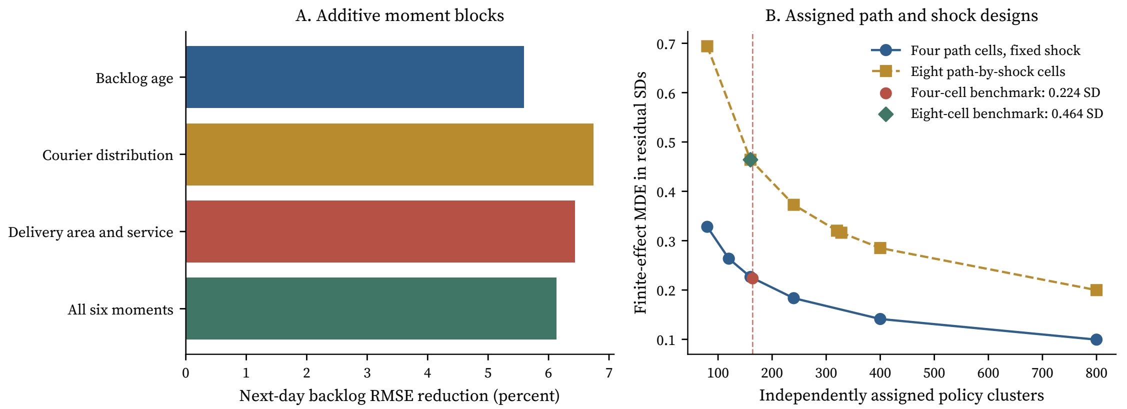

FIGURE 5. Gorman-coordinate blocks and prospective Express assignment

Panel A decomposes the next-day backlog RMSE reduction beyond report history and accepted flow. Each two-coordinate block enters with current-report interactions; the final bar uses all six coordinates. Panel B gives two-sided five-percent, 80-percent-power minimum detectable finite effects in matched-set residual standard deviations. The four-cell curve holds the terminal shock fixed across orientation-by-mirror arms. The eight-cell curve estimates the difference between high-shock and low-shock four-arm contrasts. Division by the shock gap converts this difference into a secant threshold. Markers use the largest balanced designs supported by 166 observed regions. LaDe residuals calibrate outcome dispersion; assigned paths and shocks identify the causal effects.

allocation, and service promises as policy coordinates. Graph-separated operational clusters would enter matched sets of four and receive orientation-by-mirror paths centered on their initial reports. An independently randomized terminal demand, priority, or capacity probe would measure backlog clearance and delivery degradation. A second path–environment design would identify the connection comparative static in equation (A89). Path assignment identifies curvature, terminal-probe assignment identifies the response, and baseline-environment assignment identifies its comparative static. A constructed marketplace benchmark evaluates the assigned-path estimator and bridge.

6. Conclusion

A closed path in reported outcomes can leave the economy in a different state for the next policy. The policy technology and implementation metric specify how a feasible report path moves the underlying state. A terminal probe then measures the part of this displacement that matters for a counterfactual target. Curvature gives the local orientation contrast; mirrored paths and the stated assignment conditions connect it to a feasible experiment.

The transition economy illustrates the mechanism in a solved equilibrium. A closed output–consumption path changes a later consumption response by as much as 6.04 percent of its ordinary value. The LaDe exercise addresses a separate observational question: backlog-age and concentration records improve predictive closure beyond report history and accepted flow. The predictive gains depend on the target and horizon: several throughput and lead-time comparisons favor the report-only benchmark. Its matched residuals calibrate a prospective experiment. Assigned paths and terminal probes would supply the causal comparison required for target holonomy.

The measurement and control results are relative to the target and the admissible technology. Transported target covectors determine which additional measurements must be retained; an integrable additive frame gives a local common-slope representation. Curvature directions determine which defects a closing path can repair. The implementation, rank, and uncertainty conditions therefore remain part of each construction.

The practical implication is to assess a report by the future policy question it is expected to support. Restoring its current value and restoring the economic distinctions behind the next response are separate tasks. The proposed loop designs make their difference measurable across path-space, environment, and finite-loop formulations.

References

Abadir, Karim M. and Gabriel Talmain, "Aggregation, Persistence and Volatility in a Macro Model," *The Review of Economic Studies*, 2002, *69* (4), 749–779.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub, "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 2021, *89* (5), 2375–2408.

Bajari, Patrick, Brian Burdick, Guido W. Imbens, Lorenzo Masoero, James McQueen, Thomas S. Richardson, and Ido M. Rosen, "Experimental Design in Marketplaces," *Statistical Science*, 2023, *38* (3), 458–476.

Bloch, Anthony M., P. S. Krishnaprasad, Jerrold E. Marsden, and Richard M. Murray, "Nonholonomic Mechanical Systems with Symmetry," *Archive for Rational Mechanics and Analysis*, 1996, *136* (1), 21–99.

Gorman, W. M., "Community Preference Fields," *Econometrica*, 1953, *21* (1), 63–80.

___ , "On a Class of Preference Fields," *Metroeconomica*, 1961, *13* (2), 53–56.

___ , "The Structure of Utility Functions," *The Review of Economic Studies*, 1968, *35* (4), 367–390.

Hermann, Robert and Arthur J. Krener, "Nonlinear Controllability and Observability," *IEEE Transactions on Automatic Control*, 1977, *22* (5), 728–740.

Holtz, David, Felipe Lobel, Ruben Lobel, Inessa Liskovich, and Sinan Aral, "Reducing Interference Bias in Online Marketplace Experiments Using Cluster Randomization: Evidence from a Pricing Meta-Experiment on Airbnb," *Management Science*, 2025, *71* (1), 390–406.

Hu, Yingyao and Matthew Shum, "Nonparametric Identification of Dynamic Models with Unobserved State Variables," *Journal of Econometrics*, 2012, *171* (1), 32–44.

Johari, Ramesh, Hannah Li, Inessa Liskovich, and Gabriel Y. Weintraub, "Experimental Design in Two-Sided Platforms: An Analysis of Bias," *Management Science*, 2022, *68* (10), 7069–7089.

Kelly, Scott D. and Richard M. Murray, "Geometric Phases and Robotic Locomotion," *Journal of Robotic Systems*, 1995, *12* (6), 417–431.

Kobayashi, Shoshichi and Katsumi Nomizu, *Foundations of Differential Geometry*, Vol. 1, New York: Wiley, 1963.

Liang, Annie, Xiaosheng Mu, and Vasilis Syrgkanis, "Dynamically Aggregating Diverse Information," *Econometrica*, 2022, *90* (1), 47–80.

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

Perdomo, Juan, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt, "Performative Prediction," in "Proceedings of the 37th International Conference on Machine Learning," Vol. 119 of *Proceedings of Machine Learning Research* PMLR 2020, pp. 7599–7609.

Reiter, Michael, "Solving Heterogeneous-Agent Models by Projection and Perturbation," *Journal of Economic Dynamics and Control*, 2009, *33* (3), 649–665.

Rothenberg, Thomas J., "Identification in Parametric Models," *Econometrica*, 1971, *39* (3), 577–591.

Sobolev, Yaroslav I., Ran Dong, Tsvi Tlusty, Jean-Pierre Eckmann, and Bartosz A. Grzybowski, "Solid-body trajectoids shaped to roll along desired pathways," *Nature*, 2023, *620* (7973), 310–315.

Wu, Keyi, Peng Chen, and Omar Ghattas, "An Offline-Online Decomposition Method for Efficient Linear Bayesian Goal-Oriented Optimal Experimental Design: Application to Optimal Sensor Placement," *SIAM Journal on Scientific Computing*, 2023, *45* (1), B57–B77.

Wu, Lixia, Haomin Wen, Haoyuan Hu, Xiaowei Mao, Yutong Xia, Ergang Shan, Jianbin Zheng, Junhong Lou, Yuxuan Liang, Liuqing Yang, Roger Zimmermann, Youfang Lin, and Huaiyu Wan, "LaDe: The First Comprehensive Last-Mile Delivery Dataset from Industry," *arXiv preprint arXiv:2306.10675*, 2025.