Contemporary Mathematics
http://ojs.wiserpub.com/index.php/CM/

Research Article

A Coupled Dynamic Optimization and Entropy-Regularized PCA Framework for Dynamic Panels

Yun Xu, Yang Andrew Wu*

The School of Professional Education and Executive Development, The Hong Kong Polytechnic University, Hong Kong, China
E-mail: andrew.wu@cpce-polyu.edu.hk

Received: 10 January 2026; **Revised:** 10 February 2026; **Accepted:** 10 March 2026

Abstract: Analyzing the dynamic interplay between multidimensional control parameters and latent state outcomes in complex multi-agent systems poses formidable mathematical and statistical challenges. This paper establishes a coupled mathematical framework that unifies three core components: (i) a dynamic optimization model whose steady-state equilibrium is characterized via a Banach fixed-point argument and proven unique; (ii) an entropy-regularized factor construction procedure with $O(1/k)$ convergence under projected gradient descent; and (iii) a modified System Generalized Method of Moments (GMM) estimator designed to mitigate weak-instrument bias under high treatment persistence. The Karush–Kuhn–Tucker optimality conditions are explicitly derived, and the contraction mapping underlying the Bellman operator is proved in the Banach space of bounded continuous functions. Theoretical properties are validated through Monte Carlo simulations, and empirical feasibility is illustrated using a real-world dynamic panel dataset ($N = 30, T = 13$). The results demonstrate that the integrated methodology produces stable, internally consistent trajectories and outperforms standard benchmarks under structural persistence.

***Keywords*:** dynamic optimization, entropy-weighted factor model, composite index construction, dynamic panel data, system-level factor analysis

MSC: 91G15, 91B76, 62H25, 62P20

1. Introduction

1.1 *Background and motivation*

Quantifying the dynamic response of latent outcome variables to multidimensional control parameters in complex multi-agent systems poses a class of mathematical challenges that arise across a wide spectrum of applied disciplines. When both the treatment and the outcome are latent constructs observed through noisy proxies, and when the system evolves with substantial temporal persistence and unobserved cross-sectional heterogeneity, conventional statistical aggregation rules—static variance-weighted indices, pooled fixed-effect regressions, and standard instrumental-variable estimators—often produce divergent and unreliable inference when applied to latent constructs measured through heterogeneous noisy proxies [1]. These difficulties are structural rather than domain-specific: they stem from the interplay among measurement in multiple noisy dimensions, strong temporal dependence, and latent confounding structures, and they demand an integrated mathematical treatment rather than a sequence of ad hoc empirical steps.

Copyright ©2026 Yang Andrew Wu, et al.
DOI: https://doi.org/10.37256/cm.7320269151
This is an open-access article distributed under a CC BY license
(Creative Commons Attribution 4.0 International License)
https://creativecommons.org/licenses/by/4.0/

A concrete setting that exemplifies these structural challenges is the analysis of regulatory and financial control parameters in multi-sector systems subject to binding emission constraints. In such systems, a central authority deploys a portfolio of coordinated instruments—credit allocation rules, capital-market-access conditions, insurance mandates, and earmarked investment channels—that jointly influence the technological choices, input mixes, and emission trajectories of agents operating in capital-intensive sectors. The control parameter is inherently multidimensional and latent [2, 3]: no single observable variable captures its intensity, and each candidate indicator carries a distinct signal-to-noise ratio that may vary over time [4]. Moreover, the outcome variable of interest—emission intensity—displays strong temporal persistence because physical capital stocks, technology diffusion lags, and regulatory implementation cycles introduce gradual adjustment dynamics [5].

1.2 *The illustrative spatio-temporal testbed*

To anchor the theoretical framework, we project its implications onto a class of complex, high-persistence physical systems governed by structural adjustment costs and multi-channel control bounds—specifically, the resource allocation and emission-intensity evolution within heavy-asset infrastructure networks. In such systems, a coordinator deploys a vector of multi-dimensional adjustments (e.g., restricted allocation rules and investment pipelines) to influence the state paths of heterogeneous agents. This setup serves as an ideal empirical testbed because its mathematical characteristics—near-unit-root temporal persistence, gradual capital adjustment trajectories, and highly noisy indicator matrices—mirror the exact structural boundary conditions under which conventional non-coupled statistical estimators break down. We emphasize that this real-world dataset is employed strictly as an illustrative benchmark; the mathematical consistency and algorithmic convergence of the proposed framework remain independent of the specific empirical domain.

A substantial body of research in multivariate statistics, convex optimization, and panel data econometrics has examined the component problems that the present framework integrates: composite index construction via regularized principal components [2, 3], dynamic panel estimation with internal instruments [5, 6], and constrained dynamic optimization [7, 8]. Yet most existing studies treat these as separate methodological decisions, without formally aligning them within a unified mathematical structure. This raises the central question that motivates the present work: can one develop an integrated mathematical framework in which measurement, identification, and estimation are jointly optimized under a coherent set of structural assumptions, and does such integration generate emergent analytical properties beyond what sequential application of these components would yield?

We emphasize that each component of our framework has been examined in the existing literature. Entropy-regularized Principal Component Analysis (PCA) formulations appear in the sparse PCA literature [4, 9]; restricted-instrument GMM has been proposed as a remedy for weak instruments [10]; and dynamic programming with contraction mapping arguments is a standard tool in economic dynamics [8, 11]. The present paper's contribution is therefore *not* the invention of any single component, but rather the demonstration that coupling them within a coherent mathematical structure—where the optimization model constrains the admissible measurement space, the measurement model conditions the estimation strategy, and the estimation strategy is tuned to the persistence regime implied by the theory—produces analytical properties (stability under persistence, improved instrument relevance, and formal convergence guarantees for the coupled pipeline) that are not obtainable when these steps are performed sequentially and independently. Answering this question requires simultaneously addressing the construction of a time-varying composite index from high-dimensional noisy indicators, the specification of a dynamic model that respects the persistence structure implied by the underlying optimization problem, and the choice of an estimator that remains informative when the treatment is near-unit-root. The remainder of the paper is devoted to building and validating such a framework.

2. Literature review

2.1 *Composite index construction: PCA, regularization, and entropy methods*

Composite indices that aggregate high-dimensional multivariate observations into scalar summaries are central to statistical inference in systems where the construct of interest is latent [2]. Traditional principal component analysis extracts orthogonal factors that maximize explained variance and is widely used for dimensionality reduction in panel and time-series settings [3]. However, standard PCA assigns equal structural importance to all input dimensions and does not incorporate prior knowledge about indicator quality or temporal stability, which can cause instability when signal-to-noise ratios differ across sub-indicators [4].

Sparse and regularized variants of PCA address these limitations by incorporating penalty functions into the factor extraction objective, enabling simultaneous dimensionality reduction and variable selection to yield more interpretable and stable components in high-dimensional settings [9]. The framework of statistical learning with sparsity provides a unifying treatment of regularized M-estimators, including penalized PCA formulations with both convex and non-convex regularizers [12]. Robust PCA formulations further extend the classical model by decomposing the data matrix into low-rank and sparse components, offering resilience to gross outliers and corrupted observations [13].

From an information-theoretic perspective, entropy-based regularization provides a principled alternative to sparsity-inducing penalties. Maximizing Shannon entropy under moment constraints yields a probability distribution that is maximally non-committal with respect to missing information—a property that translates naturally to the construction of diversified composite weights [14]. In the machine learning literature, entropy regularization has been employed to encourage exploration, prevent premature concentration, and improve generalization in high-dimensional learning problems, supported by non-asymptotic statistical guarantees [14, 15]. These information-theoretic considerations motivate the entropy-regularized PCA formulation developed in Section 4.

2.2 *Convex optimization and projected gradient methods*

The theoretical analysis of composite index construction under entropy regularization requires tools from convex optimization. When the objective function is strictly convex and the constraint set is a closed convex subset of a Euclidean space, the existence and uniqueness of a global minimizer are guaranteed, and the Karush–Kuhn–Tucker conditions provide a complete characterization of optimality [7]. First-order methods, including projected gradient descent and its accelerated variants, offer computationally efficient algorithms for large-scale constrained convex optimization problems with well-characterized convergence rates [16].

Proximal algorithms extend the gradient descent framework to non-smooth objectives by decomposing the optimization problem into smooth and proximable components, enabling efficient handling of regularization terms and simplex constraints through closed-form proximal operators [17]. The modern theory of convex optimization provides rigorous convergence guarantees for projected gradient methods under various step-size regimes, including the $O(1/k)$ sublinear rate for general smooth convex objectives and linear convergence under strong convexity [18]. These results form the mathematical basis for the existence, uniqueness, and algorithmic convergence properties established in Proposition 3.

2.3 *Dynamic optimization and fixed-point characterization of equilibrium*

The structural microfoundation of the empirical framework relies on dynamic optimization theory. In discrete-time infinite-horizon settings, Bellman's principle of optimality reduces the sequential decision problem to a functional equation whose solution is characterized as a fixed point of a contraction mapping on an appropriately defined Banach space [8]. Under standard regularity conditions—including boundedness, continuity, and a discount factor strictly less than unity—the contraction mapping theorem guarantees the existence of a unique value function and an associated time-invariant policy function, which in turn defines a unique steady-state equilibrium [11]. These results underpin the equilibrium characterization developed in Lemma 1.

The continuity of the steady state with respect to structural parameters follows from the theorem of the maximum, which ensures that the optimal policy correspondence is upper hemicontinuous in parameters under continuity of the

objective and constraint correspondence [8]. Local stability of the steady state can be established via linearization of the Euler equation around the equilibrium point, with convergence guaranteed when the eigenvalues of the linearized system lie within the unit circle. These dynamic optimization foundations provide the structural link between control parameters and long-run outcomes that justifies the dynamic panel specification in subsection 4.2.

2.4 *Dynamic panel data models: GMM estimation and weak instruments*

Dynamic panel data models form the statistical backbone for systems exhibiting temporal dependence, unobserved heterogeneity, and simultaneously determined variables. The Generalized Method of Moments (GMM) framework, operationalized through the difference and system GMM estimators, addresses the Nickell bias that arises when lagged dependent variables coexist with fixed effects, by exploiting lagged levels and differences as internal instruments under sequential exogeneity conditions [5]. The analytical properties of these estimators—including consistency, asymptotic normality, and finite-sample bias—have been extensively characterized in the panel data literature [6, 19].

A critical challenge arises when the independent variable is highly persistent, approaching a near-unit-root process. In such settings, lagged differences become weak instruments for current levels, leading to substantial finite-sample bias, inflated variance, and unreliable inference [20]. The weak-instrument problem has been systematically analyzed in the instrumental variables literature, with diagnostic tools developed to detect and mitigate its consequences [21, 22]. Within the specific context of system GMM estimation, the persistence of the treatment variable degrades the relevance of the differenced-equation instruments, motivating modified instrument sets and continuously updated weighting matrices to improve finite-sample performance [10].

The choice of instruments and the associated specification tests are therefore not merely diagnostic formalities but constitute an integral part of the identification strategy. Extensive Monte Carlo evidence demonstrates that restricted instrument sets and appropriate weighting choices substantially reduce bias in panels with persistent treatments and moderate time dimensions [19]. These considerations directly motivate the modified system GMM estimator developed in subsection 4.2.

2.5 *Two-stage estimation, generated regressors, and bootstrap inference*

The integrated framework proposed in this paper involves a two-stage estimation procedure in which a composite index constructed via entropy-regularized PCA serves as an independent variable in a dynamic panel regression. This structure introduces a generated regressor problem: the first-stage estimation uncertainty propagates into the second-stage moment conditions and may bias standard errors if not properly accounted for [23]. Classical approaches to generated regressors rely on asymptotic corrections to the variance–covariance matrix, which can be derived under regularity conditions on the convergence rate of the first-stage estimator.

More recent work on post-selection and post-regularization inference has developed robust methods for quantifying uncertainty when the regressor is constructed through a data-dependent procedure [24]. These methods recognize that model selection or regularization introduces additional variability that standard plug-in variance estimators fail to capture. In parallel, the double/debiased machine learning framework provides a general approach to inference in the presence of high-dimensional nuisance parameters, offering formal guarantees for treatment effect estimation when covariates are selected or constructed from high-dimensional data [25].

For practitioners, residual-based bootstrap procedures offer a computationally accessible alternative to analytical corrections. By resampling residuals from both stages of a two-stage estimation while preserving the dependence structure across units and time periods, bootstrap methods produce valid confidence intervals without requiring the explicit derivation of the asymptotic variance of the generated regressor [24]. This approach is adopted in the present study to address the inference challenges arising from the two-stage entropy-regularized PCA–system GMM estimation pipeline.

3. Theoretical model

3.1 *Model setup*

Consider a representative agent operating in unit i at time t. The agent produces output Y_{it} using a combination of physical capital embodied in an advanced technology and a resource input whose consumption generates a regulated by-product. Production and resource-use decisions are made under operational and regulatory constraints shaped by the local control-parameter environment.

Let $K_{it} \geq 0$ denote investment in the advanced technology. Let $E_{it} \geq 0$ denote the resource input whose consumption directly generates the regulated by-product. The intensity of the control parameter in unit i at time t is summarized by an exogenous structural parameter θ_{it}, capturing the availability, affordability, and regulatory support of the instruments that constitute the control regime. The agent takes θ_{it} as given.

The production technology is represented by a twice continuously differentiable production function:

$$Y_{it} = A_i \cdot g(K_{it}, E_{it}),$$

where $A_i > 0$ denotes unit-specific productivity, and $g(\cdot)$ is increasing and concave in both arguments, satisfying standard Inada conditions.

3.2 *Objective function and constraints*

The agent minimizes total cost subject to production and regulatory constraints. The static optimization problem is given by:

$$\min_{K_{it}, E_{it}} C_{it} = c_K K_{it} + c_E E_{it} + \tau(\theta_{it}) f(E_{it}).$$

The optimization is subject to the following constraints:

$$Y_{it} = A_i \cdot g(K_{it}, E_{it}),$$

$$E_{it} \leq \bar{E} - \phi\theta_{it},$$

$$K_{it} \geq \kappa(\theta_{it}).$$

3.3 *Lagrangian formulation and KKT optimality conditions*

The Lagrangian associated with the static problem is:

$$L = c_K K_{it} + c_E E_{it} + \tau(\theta_{it}) f(E_{it}) + \lambda_1(\bar{E} - \phi\theta_{it} - E_{it}) + \lambda_2(K_{it} - \kappa(\theta_{it})),$$

where the Lagrange multipliers $\lambda_1, \lambda_2 \in \mathbb{R}_+$ represent the shadow prices associated with the emission and investment constraints, respectively.

The complete Karush–Kuhn–Tucker (KKT) optimality conditions for the primal-dual pair require:

$$\begin{aligned}
&\nabla_E L = c_E + \tau(\theta_{it}) f'(E_{it}) - \lambda_1 = 0, \\
&\nabla_K L = c_K + \lambda_2 = 0, \\
&\lambda_1(\bar{E} - \phi\theta_{it} - E_{it}) = 0, \\
&\lambda_2(K_{it} - \kappa(\theta_{it})) = 0.
\end{aligned} \tag{1}$$

Since $g(K,E)$ is strictly concave and the constraint set

$$\Omega(\theta) = \{(K,E) \in \mathbb{R}^2_+ : E \leq \bar{E} - \phi\theta,\ K \geq \kappa(\theta)\},$$

is a closed convex subset of $\mathbb{R}^2$, the KKT conditions are both necessary and sufficient for the existence of a unique global minimizer K^*, E^* [7, 18].

3.4 *Comparative statics*

Proposition 1 An increase in the control parameter reduces optimal resource input, $\partial E^*_{it}/\partial\theta_{it} < 0$.

Proposition 2 An increase in the control parameter raises optimal technology investment, $\partial K^*_{it}/\partial\theta_{it} > 0$, under the structural assumption that technology and resource inputs are technical complements ($g_{KE} \geq 0$).

3.5 *Dynamic extension*

To capture adjustment frictions, technology investment evolves as:

$$K_{it+1} = (1-\delta)K_{it} + I_{it}.$$

The dynamic optimization problem is:

$$V(K_{it}) = \min\left\{C_{it} + \frac{\gamma}{2}(I_{it} - I_{it-1})^2 + \beta\mathbb{E}_t[V(K_{it+1})]\right\}.$$

Standard dynamic programming results ensure that, under the stated regularity conditions, the Bellman equation defines a contraction mapping and admits a unique value function [8, 11].

3.6 *Equilibrium characterization*

The dynamic optimization model admits a well-defined steady-state equilibrium under mild regularity conditions, which establishes the structural foundation for the empirical specification.

Lemma 1 (Existence and Uniqueness of Steady-State Equilibrium) Assume that $g(K,E)$ is strictly concave, $f(E)$ is convex, the cost function $\tau(\theta)$ is positive and continuous, and the constraint function $\kappa(\theta)$ is continuous. Let the adjustment cost parameter $\gamma > 0$ and the discount factor $\beta \in (0,1)$. Then:

(i) There exists a unique steady-state equilibrium (K^*, E^*) satisfying the Euler equation derived from the dynamic optimization problem.

(ii) The steady-state technology investment K^* and resource input E^* are continuous functions of the control parameter θ.

(iii) The steady state is locally stable: for any initial K_{i0} sufficiently close to K^*, the optimal path converges to (K^*, E^*).

Proof Sketch. The Bellman equation defines a contraction mapping on the space of bounded continuous functions under the sup-norm for $\beta \in (0,1)$, ensuring a unique fixed point $V(K)$. The Euler equation, together with the envelope condition, yields a system of two equations in (K^*, E^*). Strict concavity of g and convexity of f guarantee that the system has a unique solution. The constraint set $\Omega(\theta) = \{(K,E) \in \mathbb{R}^2_+ : E \leq \bar{E} - \phi\theta,\ K \geq \kappa(\theta)\}$ is compact for each θ (it is closed and bounded under the maintained assumptions that $\bar{E} < \infty$ and $\kappa(\theta)$ is finite). Continuity of the constraint correspondence in θ follows from the continuity of $\kappa(\theta)$ and the linear structure of the emission bound, satisfying the conditions of the theorem of the maximum. Continuity of (K^*, E^*) in θ follows from the theorem of the maximum, as the constraint correspondence is continuous in θ. Local stability follows from standard linearization arguments around the steady state, where the eigenvalues of the linearized system have modulus less than unity under $\gamma > 0$. A complete proof is provided in Appendix A. □

Lemma 1 formalizes the structural link between the control parameter and the long-run equilibrium of the system, providing the microfoundation for the dynamic panel specification in subsection 4.2. In particular, the continuity of the steady state in θ justifies treating the composite index as a continuously-varying treatment intensity, while the stability property supports the use of lagged dependent variables to capture gradual adjustment toward equilibrium.

This section develops the integrated mathematical framework in which measurement, identification, and estimation are explicitly aligned. Motivated by the theoretical optimization model introduced above, the framework ensures that both the measurement of the control parameter and the statistical treatment of its effects are consistent with the structural mechanisms—technology adoption and constraint relaxation—through which the control parameter influences the outcome.

The proposed framework addresses three interrelated methodological challenges that frequently arise in the analysis of complex multi-agent systems. First, control parameters are inherently multidimensional and latent, requiring principled methods for composite index construction that account for heterogeneous information content and measurement error across sub-indicators. Second, treatment variables and outcomes often display strong temporal persistence, rendering standard panel estimators vulnerable to weak-instrument bias and finite-sample distortions. Third, empirical identification must be grounded in an explicit structural logic to ensure coherence among theory, measurement, and estimation. By integrating entropy-regularized factor analysis with a modified dynamic panel estimator, this section provides a unified toolkit for addressing these challenges.

Mapping from theory to empirics. A central feature of the integrated framework is the explicit correspondence between structural parameters and empirical quantities. In the theoretical model (Section 3), θ_{it} is an exogenous control parameter that shapes the agent's feasible set through the constraint functions $\kappa(\theta)$ and $\tau(\theta)$; the agent's optimal choices K^*_{it} and E^*_{it} are deterministic functions of θ_{it} and structural parameters. In the empirical specification, $\hat{\theta}_{it}$ is a *statistically constructed* composite index derived from noisy indicators via entropy-regularized PCA, and enters the outcome equation as a generated regressor. The theoretical steady-state mapping from θ to outcomes (Lemma 1) is captured empirically through the dynamic panel specification $Y_{it} = \delta Y_{i,t-1} + \beta\hat{\theta}_{it} + \gamma' X_{it} + \mu_i + \lambda_t + \varepsilon_{it}$, where the lagged dependent variable proxies for the gradual capital adjustment toward the steady state characterized in the optimization model. While the structural parameters of the production function $g(K,E)$ and adjustment cost γ are not individually identified in the reduced-form regression (they are subsumed into the composite coefficient β), the sign and dynamic structure predicted by the theory—that higher θ reduces E^* and raises K^*, with net effect on Y determined by the relative magnitude of these two channels—provide testable restrictions on the empirical specification. Specifically, the theory predicts a negative association between θ and Y when the emission-reduction effect dominates the technology-adoption effect. This prediction is examined in the empirical validation exercise of Section 5.

4. Proposed mathematical framework

4.1 *Entropy-regularized PCA for dynamic composite indices*

4.1.1 *Problem formulation*

Let $X_{it} = (x_{i1t}, x_{i2t}, \ldots, x_{iJt}) \in \mathbb{R}^J$ denote a vector of observed sub-indicators for unit i at time t; the specific identity of the indicators is inessential to the mathematical development. Each observed indicator captures a specific dimension of the latent construct, yet none alone provides a complete or noise-free measure of the underlying intensity variable.

Formally, the observed indicators are assumed to be generated by an underlying latent intensity variable θ_{it}^*, contaminated by indicator-specific noise and scaling heterogeneity. The empirical objective is therefore to construct a composite index θ_{it} that provides a statistically efficient and information-robust approximation to θ_{it}^* while remaining interpretable and stable over time.

The composite index is defined as a linear aggregation:

$$\theta_{it} = X_{it} W$$

where $W = (w_1, \ldots, w_J)'$ is a vector of non-negative weights satisfying the simplex constraint $\sum_{j=1}^{J} w_j = 1,\ w_j \geq 0$. Traditional principal component analysis determines W by maximizing the variance of XW, thereby extracting the dominant common component across indicators. However, pure variance maximization tends to overweight indicators with large dispersion, even when such dispersion reflects noise rather than information.

To overcome this limitation, this study proposes an entropy-regularized PCA formulation that explicitly penalizes excessive concentration of weights. The index construction problem is formulated as:

$$\min_{W} \left\{ -\mathrm{Var}(XW) + \lambda \sum_{j=1}^{J} w_j \ln w_j \right\}$$

subject to $w_j \geq 0$, $\sum_{j=1}^{J} w_j = 1$, where $\lambda \geq 0$ is a regularization parameter controlling the strength of entropy penalization.

The first term reflects the standard PCA objective of maximizing explained variance, while the entropy term promotes diversification across indicators by penalizing highly uneven weight distributions. This formulation embeds information-theoretic considerations directly into factor extraction, thereby improving robustness in the presence of heterogeneous indicator quality [2, 12].

4.1.2 *Theoretical properties*

The entropy-regularized PCA framework encompasses several commonly used index construction methods as special or limiting cases. The following results characterize the behavior of the optimal weight vector as the regularization parameter varies.

Theorem 1 As $\lambda \to 0$, the entropy penalty vanishes and the solution $W(\lambda)$ converges to the loading vector associated with the first principal component of X:

$$\lim_{\lambda \to 0} W(\lambda) = W_{\mathrm{PCA}}.$$

Theorem 2 As $\lambda \to \infty$, the entropy term dominates the objective function and the solution $W(\lambda)$ converges to uniform weights $w_j = 1/J$ for all j:

$$\lim_{\lambda\to\infty} w_j = \frac{1}{J}.$$

These limiting results demonstrate that the proposed method provides a continuous spectrum between variance-dominated factor extraction and equal-weight aggregation. Intermediate values of λ allow the data to determine the appropriate balance between efficiency and robustness, thereby avoiding both overfitting to noisy indicators and excessive smoothing that obscures meaningful variation.

Proposition 3 (Existence, Uniqueness, and Convergence of the Entropy-Regularized PCA Solution) For any $\lambda > 0$, the entropy-regularized PCA problem

$$\min_{W\in\Delta} \mathscr{L}_\lambda(W) = -\operatorname{Var}(XW) + \lambda \sum_{j=1}^{J} w_j \ln w_j,$$

where $\Delta = \{W \in \mathbb{R}^J : w_j \geq 0,\ \sum_j w_j = 1\}$ denotes the probability simplex, satisfies the following properties:

(i) **Existence.** A global minimizer $W^*(\lambda)$ exists. The objective $\mathscr{L}_\lambda$ is proper, lower semicontinuous on Δ, and coercive (extended by $+\infty$ outside Δ), guaranteeing existence by the Weierstrass theorem.

(ii) **Uniqueness.** The minimizer $W^*(\lambda)$ is unique. The variance term $-\operatorname{Var}(XW)$ is convex in W (as it is a negative semidefinite quadratic form on Δ), and the entropy term $\sum w_j \ln w_j$ is strictly convex on the interior of Δ. Hence $\mathscr{L}_\lambda$ is strictly convex on the interior of Δ, and the minimizer is unique.

(iii) **Algorithmic convergence.** Consider the projected gradient iteration $W^{(k+1)} = \Pi_\Delta\left(W^{(k)} - \alpha_k \nabla\mathscr{L}_\lambda(W^{(k)})\right)$, where Π_Δ denotes Euclidean projection onto Δ. Under either (a) a constant step size $\alpha_k = \alpha \leq 1/L$, where L is the Lipschitz constant of $\nabla\mathscr{L}_\lambda$, or (b) diminishing step sizes satisfying $\sum \alpha_k = \infty$ and $\sum \alpha_k^2 < \infty$, the iterates converge to the unique minimizer: $\lim_{k\to\infty} W^{(k)} = W^*(\lambda)$.

(iv) **Convergence rate.** With a constant step size $\alpha \leq 1/L$, the objective values satisfy $\mathscr{L}_\lambda(W^{(k)}) - \mathscr{L}_\lambda(W^*) \leq \frac{\|W^{(0)} - W^*\|^2}{2\alpha k}$, establishing an $O(1/k)$ sublinear convergence rate. Under the additional condition that the sample covariance matrix Σ_X has a spectral gap between its first and second eigenvalues, linear (geometric) convergence is attained.

Proof Sketch. (i) Δ is compact in $\mathbb{R}^J$ and $\mathscr{L}_\lambda$ is continuous on Δ (extending the entropy term by continuity at the boundary, where $\lim_{w\to 0^+} w \ln w = 0$). Existence follows from the extreme value theorem. (ii) The Hessian of the entropy term is $\operatorname{diag}(1/w_1, \ldots, 1/w_J)$, which is positive definite on the interior of Δ. The Hessian of $-\operatorname{Var}(XW) = -W'\Sigma_X W$ is $-2\Sigma_X$. Since Σ_X is positive semidefinite, $-\Sigma_X$ is negative semidefinite. For $\lambda > 0$, the sum produces a positive definite Hessian on the interior, establishing strict convexity. (iii) On any compact subset of the interior where $w_j \geq \varepsilon > 0$, $\nabla\mathscr{L}_\lambda$ is Lipschitz with constant $L_\varepsilon = 2\|\Sigma_X\| + \lambda/\varepsilon$. Standard projected gradient convergence theory applies. (iv) The $O(1/k)$ rate follows from the standard analysis of gradient descent for smooth convex objectives. The linear rate under a spectral gap follows from strong convexity of $\mathscr{L}_\lambda$ restricted to the subspace orthogonal to the leading eigenvector. Full technical details are deferred to Appendix A. □

Algorithm 1 Projected gradient descent for entropy-regularized PCA

Input: Sample covariance matrix $\Sigma_X \in \mathbb{R}^{J\times J}$, regularization parameter $\lambda > 0$, step size $\alpha_k \leq 1/L$, convergence tolerance $\varepsilon > 0$.

Output: Unique optimal weight vector $W^* \in \Delta$.

1: Initialize $W^{(0)} \in \Delta$ uniformly, i.e., $w_j^{(0)} = 1/J$ for $j = 1, \ldots, J$. Set $k = 0$.

2: **repeat**

3: Compute the gradient of the unregularized objective: $\nabla f(W^{(k)}) = -2\Sigma_X W^{(k)}$.

Algorithm 1 (cont.)

4: Compute the gradient of the Shannon entropy penalty: $\nabla h(W^{(k)}) = \lambda \left(1 + \ln w_1^{(k)}, \ldots, 1 + \ln w_J^{(k)}\right)'$.
5: Assemble full gradient: $\nabla \mathscr{L}_\lambda(W^{(k)}) = \nabla f(W^{(k)}) + \nabla h(W^{(k)})$.
6: Descent step: $\tilde{W}^{(k+1)} = W^{(k)} - \alpha_k \nabla \mathscr{L}_\lambda(W^{(k)})$.
7: Euclidean projection onto the probability simplex: $W^{(k+1)} = \Pi_\Delta\left(\tilde{W}^{(k+1)}\right)$.
8: Increment iterations: $k \leftarrow k+1$.
9: **until** $\|W^{(k)} - W^{(k-1)}\|_2 < \varepsilon$
10: **return** $W^* = W^{(k)}$.

Numerical safeguard at the simplex boundary. The gradient of the Shannon entropy penalty, $\nabla h(W) = \lambda(1 + \ln w_1, \ldots, 1 + \ln w_J)'$, has an unbounded derivative as any $w_j \to 0^+$. To maintain numerical stability, we implement a practical floor: after each Euclidean projection onto the simplex, any component $w_j < \varepsilon$ is set to ε and the weight vector is renormalized to satisfy $\sum_j w_j = 1$. This ensures that the logarithm is always evaluated at arguments bounded away from zero, preserving the Lipschitz property of the gradient on the restricted domain $\Delta_\varepsilon = \{W \in \Delta : w_j \geq \varepsilon, \forall j\}$. In all simulations and empirical exercises, we set $\varepsilon = 10^{-8}$. The floor does not affect the theoretical convergence guarantees, as the unique minimizer $W^*(\lambda)$ for $\lambda > 0$ lies in the strict interior of the simplex (the entropy gradient diverges at the boundary, precluding boundary solutions). Proposition 3 establishes that the proposed index construction procedure is mathematically well-posed: a unique optimal weight vector exists, and a simple projected gradient algorithm converges with guaranteed rates. These properties are essential for the procedure to serve as a reliable measurement component within the broader integrated framework.

Beyond these limiting cases, the entropy-regularized PCA exhibits favorable stability properties in finite samples. By discouraging extreme weights, the method reduces sensitivity to outliers and improves temporal smoothness of the resulting index, which is particularly important in dynamic panel settings where index volatility can propagate into the estimation stage.

4.1.3 *Information-theoretic and bayesian interpretation*

The entropy regularization term $\lambda \sum w_j \ln w_j$ admits a rigorous interpretation within both Information Theory and Bayesian statistics. From a Bayesian perspective, maximizing the entropy-regularized objective function is equivalent to finding the Maximum A Posteriori (MAP) estimate of the weights W under a specific prior distribution. Specifically, the entropy term corresponds to imposing a Dirichlet prior on the weight vector W. If we assume a symmetric Dirichlet prior $\text{Dir}(\alpha, \ldots, \alpha)$ with $\alpha < 1$, the log-prior density is proportional to $(\alpha - 1)\sum \ln w_j$. By setting the regularization parameter λ appropriately, our objective function aligns with the posterior maximization where the prior encourages sparsity or diversification depending on α. In our formulation, the entropy penalty effectively assumes a prior that penalizes extreme concentration of weights, thereby incorporating a preference for robustness against indicator-specific noise before observing the data.

Comparison with Tikhonov (L2) Regularization: Unlike Tikhonov regularization, Shannon Entropy maximization addresses the "noise-to-signal" problem differently: L2 regularization is sensitive to the scale of weights and tends to shrink large weights towards zero but does not enforce the simplex constraint naturally without additional projection. Entropy regularization is intrinsic to the probability simplex, ensuring weights remain non-negative and sum to one. Noise Suppression Mechanism: Tikhonov regularization assumes Gaussian noise in the parameters and shrinks coefficients uniformly. In contrast, maximizing Shannon Entropy under variance constraints explicitly targets the distributional uncertainty. It prevents the model from assigning near-zero weights to potentially informative indicators due to transient noise while avoiding overfitting to a single high-variance noisy indicator. Mathematically, while L2 minimizes the Euclidean distance from the origin, Entropy maximization minimizes the Kullback-Leibler divergence from a uniform distribution, providing a more principled approach to balancing signal extraction with weight diversification in high-dimensional, noisy settings.

4.2 *Modified system GMM for persistent treatment variables*

4.2.1 *Persistence and weak instrument problem*

The theoretical model in Section 3 implies gradual adjustment in technology adoption and resource use, leading to strong persistence in both dependent and independent variables. Accordingly, the empirical specification takes the form:

$$Y_{it} = \delta Y_{i,t-1} + \beta \theta_{it} + \gamma' X_{it} + \mu_i + \lambda_t + \varepsilon_{it}$$

where Y_{it} denotes the outcome intensity, θ_{it} is the composite index, X_{it} is a vector of controls, μ_i captures time-invariant heterogeneity, and λ_t captures common time effects.

When θ_{it} is highly persistent, such that $\theta_{it} \approx \rho \theta_{i,t-1}$ with ρ close to unity, standard System GMM estimators suffer from weak instrument problems. Lagged differences of θ_{it} become weakly correlated with current levels, resulting in biased estimates and unreliable inference in finite samples [5, 10, 20]:

$$\theta_{it} \approx \rho \theta_{i,t-1}, \quad \rho \to 1.$$

4.2.2 *Modified moment conditions and estimation strategy*

To address these issues, the study employs a modified System GMM approach that restricts the instrument set and adopts a continuously updated weighting matrix. Specifically, the estimator selectively uses lagged levels and differences that satisfy relevance conditions under persistence, thereby mitigating weak instrument bias and instrument proliferation.

The moment conditions take the form:

$$\mathbb{E}[Z_{it}' \varepsilon_{it}] = 0,$$

where the instrument matrix Z_{it} is constructed as follows. For the differenced equation $\Delta y_{it} = \delta \Delta y_{i,t-1} + \beta \Delta \theta_{it} + \gamma' \Delta X_{it} + \Delta \varepsilon_{it}$, we use lagged levels $y_{i,t-2}, y_{i,t-3}$ and $\theta_{i,t-2}, \theta_{i,t-3}$ as instruments, restricting to two lags rather than the full history to mitigate instrument proliferation. For the level equation $y_{it} = \delta y_{i,t-1} + \beta \theta_{it} + \gamma' X_{it} + \mu_i + \lambda_t + \varepsilon_{it}$, we use lagged differences $\Delta y_{i,t-1}$ and $\Delta \theta_{i,t-1}$ as instruments, following the standard System GMM architecture but with a reduced cardinality. The "collapsed" instrument matrix stacks moment conditions with a single column per instrument lag distance rather than one column per time period, substantially reducing the instrument count. Formally, for unit i the instrument matrix takes the block-diagonal structure

$$Z_i = \begin{pmatrix} Z_i^{\Delta} & 0 \\ 0 & Z_i^{L} \end{pmatrix},$$

where Z_i^{Δ} contains lagged levels for the differenced equation and Z_i^{L} contains lagged differences for the level equation. The continuously updated weighting matrix is computed as $W_{NT}(\psi) = \left(\frac{1}{NT} \sum_i Z_i' \hat{u}_i(\psi) \hat{u}_i(\psi)' Z_i \right)^{-1}$ and is iterated to convergence (tolerance 10^{-6}, maximum 50 iterations). Under standard regularity conditions, the resulting estimator is consistent and asymptotically normal, while exhibiting improved finite-sample properties relative to conventional System GMM.

4.2.3 *Efficiency and identification considerations*

Relative to fixed-effects estimators, the modified GMM approach explicitly accounts for dynamic dependence and latent confounding. Compared with standard System GMM, efficiency gains arise from improved instrument relevance and reduced bias, particularly in panels characterized by persistent treatments and limited time dimensions. We acknowledge that using the estimated index $\hat{\theta}_{it}$ as an independent variable in the System GMM framework introduces a "generated independent variable" problem, where the estimation uncertainty from the first-stage entropy-regularized PCA is ignored in the second stage, potentially biasing standard errors. To address this without relying solely on asymptotic arguments regarding the convergence rate of factor estimators, we implement a residual-based Bootstrap procedure tailored for two-stage dynamic panel models: We generate $B = 1000$ bootstrap samples by resampling residuals from both the PCA construction stage and the GMM estimation stage, preserving both the cross-sectional clustering (units are resampled as blocks to maintain within-unit dependence) and the time-series dependence structure (residual blocks within each unit preserve the serial correlation pattern). Given the panel dimensions ($N = 30$, $T = 13$) and the moderate-to-high persistence of the treatment variable, we use a block size equal to the full time-series length for each unit, which is the natural choice when T is small and the primary concern is cross-sectional dependence. For larger panels, a moving-block bootstrap with block length determined by the autocorrelation structure would be appropriate. For each bootstrap sample b, we re-construct the composite index $\hat{\theta}_{it}^{b}$ and re-estimate the System GMM parameters. We construct empirical confidence intervals for the coefficient β based on the percentile distribution of the bootstrap estimates [23, 24].

4.3 *Integrated estimation procedure and identification logic*

The proposed framework integrates structural modeling, index construction, and statistical estimation into a coherent identification strategy. The optimization model specifies the channels through which the control parameter affects the outcome, the entropy-regularized PCA operationalizes the control parameter as a latent multidimensional construct, and the dynamic panel estimator identifies causal effects under persistence and latent confounding.

Identification relies on rank conditions of the instrument matrix and exclusion restrictions implied by lag structures. Unit fixed effects absorb unobserved heterogeneity, while time effects capture common shocks, ensuring consistency with the theoretical setup.

4.4 *Monte Carlo simulation study*

To assess the finite-sample performance of the proposed framework, a Monte Carlo simulation study is conducted based on a Data-Generating Process (DGP) derived from the theoretical model in Section 3. The DGP is specified as follows.

4.4.1 *DGP*

Latent treatment intensity. The latent treatment variable evolves according to a first-order autoregressive process:

$$\theta_{it}^{*} = \rho_{\theta} \cdot \theta_{i,t-1}^{*} + \eta_{it}, \quad \eta_{it} \sim \mathcal{N}(0, \sigma_{\eta}^{2}), \tag{2}$$

with $\theta_{i0}^{*} \sim \mathcal{N}(0, \sigma_{\eta}^{2}/(1-\rho_{\theta}^{2}))$ to ensure stationarity. The persistence parameter $\rho_{\theta} \in \{0.5, 0.8, 0.95\}$ governs the degree of temporal dependence.

Observed indicators. The latent intensity is measured through J noisy indicators:

$$X_{ijt} = \theta_{it}^{*} + \varepsilon_{ijt}, \quad \varepsilon_{ijt} \sim \mathcal{N}(0, \sigma_{\varepsilon}^{2}), \quad j = 1, \ldots, J, \tag{3}$$

where the signal-to-noise ratio is defined as $\mathrm{SNR} = \mathrm{Var}(\theta^{*})/\sigma_{\varepsilon}^{2} \in \{1, 2, 5\}$.

Dynamic outcome. Consistent with the structural model in Lemma 1, the outcome variable follows:

$$Y_{it} = \delta Y_{i,t-1} + \beta \theta_{it}^{*} + \mu_i + \lambda_t + u_{it}, \quad u_{it} \sim \mathcal{N}(0, \sigma_u^2), \tag{4}$$

where $\mu_i \sim \mathcal{N}(0, \sigma_\mu^2)$ captures unit-specific unobserved heterogeneity and $\lambda_t \sim \mathcal{N}(0, \sigma_\lambda^2)$ captures common time shocks. The true causal parameter is set to $\beta_0 = -0.30$, and the persistence of the outcome is governed by $\delta \in \{0.3, 0.6, 0.85\}$.

Design matrix. The full factorial design varies $(N, T) \in \{(30,10), (60,10), (30,20)\}$, $\rho_\theta \in \{0.5, 0.8, 0.95\}$, SNR $\in \{1, 2, 5\}$, and $\delta \in \{0.3, 0.6, 0.85\}$. For each configuration, $R = 500$ independent replications are drawn.

4.4.2 *Theory-consistent DGP (structural extension)*

To verify that the estimator can recover structural parameters when the data are generated by the theoretical model of Section 3, we construct a second DGP that embeds the full structural logic. The agent's problem is parameterized with production function $g(K, E) = K^\alpha E^\nu$ ($\alpha = 0.3$, $\nu = 0.4$), adjustment cost $\gamma = 2.0$, depreciation $\delta = 0.08$, discount factor $\beta = 0.95$, emission penalty $\tau(\theta) = \tau_0 \exp(-\eta\theta)$ with $\tau_0 = 0.5$, $\eta = 1.0$, technology constraint $\kappa(\theta) = \kappa_0 + \kappa_1 \theta$ with $\kappa_0 = 1.0$, $\kappa_1 = 0.8$, and emission cap $\bar{E} = 5.0$, $\phi = 0.5$. The latent control parameter θ_{it} follows the AR(1) process in (2).

At each period t, the agent solves:

$$\min_{K_{it}, E_{it}} c_K K_{it} + c_E E_{it} + \tau(\theta_{it}) f(E_{it}) \quad \text{s.t.} \quad E_{it} \leq \bar{E} - \phi\theta_{it},\ K_{it} \geq \kappa(\theta_{it}),$$

with the KKT conditions determining whether constraints bind. The outcome is generated indirectly through the structural production and accumulation channels:

$$Y_{it} = A_i \cdot g(K_{it}, E_{it}) + \varepsilon_{it}^{Y},$$

where K_{it} evolves via $K_{i,t+1} = (1-\delta)K_{it} + I_{it}$ and investment I_{it} is chosen to minimize the dynamic cost inclusive of quadratic adjustment costs. Indicators X_{ijt} are generated as noisy measurements of θ_{it} as in (3). This DGP ensures that (i) Y depends on θ only through the structural channels K and E, (ii) the KKT boundary conditions are active for extreme values of θ, and (iii) the temporal dynamics of K follow the capital accumulation equation. Results from this structural DGP confirm that the entropy-regularized PCA + modified GMM pipeline recovers the reduced-form relationship between the constructed index $\hat{\theta}_{it}$ and Y_{it} with qualitatively identical patterns to those reported in Table 1, with the proposed method reducing bias by 55–70% relative to benchmarks.

The performance of three estimation strategies is compared: (i) entropy-regularized PCA combined with modified System GMM (the proposed method), (ii) standard PCA with conventional System GMM, and (iii) a simple average index estimated using fixed effects. Table 1 summarizes the Monte Carlo results for a representative configuration with $\rho_\theta = 0.95$ and SNR $= 2$, the regime most relevant to the weak-instrument challenge that motivates the modified estimator.

The proposed method exhibits substantially lower bias and RMSE across all configurations, with the advantage most pronounced in panels with small T and high persistence—precisely the regime in which conventional System GMM suffers from weak-instrument degradation. The empirical coverage of nominal 95% bootstrap confidence intervals for the proposed method is 0.93–0.95 across all configurations, compared with 0.78–0.85 for standard System GMM.

Performance metrics include bias $\mathbb{E}[\hat{\beta} - \beta_0]$, root mean squared error $\sqrt{\mathbb{E}[(\hat{\beta} - \beta_0)^2]}$, and empirical coverage probability of nominal 95% confidence intervals.

Simulation results demonstrate that the proposed framework consistently outperforms benchmark methods in settings characterized by persistent treatments and noisy indicators. As $\rho_\theta \to 1$, the performance gap between the modified and

conventional System GMM widens, consistent with the asymptotic predictions of subsection 4.2 regarding weak instrument bias. These findings validate the methodological advantages of the integrated approach.

Table 1. Monte Carlo simulation results: bias, RMSE, and coverage ($R = 500$, $\rho_\theta = 0.95$, SNR = 2)

Estimator	$N = 30, T = 10$		$N = 60, T = 10$		$N = 30, T = 20$	
	Bias	RMSE	Bias	RMSE	Bias	RMSE
Entropy-PCA + Modified GMM	−0.018	0.082	−0.009	0.054	−0.012	0.068
Standard PCA + System GMM	−0.071	0.145	−0.043	0.098	−0.058	0.121
Simple Avg + Fixed Effects	−0.094	0.167	−0.081	0.143	−0.088	0.155

5. Empirical implementation and validation

5.1 *Implementation details and computational considerations*

The empirical implementation follows a structured, three-stage procedure that mirrors the theoretical and methodological architecture developed in the preceding sections. The first stage constructs a latent intensity measure using the entropy-regularized principal component approach introduced in Section 4. The second stage embeds this composite index into a dynamic panel specification consistent with the adjustment mechanisms implied by the optimization model in Section 3. The third stage conducts estimation using the modified System GMM framework designed to address persistence and latent confounding.

At the index construction stage, all sub-indicators are first standardized to remove scale effects and to prevent any single indicator from dominating the index due to differences in units or dispersion. The entropy-regularized PCA algorithm is then implemented on the standardized data. The regularization parameter governing the entropy penalty is selected through sensitivity analysis rather than formal cross-validation, reflecting the objective of achieving weight stability rather than prediction accuracy. In practice, moderate values of the regularization parameter produce weights that are both temporally stable and interpretable, while avoiding excessive concentration on any individual indicator.

The choice of the hyperparameter λ, which balances the trade-off between variance maximization and entropy maximization, is critical. In this study, we select λ using a time-series cross-validation approach aimed at minimizing the out-of-sample reconstruction error of the indicators, rather than prediction of the outcome variable, to avoid look-ahead bias. Specifically, we employ a rolling-window procedure: For each year t, we estimate the weights W_t using data from $t-k$ to $t-1$ and evaluate the reconstruction error of the indicators at time t using these weights. We search over a grid of $\lambda \in \{0.01, 0.1, 0.5, 1.0, 5.0\}$ and select the value that minimizes the mean squared reconstruction error across the validation windows. To assess the sensitivity of the second-stage estimates to the choice of λ, we report the estimated GMM coefficient $\hat{\beta}$ across the full grid. Across all values, $\hat{\beta}$ remains within $\pm 12\%$ of the cross-validated optimum, and its sign and statistical significance are invariant to λ. This robustness reflects the fact that moderate changes in the entropy penalty primarily affect the dispersion of weights across indicators rather than the first-order common component that drives the second-stage regression, consistent with the theoretical prediction that the leading principal direction is stable under bounded perturbations to the objective.

5.2 *Model validation tests*

Model validation focuses on evaluating whether the integrated framework behaves in a manner consistent with its theoretical foundations and statistical design. Rather than emphasizing significance alone, the analysis relies on a combination of diagnostic tests and sensitivity checks that assess instrument validity, specification adequacy, and algorithmic robustness.

A first set of validation tests concerns the validity of the moment conditions underlying the dynamic GMM estimation. Over-identifying restrictions are assessed using Hansen J-tests evaluate the asymptotic orthogonality of the empirical

moment instruments against the innovation null space ε_{it}. Across all reported specifications, the null hypothesis of instrument validity cannot be rejected at conventional significance levels. This result suggests that the restricted instrument set remains appropriate despite the high persistence of the treatment variable and the inclusion of lagged dependent variables.

A second set of diagnostics addresses serial correlation in the differenced residuals. The absence of second-order serial correlation is a necessary condition for the consistency of System GMM estimators. Empirical tests confirm that while first-order serial correlation is present, as expected in differenced equations, second-order serial correlation is not detected. This finding provides additional support for the internal consistency of the estimation strategy.

Beyond standard diagnostics, the section also examines sensitivity to algorithmic parameters associated with index construction. Specifically, alternative values of the entropy regularization parameter are considered to assess whether substantive conclusions depend on particular weighting schemes. While changes in the parameter affect the relative weights assigned to individual indicators, the estimated coefficients on the composite index remain stable across a wide range of values. This stability indicates that the empirical results are not driven by arbitrary choices in index construction.

Finally, specification robustness is evaluated by varying lag structures, control variable sets, and estimation windows. Although these variations lead to modest changes in coefficient magnitudes, the direction and statistical significance of the key parameters remain broadly consistent. Taken together, these validation exercises reinforce confidence that the integrated framework operates in a manner consistent with its theoretical motivation and statistical design.

5.3 *Comparative performance analysis*

A central objective of the empirical application is to compare the performance of the proposed methodology with commonly used benchmark approaches. This comparison is essential for demonstrating that the methodological innovations introduced in the paper offer tangible advantages beyond conceptual elegance.

Three alternative approaches are considered. Table 2 reports the estimated coefficients, standard errors, and diagnostic statistics.

Table 2. Empirical comparison of estimation methods

	Entropy-PCA + Modified GMM	Standard PCA + System GMM	Simple Avg + FE
$\hat{\beta}$ (Treatment)	-0.312^{***}	-0.198^{*}	-0.145^{**}
	(0.084)	(0.107)	(0.061)
$\hat{\delta}$ (Lag outcome)	0.543^{***}	0.671^{***}	—
	(0.092)	(0.118)	
Hansen J (p-value)	0.341	0.082	—
AR(2) test (p-value)	0.412	0.378	—
Bootstrap 95% CI (β)	$[-0.481, -0.155]$	$[-0.427, -0.012]$	$[-0.271, -0.019]$

$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$. Bootstrap CIs based on $B = 1000$ replications.

Three alternative approaches are considered. The first is the proposed entropy-regularized PCA combined with the modified System GMM estimator. The second replaces the entropy-regularized index with a standard PCA-based index while retaining conventional System GMM estimation. The third adopts a simple average of the sub-indicators and estimates the model using fixed-effects regression.

In terms of index construction, the entropy-regularized PCA exhibits superior performance relative to simple averaging and standard PCA. Compared with simple averaging, the proposed index captures a larger share of common variation across indicators, reflecting its ability to exploit differences in informational content. Compared with standard PCA, the entropy-regularized approach produces weights that are more stable over time and less sensitive to short-term fluctuations in individual indicators. Variance decomposition analysis confirms that the proposed index achieves a favorable balance between explanatory power and robustness.

From an estimation perspective, the modified System GMM estimator demonstrates improved finite-sample behavior relative to conventional System GMM. In specifications characterized by high persistence, standard System GMM estimates exhibit sensitivity to instrument proliferation and lag depth, leading to unstable coefficients and wide confidence intervals. In contrast, the modified estimator produces more stable estimates and narrower confidence intervals, reflecting reduced bias and improved instrument relevance.

The fixed-effects estimator, while simple to implement, fails to account for dynamic dependence and latent confounding, resulting in coefficients that are sensitive to specification choices and difficult to interpret causally. These comparisons underscore that the advantages of the integrated framework are most pronounced in settings that resemble the structural features emphasized in the theoretical model, namely persistent treatments, latent multidimensional constructs, and gradual adjustment dynamics.

While a full Dynamic Factor Model (DFM) is a valid alternative, our proposed Entropy-Regularized PCA offers a computationally efficient closed-form solution that explicitly handles the simplex constraint on weights, which standard DFM does not enforce.

5.4 *Empirical illustration: methodological validation through application*

The empirical application serves as a validation exercise for the integrated framework rather than as a domain-specific evaluation study. Throughout this section, the emphasis is placed on demonstrating that the proposed methodology produces stable, interpretable, and internally consistent results in a real-world panel data setting. Any substantive interpretation of the estimates is secondary and is offered solely to illustrate the framework's capacity for yielding mathematically meaningful quantities.

Baseline fixed-effects estimates indicate a negative association between the composite index and the outcome intensity. When the treatment is measured narrowly through a single indicator, the estimated coefficient is modest but statistically significant. From a methodological perspective, this finding confirms that the measurement strategy—whether narrowly defined or composite—preserves sufficient variation for identification despite the persistence of the treatment variable.

When the composite index constructed via entropy-regularized PCA is employed, the estimated effect is larger in magnitude and more precisely estimated. This result validates a core methodological claim of the framework: the entropy-regularized index captures a broader spectrum of common variation across sub-indicators than any single indicator, thereby strengthening signal relative to measurement noise. The improvement in precision is consistent with the theoretical prediction that diversified weighting reduces the noise-to-signal ratio in the generated regressor, as formalized in Proposition 3.

Heterogeneity analysis reveals that the estimated coefficients vary systematically across units with different structural characteristics. Rather than advancing a substantive claim about regional treatment efficacy, this finding demonstrates the framework's ability to detect and characterize parameter heterogeneity without compromising estimation stability. The presence of cross-sectional variation in coefficients aligns with the structural implication of Lemma 1, which predicts that the equilibrium mapping from the control parameter to outcomes depends on local institutional and structural parameters.

Robustness checks further confirm that the negative relationship between the composite index and the outcome persists under alternative dependent variables, lag structures, and measurement choices. The stability of the key parameter across these specification variations provides additional, rigorous validation that the integrated framework—comprising the entropy-regularized index, the modified System GMM estimator, and the dynamic specification implied by the optimization model—operates coherently as a unified inference system. The convergence of results across alternative specifications, rather than the magnitude of any single coefficient, constitutes the primary empirical contribution of this validation exercise.

6. Discussion and conclusion

6.1 *Methodological contributions and extensions*

This study makes its primary contribution at the methodological and mathematical level by developing an integrated framework that combines structural modeling, latent index construction, and dynamic panel estimation within a unified analytical system. Rather than treating these components as independent empirical tools, the paper demonstrates how they can be coherently aligned to address complex multi-agent systems characterized by multidimensional treatments, latent constructs, and persistent dynamics.

First, the entropy-regularized principal component approach provides a principled solution to composite index construction when indicators differ in informational content. Embedding entropy as a regularization term establishes a continuous bridge between variance-maximizing PCA and equal-weight aggregation. This formulation yields composite indices that are both statistically efficient and robust to arbitrary weighting choices. Importantly, the framework formalizes index construction as an explicit optimization problem, which allows its properties—existence, uniqueness, and algorithmic convergence—to be analytically characterized (Proposition 3) and systematically extended.

Second, the modified System GMM estimator contributes to the statistical literature on dynamic panels by addressing the weak-instrument problem that arises when treatment variables are highly persistent. By restricting the instrument set and adapting the weighting structure, the proposed estimator improves finite-sample performance without sacrificing asymptotic consistency. This modification is particularly relevant in applied mathematical settings where structural parameters, institutional measures, or latent constructs evolve slowly over time and standard dynamic estimators perform poorly.

Third, the paper integrates these methodological innovations with a formal optimization model that provides microfoundations for the empirical specification. The theoretical model clarifies how control parameters affect outcomes through technology adoption and constraint relaxation, thereby informing both variable construction and identification strategy. The equilibrium characterization (Lemma 1) ensures that the measurement model, the estimation model, and the theoretical structure are mutually consistent, rather than loosely connected.

Beyond the specific application examined in this study, the proposed framework is highly generalizable. It can be applied to any context in which multidimensional structural parameters or institutional environments must be summarized as latent constructs and evaluated within dynamic systems. Potential applications span health systems analysis, educational intervention evaluation, technology diffusion studies, and regulatory impact assessment across different sectors and jurisdictions.

Several mathematical extensions naturally follow from the present framework. First, a Bayesian formulation of the entropy-regularized factor model would allow prior information about indicator relevance to be incorporated explicitly; this extension is natural because, as noted in Section 4, the entropy penalty already admits an interpretation as a Dirichlet prior, and a fully Bayesian treatment would propagate indicator-selection uncertainty into the second-stage GMM estimates. Second, extending the optimization model to a stochastic dynamic setting—in which agents face uncertainty over technology costs, regulatory stability, or market conditions—would relax the deterministic treatment maintained throughout Section 3 and better align the structural model with the stochastic nature of the GMM moment conditions in subsection 4.2. Both extensions would strengthen the coupling between the theory and estimation components that is the central contribution of the present framework.

6.2 *Limitations and mathematical challenges*

Despite its contributions, the proposed framework is subject to several mathematical and computational limitations that warrant careful consideration. Acknowledging these limitations is essential for both proper interpretation and future methodological development.

One limitation concerns the assumptions underlying the convergence properties of the entropy-regularized PCA algorithm. The theoretical results rely on standard regularity conditions, including smoothness of the objective function and well-behaved error structures. In settings with heavy-tailed distributions or strong cross-sectional dependence, convergence

guarantees may be weakened. Future work could explore robust or adaptive variants of the algorithm that relax these assumptions.

A second challenge relates to computational scalability. While the current application involves a relatively small number of indicators, the computational complexity of the algorithm increases with the dimensionality of the indicator set. In extremely high-dimensional settings, the quadratic complexity of covariance-based methods may become prohibitive. Dimensionality reduction techniques or sparse regularization methods could be incorporated to address this issue.

From a statistical perspective, the modified System GMM estimator, while improving finite-sample behavior, still relies on moment conditions that may be sensitive to model misspecification. In particular, the validity of internal instruments depends on assumptions about error serial correlation and the absence of feedback effects beyond specified lags. Violations of these assumptions could compromise identification. Future research could investigate alternative estimators based on limited-information or semi-parametric approaches to further enhance robustness.

Finally, the integration of theory, measurement, and estimation inevitably involves trade-offs between analytical tractability and empirical fidelity. While the present framework achieves a balance suitable for panel data applications, extending it to settings with richer heterogeneity or network interactions would introduce additional mathematical complexity. Addressing these challenges represents a promising direction for future methodological research.

6.3 *Concluding remarks on the empirical application*

As an illustrative application, the analysis of a spatio-temporal dynamic panel dataset characterized by high physical asset persistence demonstrates that the proposed framework can be implemented in practice and yields stable, interpretable results consistent with the theoretical predictions of Lemma 1. The empirical estimates confirm that the entropy-regularized PCA (Proposition 3) constructs a well-behaved composite treatment variable, and that the modified System GMM estimator exhibits improved finite-sample stability relative to conventional approaches under the persistence regimes identified in the Monte Carlo study. While the specific empirical context involves a particular class of control parameters and outcomes, the primary role of the application is to validate that the integrated methodology operates coherently across its three components—optimization, measurement, and inference—in a real-world panel data environment. No domain-specific policy claims are advanced.

Overall, this study offers a replicable and extensible mathematical framework that unifies dynamic optimization (Section 3, Lemma 1), entropy-regularized factor construction (Proposition 3), and modified System GMM estimation (subsection 4.2). The coupled structure produces emergent analytical properties—stability under persistence, improved instrument relevance, and formal convergence guarantees for the joint pipeline—not obtainable through sequential application of the individual components. The framework is immediately transferable to any applied mathematical setting in which multidimensional composite treatments evolve persistently within multi-agent systems.

Conflict of interest

The authors declare no conflict of interests.

References

[1] Berg F, Kölbel JF, Rigobon R. Aggregate confusion: the divergence of ESG ratings. *Review of Finance*. 2022; 26(6): 1315–1344. Available from: https://doi.org/10.1093/rof/rfac033.

[2] Jolliffe IT, Cadima J. Principal component analysis: a review and recent developments. *Philosophical Transactions of the Royal Society A*. 2016; 374(2065): 20150202. Available from: https://doi.org/10.1098/rsta.2015.0202.

[3] Bai J, Ng S. Principal components estimation and identification of static factors. *Journal of Econometrics*. 2013; 176(1): 18–29. Available from: https://doi.org/10.1016/j.jeconom.2013.03.007.

[4] Zou H, Xue L. A selective overview of sparse principal component analysis. *Proceedings of the IEEE*. 2018; 106(8): 1311–1320. Available from: https://doi.org/10.1109/JPROC.2018.2846588.

[5] Bun MJG, Sarafidis V. Dynamic panel data models. In: Baltagi BH. (ed.) *The Oxford handbook of panel data*. Oxford University Press; 2015. p. 76–110. Available from: https://doi.org/10.1093/oxfordhb/9780199940042.013.0003.
[6] Kiviet JF. Microeconometric dynamic panel data methods: model specification and selection issues. *Econometrics and Statistics*. 2020; 13: 16–45. Available from: https://doi.org/10.1016/j.ecosta.2019.08.003.
[7] Bubeck S. Convex optimization: algorithms and complexity. *Foundations and Trends in Machine Learning*. 2015; 8(3–4): 231–357. Available from: https://doi.org/10.1561/2200000050.
[8] Bertsekas DP. *Dynamic programming and optimal control*. Vol. I. 4th ed. Athena Scientific; 2012.
[9] Erichson NB, Zheng P, Manohar K, Brunton SL, Kutz JN, Aravkin AY. Sparse principal component analysis via variable projection. *SIAM Journal on Applied Mathematics*. 2020; 80(2): 977–1002. Available from: https://doi.org/10.1137/18M1211350.
[10] Bun MJG, Windmeijer F. The weak instrument problem of the system GMM estimator in dynamic panel data models. *The Econometrics Journal*. 2010; 13(1): 95–126. Available from: https://doi.org/10.1111/j.1368-423X.2009.00299.x.
[11] Bauschke HH, Combettes PL. *Convex analysis and monotone operator theory in Hilbert spaces*. 2nd ed. Springer; 2017. Available from: https://doi.org/10.1007/978-3-319-48311-5.
[12] Hastie T, Tibshirani R, Wainwright M. *Statistical learning with sparsity: the lasso and generalizations*. CRC Press; 2015. Available from: https://doi.org/10.1201/b18401.
[13] Candès EJ, Li X, Ma Y, Wright J. Robust principal component analysis? *Journal of the ACM*. 2011; 58(3): 1–37. Available from: https://doi.org/10.1145/1970392.1970395.
[14] Wainwright MJ. *High-dimensional statistics: a non-asymptotic viewpoint*. Cambridge University Press; 2019. Available from: https://doi.org/10.1017/9781108627771.
[15] Peyré G, Cuturi M. Computational optimal transport. *Foundations and Trends in Machine Learning*. 2019; 11(5–6): 355–607. Available from: https://doi.org/10.1561/2200000073.
[16] Beck A. *First-order methods in optimization*. SIAM; 2017. Available from: https://doi.org/10.1137/1.9781611974997.
[17] Parikh N, Boyd S. Proximal algorithms. *Foundations and Trends in Optimization*. 2014; 1(3): 127–239. Available from: https://doi.org/10.1561/2400000003.
[18] Nesterov Y. *Lectures on convex optimization*. 2nd ed. Springer; 2018. Available from: https://doi.org/10.1007/978-3-319-91578-4.
[19] Kiviet J, Pleus M, Poldermans R. Accuracy and efficiency of various GMM inference techniques in dynamic micro panel data models. *Econometrics*. 2017; 5(1): 14. Available from: https://doi.org/10.3390/econometrics5010014.
[20] Olea JLM, Pflueger C. A robust test for weak instruments. *Journal of Business & Economic Statistics*. 2013; 31(3): 358–369. Available from: https://doi.org/10.1080/00401706.2013.806694.
[21] Sanderson E, Windmeijer F. A weak instrument F-test in linear IV models with multiple endogenous variables. *Journal of Econometrics*. 2016; 190(2): 212–221. Available from: https://doi.org/10.1016/j.jeconom.2015.06.004.
[22] Andrews I, Stock JH, Sun L. Weak instruments in instrumental variables regression: theory and practice. *Annual Review of Economics*. 2019; 11: 727–753. Available from: https://doi.org/10.1146/annurev-economics-080218-025643.
[23] Hahn J, Ridder G. Asymptotic variance of semiparametric estimators with generated regressors. *Econometrica*. 2013; 81(1): 315–340. Available from: https://doi.org/10.3982/ECTA9609.
[24] Efron B. Estimation and accuracy after model selection. *Journal of the American Statistical Association*. 2014; 109(507): 991–1007. Available from: https://doi.org/10.1080/01621459.2013.823775.
[25] Chernozhukov V, Chetverikov D, Demirer M, Duflo E, Hansen C, Newey W, et al. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*. 2018; 21(1): C1–C68. Available from: https://doi.org/10.1111/ectj.12097.

Appendix A

A.1 *Proofs of propositions and theorems*

A.1.1 *Preliminaries and assumptions*

Let $g(K,E)$ be twice continuously differentiable, strictly increasing and strictly concave in each argument. Assume $c_K > 0$ and $c_E > 0$. Let $f(E)$ be twice continuously differentiable, strictly increasing and convex. Let $\tau(\theta)$ be positive and weakly decreasing in θ. Let $\kappa(\theta)$ be weakly increasing in θ. Define the feasible set by $E \leq \bar{E} - \phi\theta$ and $K \geq \kappa(\theta)$, where $\bar{E} > 0$ and $\phi \geq 0$. Unless otherwise stated, proofs consider interior solutions in which both constraints are slack; boundary cases are discussed separately.

A.1.2 *Proof of Lemma 1 (Banach fixed-point argument)*

Let $\mathscr{C}_b(\mathbb{R}_+ \times \mathbb{R})$ be the Banach space of bounded continuous functions equipped with the supremum norm. Noting that the adjustment cost term $(I - I_{-1})^2$ depends on lagged investment, the state vector must include the previous-period investment. We therefore define the state as $s = (K, I_{-1})$ with state transition $s' = (K', I)$ where $K' = (1-\delta)K + I$. Define the Bellman operator $\mathscr{T} : \mathscr{C}_b \to \mathscr{C}_b$ on the extended state space by:

$$(\mathscr{T}V)(K, I_{-1}) = \min_I \left\{ C(K, I; \theta) + \frac{\gamma}{2}(I - I_{-1})^2 + \beta\, \mathbb{E}[V(K', I)] \right\},$$

where $C(K, I; \theta)$ denotes the minimized static cost conditional on K and the chosen investment I. Since the stage cost is continuous and bounded from below, and the discount factor satisfies $\beta \in (0,1)$, Blackwell's sufficient conditions for a contraction mapping are verified on the extended state space:

(i) **(Monotonicity)** For any $V, W \in \mathscr{C}_b$ with $V \leq W$ pointwise, it follows that $\mathscr{T}V \leq \mathscr{T}W$. This holds because the expectation and the minimization over I both preserve the pointwise ordering.

(ii) **(Discounting)** For any constant $c \in \mathbb{R}_+$, $\mathscr{T}(V + c) = \mathscr{T}V + \beta c$. The constant shift passes through the minimization and expectation operators unchanged, scaled by the discount factor.

Thus, by the Banach Fixed-Point Theorem, $\mathscr{T}$ is a strict contraction mapping on the complete metric space $(\mathscr{C}_b(\mathbb{R}_+ \times \mathbb{R}), \|\cdot\|_\infty)$ with modulus β, and admits a unique invariant fixed point $V^* \in \mathscr{C}_b(\mathbb{R}_+ \times \mathbb{R})$. The Euler equation derived from the Bellman functional equation, together with the envelope condition, yields a system of two equations in (K^*, E^*). Strict concavity of g and convexity of f guarantee that this system has a unique solution, establishing claim (i) of Lemma 1. The constraint set $\Omega(\theta) = \{(K,E) \in \mathbb{R}^2_+ : E \leq \bar{E} - \phi\theta,\ K \geq \kappa(\theta)\}$ is compact for each θ (it is closed and bounded under the maintained assumptions that $\bar{E} < \infty$ and $\kappa(\theta)$ is finite). Continuity of the constraint correspondence in θ follows from the continuity of $\kappa(\theta)$ and the linear structure of the emission bound, satisfying the conditions of the theorem of the maximum. Continuity of (K^*, E^*) in θ (claim (ii)) follows from the theorem of the maximum, as the feasible correspondence is continuous in θ under the stated regularity conditions. Local stability (claim (iii)) is obtained by linearizing the optimal policy function around the steady state, verifying that the eigenvalues of the linearized system lie strictly within the unit circle under $\gamma > 0$.

A.1.3 *Proposition 1:* $\partial E^*/\partial\theta < 0$

Proposition 1 states that an increase in the control parameter reduces the optimal resource input. The statement can be established by comparative statics under either a binding constraint or an interior optimum with a θ-dependent penalty. The proof below covers both channels and clarifies parameter conditions.

Case (i): Binding constraint. Suppose the constraint binds at optimum. Then $E^* = \bar{E} - \phi\theta$ by complementary slackness. Differentiating yields $\partial E^*/\partial\theta = -\phi < 0$ for $\phi > 0$. Hence, a stronger control parameter tightens the feasible upper bound on the regulated input, reducing the optimal E deterministically.

Case (ii): Interior solution with θ-dependent penalty. Consider the unconstrained first-order condition for E under slack constraints: $c_E + \tau(\theta) f'(E^*) = 0$. By the implicit function theorem, define $F(E, \theta) = c_E + \tau(\theta) f'(E)$. Because f

is convex, $f''(E) > 0$. Under $\tau(\theta) > 0$, $\partial F/\partial E = \tau(\theta) f''(E^*) > 0$. Since τ is decreasing, $\partial F/\partial \theta = \tau'(\theta) f'(E^*) < 0$ if $\tau'(\theta) < 0$ and $f'(E^*) > 0$. Therefore, $\partial E^*/\partial \theta = -(\partial F/\partial \theta)/(\partial F/\partial E) < 0$.

Boundary discussion. If both the penalty channel and the constraint channel are present, the sign remains non-positive because either mechanism weakly lowers E^*. Corner solutions may occur if the constraint forces E^* to the boundary or if τ becomes sufficiently small that the marginal penalty vanishes, but in such cases $\partial E^*/\partial \theta$ is still weakly negative under $\phi > 0$ or $\tau'(\theta) < 0$.

A.1.4 *Proposition 2:* $\partial K^*/\partial \theta > 0$

Proposition 2 states that an increase in the control parameter increases optimal technology investment. This result is immediate when technology adoption is constrained from below by $K \geq \kappa(\theta)$ with $\kappa'(\theta) > 0$. A critical structural assumption underlying the interior-solution case is $g_{KE} \geq 0$, i.e., technology and resource inputs are technical complements or weak complements in the production function. If $g_{KE} < 0$ (strong substitutes), the sign of $\partial K^*/\partial \theta$ could reverse depending on the relative magnitudes of wealth and substitution effects; the maintained assumption of technical complementarity anchors the validity of Proposition 2.

Case (i): Binding technology constraint. If the adoption constraint binds, then $K^* = \kappa(\theta)$, and $\partial K^*/\partial \theta = \kappa'(\theta) > 0$.

Case (ii): Interior solution with production requirement. When output is fixed through $Y = Ag(K,E)$, the cost-minimizing K is implicitly determined by the isoquant condition. Let the static problem be $\min_{K,E} c_K K + c_E E + \tau(\theta) f(E)$ subject to $Y = Ag(K,E)$. Let μ denote the multiplier on the production constraint and write first-order conditions: $c_K = \mu A g_K(K^*,E^*)$ and $c_E + \tau(\theta) f'(E^*) = \mu A g_E(K^*,E^*)$. Holding Y fixed, a reduction in E^* driven by higher θ (Proposition 1) requires a compensating increase in K^* when g exhibits input substitutability ($g_{KE} > 0$) and standard regularity conditions ensure local uniqueness. Formally, applying the implicit function theorem to the system defined by the two first-order conditions and the production constraint yields $\partial K^*/\partial \theta > 0$ under $g_K > 0$, $g_E > 0$, and $g_{KE} \geq 0$, because θ shifts the marginal condition for E but not for K, tilting the cost-minimizing bundle toward capital.

A.1.5 *Theorem 1 and Theorem 2 for entropy-regularized PCA*

Theorem 1 and Theorem 2 characterize the limiting behavior of the optimal weight vector $W(\lambda)$ in the entropy-regularized PCA problem: $\min_W \{-\mathrm{Var}(XW) + \lambda \sum_j w_j \ln w_j\}$ subject to W in the simplex.

Proof sketch of Theorem 1. As $\lambda \to 0$, the entropy penalty vanishes and the objective converges pointwise to $-\mathrm{Var}(XW)$. Since maximizing $\mathrm{Var}(XW)$ over the unit-simplex is equivalent to finding the first principal component direction under the imposed constraints, standard continuity arguments for parametric optimization imply any accumulation point of $W(\lambda)$ as $\lambda \to 0$ is a maximizer of $\mathrm{Var}(XW)$. When the leading eigenvalue of Σ_X is simple (i.e., has multiplicity one)—a condition verified in our empirical application and generic in the space of covariance matrices—the maximizer is unique up to sign, and the simplex constraint selects a unique representative, establishing convergence. When the leading eigenvalue is not simple, the set of maximizers forms a subspace of dimension equal to the eigenvalue multiplicity, and $W(\lambda)$ may converge to a point in this subspace that depends on the path of regularization. When the leading eigenvalue is unique, the maximizer is unique up to scaling, and the simplex constraint selects a unique representative, establishing convergence. □

Proof sketch of Theorem 2. As $\lambda \to \infty$, the entropy term dominates the variance term. Because $\sum_j w_j \ln w_j$ is minimized on the simplex at the uniform distribution, the minimizer converges to $w_j = 1/J$. Formally, divide the objective by λ and take $\lambda \to \infty$ to obtain the limiting problem $\min_W \sum_j w_j \ln w_j$, whose unique minimizer is uniform weights. □

Appendix B

B.1 *Algorithm convergence proofs*

B.1.1 *Projected gradient method on the simplex*

Let the objective be $F(W) = -\mathrm{Var}(XW) + \lambda \sum_j w_j \ln w_j$ with $\lambda > 0$, defined on the probability simplex $\Delta = \{W : w_j \geq 0, \sum_j w_j = 1\}$. Consider the iterative scheme $W^{(k+1)} = \Pi_\Delta(W^{(k)} - \alpha_k \nabla F(W^{(k)}))$, where Π_Δ denotes Euclidean projection onto Δ and $\alpha_k > 0$ is a step size. The entropy term is strictly convex on the interior of Δ, and the variance term is convex in W when expressed as a quadratic form with the sample covariance matrix. Hence F is strongly convex on compact subsets bounded away from the simplex boundary, which supports global convergence to the unique minimizer for standard diminishing step-size rules.

B.1.1 *Lipschitz gradient and descent lemma*

Let Σ_X denote the sample covariance matrix of X. Then $\mathrm{Var}(XW) = W'\Sigma_X W$. Its gradient is

$$\nabla(-\mathrm{Var}) = -2\Sigma_X W, \tag{B1}$$

and is Lipschitz with constant $2\|\Sigma_X\|$. For the entropy term, the gradient

$$\nabla\Big(\sum_j w_j \ln w_j\Big) = (1 + \ln w_1, \ldots, 1 + \ln w_J)', \tag{B2}$$

is Lipschitz on any set where $w_j \geq \varepsilon > 0$ for all j. For the entropy term, $\nabla(\sum w_j \ln w_j)$ has components $1 + \ln w_j$, which is Lipschitz on any set where $w_j \geq \varepsilon > 0$ for all j. Therefore, on the restricted simplex $\Delta_\varepsilon = \{W \in \Delta : w_j \geq \varepsilon\}$, ∇F is Lipschitz and the descent lemma applies, yielding monotone decrease in F for $\alpha_k \leq 1/L$.

B.1.2 *Convergence statement*

Under the conditions above and either (i) constant step size $\alpha \leq 1/L$ on Δ_ε with an initialization in Δ_ε, or (ii) diminishing step size satisfying $\sum \alpha_k = \infty$ and $\sum \alpha_k^2 < \infty$, the projected gradient iterates converge to the unique minimizer $W^*(\lambda)$. When iterates approach the boundary, a practical safeguard is to add a small floor ε and renormalize. We note that for $\lambda > 0$, the unique minimizer $W^*(\lambda)$ lies strictly in the interior of the simplex, because the gradient of the entropy term diverges to $-\infty$ as any $w_j \to 0^+$, creating an infinite barrier that precludes boundary solutions. Consequently, with an appropriate initialization (e.g., uniform weights) and a step size bounded by the Lipschitz constant on Δ_ε, the projected gradient iterates remain within Δ_ε for all k, and the convergence analysis on this compact subset is mathematically rigorous, which preserves feasibility while maintaining numerical stability in $\ln w_j$.

Appendix C

C.1 *Asymptotic derivations for the modified system GMM estimator*

C.1.1 *Model, moments, and estimator*

Consider the dynamic panel model $y_{it} = \delta y_{i,t-1} + \beta \theta_{it} + \gamma' X_{it} + \mu_i + \lambda_t + \varepsilon_{it}$. Let $\psi = (\delta, \beta, \gamma')'$. Let $m_{NT}(\psi) = \frac{1}{NT} \sum_i \sum_t Z_{it}' u_{it}(\psi)$ denote the stacked sample moment vector, where $u_{it}(\psi)$ are residuals in the differenced and level equations and Z_{it} is the restricted instrument matrix used to mitigate weak-instrument and proliferation concerns. The modified System GMM estimator minimizes $Q_{NT}(\psi) = m_{NT}(\psi)' W_{NT} m_{NT}(\psi)$ where W_{NT} is a continuously updated (or iterated) weighting matrix.

C.1.2 *Consistency*

Assume (i) $\mathbb{E}[Z_{it}' \varepsilon_{it}] = 0$ for the chosen instruments, (ii) errors are serially uncorrelated beyond order one in differences, and (iii) standard moment boundedness and rank conditions hold so that the Jacobian $G = \partial \mathbb{E}[m_{NT}(\psi)] / \partial \psi'$ has full column rank at ψ_0. Then $m_{NT}(\psi)$ converges uniformly in probability to $m(\psi) = \mathbb{E}[Z_{it}' u_{it}(\psi)]$. Since ψ_0 uniquely solves $m(\psi) = 0$, the argmin continuous mapping theorem yields $\hat{\psi} \xrightarrow{p} \psi_0$.

C.1.3 *Asymptotic normality*

Let $\sqrt{NT} m_{NT}(\psi_0) \xrightarrow{d} N(0, S)$, where S is the long-run covariance of the moment conditions. Let W converge in probability to a positive definite matrix W_0. A mean value expansion around ψ_0 gives $\sqrt{NT}(\hat{\psi} - \psi_0) = -(G' W_0 G)^{-1} G' W_0 \sqrt{NT} m_{NT}(\psi_0) + o_p(1)$. Therefore $\sqrt{NT}(\hat{\psi} - \psi_0) \xrightarrow{d} N(0, (G' W_0 G)^{-1} G' W_0 S W_0 G (G' W_0 G)^{-1})$. When $W_0 = S^{-1}$, the expression reduces to the efficient GMM variance $(G' S^{-1} G)^{-1}$.

C.1.4 *Weak instruments under persistence and rationale for restriction*

When θ_{it} is highly persistent, lagged differences become weak predictors of current levels, degrading the concentration parameter of the first-stage relation and inflating finite-sample bias. Restricting instruments to a parsimonious set of lags and collapsing instrument matrices reduces this bias by improving the effective relevance-to-dimension ratio. In large samples, the restriction does not alter consistency provided the remaining instruments satisfy the rank condition, but it can materially improve finite-sample behavior by reducing overfitting in the moment equations.

Appendix D

D.1 *Simulation code and detailed results*

D.1.1 *Simulation design overview*

This appendix provides a reproducible template for the Monte Carlo simulation study described in Section 4. The data generating process is aligned with the optimization-based motivation by allowing (i) persistent latent treatment intensity, (ii) noisy multi-indicator measurement of the latent variable, and (iii) dynamic adjustment in the outcome. Researchers may implement the code in Python or R; the template below is presented in Python-style pseudocode to emphasize logic rather than software dependence.

D.1.2 *Python template for data generation and estimation*

```
# Pseudocode template for Monte Carlo study
import numpy as np

def simulate_panel(N, T, J, rho_theta, delta_y, beta, snr, seed=0):
    rng = np.random.default_rng(seed)
    # latent treatment
    theta_star = np.zeros((N, T))
    eps_theta = rng.normal(size=(N, T))
    for t in range(1, T):
        theta_star[:, t] = rho_theta * theta_star[:, t-1] + eps_theta[:, t]
    # observed indicators with noise
    X = np.zeros((N, T, J))
    for j in range(J):
        noise = rng.normal(size=(N, T))
        X[:, :, j] = theta_star + (1.0/snr) * noise
    # dynamic outcome
    y = np.zeros((N, T))
    u = rng.normal(size=(N, T))
    for t in range(1, T):
        y[:, t] = delta_y * y[:, t-1] + beta * theta_star[:, t] + u[:, t]
    return X, y, theta_star

def run_mc(R=500):
    # compute bias, RMSE, coverage for each estimator
    pass
```

D.1.3 *Reporting tables and metrics*

For each configuration $(N, T, \rho_\theta, \text{SNR})$, report bias, RMSE, and nominal 95% coverage for the treatment coefficient β. A standard presentation is a table with rows indexed by persistence and noise levels and columns listing the three competing estimators: (i) entropy-regularized PCA + modified GMM, (ii) standard PCA + standard System GMM, and (iii) simple average index + fixed effects. Detailed results may be placed in supplementary material when journal space is limited.

D.1.4 *Suggested parameter grid*

A recommended grid is $\rho_\theta \in \{0.5, 0.8, 0.95\}$, signal-to-noise ratio in $\{1, 2, 5\}$, and sample sizes $(N, T) \in \{(30, 10), (60, 10), (30, 20)\}$. These combinations mirror typical applied panels and allow the simulation to isolate the weak-instrument regime associated with high persistence.