A Generalized Finite Dependence Framework for Dynamic Discrete Choice Models*

Yu Hao† Hiroyuki Kasahara‡ Katsumi Shimotsu§

May 4, 2026

Abstract

We characterize when finite dependence exists in dynamic discrete choice models and develop an efficient estimator for the finite-dependence class. A flow parameterization converts the weight-search problem into a convex quadratic program, and a controllability-matrix rank test—the econometric analogue of a Kalman reachability condition—gives a checkable existence diagnostic. Within the finite-dependence class, the optimally weighted estimator attains the minimum trace of the asymptotic variance (FD-class trace optimality). A conditional extension shows that, under full controllability, an instrument-rank condition, and two additional per-DGP gates, this optimum coincides with the ρ-horizon semiparametric efficiency bound; we present this as a gated extension rather than an unconditional theorem. When finite dependence holds, payoff-only counterfactual CCPs satisfy a fixed-point system that admits Bellman-free iteration with Bellman-based verification for uniqueness. Monte Carlo exercises deliver consistency on a single-agent capital-investment model and a three-player Markov-perfect entry/exit game, with substantial speed gains over Bellman-based benchmarks: $\sim$370–1,000$\times$ over CCP-2step and NFPM on the single-agent calibration, $\sim 4\times$ over NFXP on the dynamic game, and up to $\sim 136\times$ over NFXP on a state-space-scaling sweep at $|\mathcal{X}| = 5{,}000$ when action-invariant Kronecker structure is exploited.

Keywords: Dynamic Discrete Choice Models; Identification; Counterfactual Analysis; Finite Dependence; Conditional Semiparametric Efficiency; Conditional Choice Probability Estimator.

JEL Codes: C14, C25, C57, C61, C73.

*We are grateful to Vadim Marmer, Kevin Song, and Florian Hoffmann for valuable comments.
†Faculty of Business and Economics, The University of Hong Kong. Email: haoyu@hku.hk.
‡Vancouver School of Economics, The University of British Columbia. Email: hkasahar@mail.ubc.ca.
§Faculty of Economics, The University of Tokyo. Email: shimotsu@e.u-tokyo.ac.jp.

1 Introduction

Dynamic discrete choice (DDC) models are widely used to study forward-looking behavior, but their estimation requires computing the agent's expected future value function, the fixed point of a Bellman equation. Nested within an estimation routine, repeatedly solving for this fixed point is computationally prohibitive when the state space is large, and the resulting estimators can be sensitive to functional-form assumptions used to extrapolate value functions into sparsely observed regions.

Finite dependence, developed in the lineage of Altug and Miller (1998) and Arcidiacono and Miller (2011) (with the underlying ideas appearing in earlier 2008 and 2010 working-paper versions of the latter), offers an alternative. If future state distributions induced by distinct current actions can be matched after a short horizon ρ, the continuation values cancel from the estimating equations, and the value-function difference reduces to a finite sum of flow utilities and conditional choice probabilities (CCPs; Hotz and Miller, 1993). This idea underlies much of the recent applied DDC literature. Yet the scope of finite dependence has been limited by reliance on model-specific constructions—renewal actions, terminal paths—and no general conditions for its existence have been available. The question of *when* finite dependence holds for a given transition structure, and how to compute the required weights when it does, has remained open.

By *generalized* finite dependence, we mean the extension from the non-negative weighting schemes of Arcidiacono and Miller (2011) (which rely on problem-specific constructions such as renewal actions or terminal paths, at a small dependence horizon) to signed weights over arbitrary horizons $\rho \geq 1$, parameterized via a flow representation that converts the weight-finding problem into a convex quadratic program.

This paper makes five contributions. Ordered from unconditional to conditional, the first four are the existence/rank-test characterization, the payoff-only computation bridge, FD-class trace optimality, and the conditional extension to a stronger economic target, ρ-horizon semiparametric efficiency. The fifth contribution collects the structural and numerical infrastructure that makes those results operational.

First, we characterize existence. A flow parameterization converts the search for valid weighting schemes into a convex quadratic program, and Theorem 3 shows that finite dependence holds if and only if a baseline distributional gap lies in the column space of a pair-specific controllability matrix $\widetilde{\mathcal{C}}_\rho$, the econometric analogue of a Kalman rank condition. This delivers a pre-estimation diagnostic: once the transition matrices are estimated, the researcher can test whether finite dependence is available at a candidate horizon ρ before committing to the estimator.

Second, we establish the payoff-only computation bridge. Theorem 5 shows that when finite dependence holds, counterfactual CCPs under payoff-only policy experiments satisfy a CCP-space fixed-point system that admits Bellman-free iteration from the finite-dependence representation. Uniqueness and exclusion of extraneous fixed points are not delivered by the iteration itself: the Bellman contraction and the Hotz–Miller bijection remain the verification tools, implemented in practice through restart and spectral-radius diagnostics. We frame this result as *Bellman-free iteration with Bellman-based verification*—a computational route within the payoff-only class—rather than as an unconditional identification theorem. KSSR identification (Kalouptsidi et al., 2021) remains the primitive compatibility condition for counterfactual analysis in DDC models; this paper adds an FD-based computational route and a validation protocol within the payoff-only class.

Third, we establish FD-class trace optimality. The finite-dependence linear system generically admits a family of feasible weights parameterized by null-space coordinates, and Theorem 8 shows that the weight vector minimizing $\operatorname{tr}(\boldsymbol{\Sigma}(\boldsymbol{\omega}))$ over a DGP-specific compact regular sub-region of that family attains the trace-minimum asymptotic variance within the finite-dependence class at horizon ρ. This is the paper's strongest closed efficiency result.

Fourth, we isolate the conditional extension to a stronger economic target. Equality to the ρ-horizon semiparametric benchmark requires additional gates beyond FD-class trace optimality: full controllability, an instrument-Jacobian rank condition, FD-Coincidence, and oracle equivalence. Section 4.5 therefore presents the semiparametric-efficiency language as a gated extension, not a front-end unconditional theorem, clearly separated from the closed FD-class result.

Fifth, we provide the supporting structural and numerical infrastructure. Proposition 10 shows that shift-register models with p-lag memory satisfy finite dependence automatically at $\rho = p$, while the Kronecker-separable structure discussed in Section 5.1 decomposes the feasibility check component by component. For the shared-action dynamic game used in our Monte Carlo exercise, Lemma 11—the action-invariant Kronecker cancellation that applies when opponents' actions appear symmetrically in the transition—delivers the concrete scaling benefit, reducing a five-player dynamic game from 38,880 joint states to 30 per player. Theorem 7 gives the two-step asymptotic theory together with the variance–weight tradeoff, and Section 3.8 extends the framework to unobserved heterogeneity through a GFD-EM scheme that nests the Arcidiacono and Miller (2011) CCP-EM iteration with a substantial per-iteration computational saving. Appendix A collects the numerical infrastructure used in the empirical sections: reachability pruning (A.1), iterative LSQR for deep-tree flow QPs (A.2), KKT-LU batched solving for wide-state $\rho = 1$ flow QPs (A.4), and the structured Kronecker FD solver that exploits action-invariant exogenous factors and underlies the $|\mathcal{X}|$-scaling results of Table 3 (A.5).

Related literature. This paper builds directly on Arcidiacono and Miller (2011) and Arcidiacono and Miller (2019), which establish finite-dependence-based CCP estimation without solving the full dynamic program. Prior to this paper, verifying finite dependence in general required problem-specific arguments: demonstrating a renewal action (Arcidiacono and Miller, 2011) or a terminal path structure (Arcidiacono and Miller, 2019) for each model specification. As a result, several influential empirical studies resort to full-solution methods precisely because their transition structures lack these special features: occupational-choice models with multi-dimensional human capital (Keane and Wolpin, 1997), industrial-dynamics models with action-dependent state persistence (Ryan, 2012), and sectoral mobility models where human capital depreciates at action-specific rates (Dix-Carneiro, 2014). Our contribution is to replace this case-by-case verification with a general algebraic framework that accommodates these settings. Arcidiacono and Miller (2020) show that counterfactual CCPs are generally not identified off short panels unless finite dependence holds, reinforcing the importance of our characterization for identification beyond estimation convenience.

Our use of "generalized" refers to the flow parameterization over arbitrary dependence horizons $\rho \geq 1$ with signed weights; this is distinct from Gayle (2021) who extends finite dependence to continuous-choice settings. An earlier working paper by two of the present authors (Hao and Kasahara, 2024) treats the two-period case using sequential weight optimization with Kronecker structure; it does not supply an arbitrary-ρ existence characterization, general asymptotic theory for signed GFD weights, or the efficiency architecture developed here.

The closest methodological comparison is Kalouptsidi et al. (2021) (KSSR), who provide a linear-representation approach to counterfactual analysis in DDC models. KSSR's central result establishes a null-space compatibility condition between a block-Kronecker matrix built from observed transitions and the Jacobian of the counterfactual payoff transformation. Our framework is complementary rather than substitutive: Theorem 5 gives a checkable computational route within the payoff-only class, with an explicit fixed-point algorithm and restart-based validation protocol once finite dependence is verified. Magnac and Thesmar (2002) establish that DDC models are generically underidentified absent normalization restrictions and study which restrictions (exclusion, parameter-sharing) restore identification at the level of the model + data distribution; our framework contributes a different operational route via finite dependence rather than "providing the constructive conditions" for the Magnac–Thesmar identification results, which rest on independent assumptions. Berry and Compiani (2023) develop IV approaches for dynamic models with serially correlated unobservables; our framework applies to the standard conditional-independence setting but extends its computational reach.

Recent work addresses DDC tractability through complementary margins. Bunting and Ura (2025) exploit index-invertibility of reduced-form CCPs whose indices combine flow-utility and transition components, and Blevins (2025) exploit continuous-time structure. Our framework contributes the rank-test characterization (Theorem 3) that applies to arbitrary transitions, the flow parameterization that converts weight-finding into a convex QP, and the efficiency result (Theorem 8). Our approach uses transition structure in discrete time. In dynamic games, Aguirregabiria and Marcoux (2021) emphasize equilibrium restrictions; our extension exploits player-specific transition decomposition to reduce computational burden under separability assumptions. Adusumilli and Eckardt (2025) use temporal-difference learning to approximate value functions in DDC models with continuous or high-dimensional state spaces, avoiding explicit computation of transition matrices; their approach works for any DDC model but introduces approximation error, while ours achieves exact elimination of continuation values when finite dependence holds. Chen (2025) develops model-adaptive sieve approximations to the Bellman equation that improve convergence rates in large state spaces. Both approaches approximate value functions more efficiently; finite dependence eliminates them entirely.

The use of negative weights connects to the broader econometrics literature on signed measures. Borusyak and Hull (2024) argue that negative weights are not a concern in design-based specifications (which include shift-share designs), because the target estimand is a convex combination of treatment effects with positive ex-ante weights even when realized weights can be negative. de Chaisemartin and D'Haultfœuille (2020) demonstrate that in two-way fixed-effects estimators, negative weights on heterogeneous treatment effects can be problematic because the estimand becomes a non-convex combination. Within synthetic control methods, Abadie et al. (2019) likewise do not require donor weights to be non-negative when computing the synthetic match, and they provide a recursive procedure for the weights; the non-linearity of that recursion makes the computation expensive. The GFD parallel is to keep the signed-weight feasible set but to replace the non-linear recursion with a single *linear* feasibility test (Theorem 3): the constraint matrix is shared across initial states, so the entire weight-search problem collapses to one factorisation applied to all right-hand sides at once (Section 6.1). Our setting differs in kind from the design-based and TWFE literatures: negative weights are applied to structural transition primitives—probability distributions over future states—not to heterogeneous treatment effects. The finite dependence condition ensures exact cancellation of the continuation value, so the resulting estimand is the true payoff difference regardless of weight signs. Under approximate finite dependence, signed weights amplify the approximation bias through the weight-norm factor, creating a variance concern analogous to weak instruments but not a sign-reversal concern.

The asymptotic theory of the GFD estimator builds on standard two-step M-estimation (Murphy and Topel, 1985; Newey and McFadden, 1994); the variance-input mapping connects to Newey (1990)'s efficient instrumental-variables framework, while the conditional ρ-horizon efficiency benchmark uses the Chamberlain (1987) minimum-chi-squared bound. The counterfactual computation route of Theorem 5 intersects with the dynamic-game NPL framework of Aguirregabiria and Mira (2007) and the two-step game estimators of Bajari et al. (2007). The unobserved-heterogeneity extension of Section 3.8 interfaces with the nonparametric mixture identification of Kasahara and Shimotsu (2009). The control-theoretic provenance of the controllability recursion follows the discrete-time free-switching reachability literature (Liberzon, 2003; Sun and Ge, 2005).

The remainder of the paper proceeds as follows. Section 2 introduces the model and the generalized finite-dependence framework. Section 3 develops the CCP estimator, the payoff-only counterfactual computation bridge (Section 3.7), and the unobserved-heterogeneity extension (Section 3.8). Section 4 establishes asymptotic normality, the variance–input mapping, optimal weight selection (Section 4.4), and the conditional semiparametric-efficiency extension (Section 4.5). Section 5 presents canonical applications and the rank-test verification on ten DDC models. Section 6 reports Monte Carlo evidence on a single-agent investment model and a three-player Markov-perfect entry/exit game. Section 7 concludes. Appendix A collects the scalable computation machinery.

2 Model and Generalized Finite Dependence

2.1 Dynamic Discrete Choice Framework

We consider a standard infinite-horizon dynamic discrete choice model. In each period t, an agent observes the state $x_t \in \mathcal{X}$ with $|\mathcal{X}| = S$ and chooses an action $d_t \in \mathcal{D} = \{0, 1, \ldots, D-1\}$ with $|\mathcal{D}| = D$. We use the slice notation $x_{a:b} := (x_a, x_{a+1}, \ldots, x_b)$ and $d_{a:b} := (d_a, d_{a+1}, \ldots, d_b)$ for any state or action sub-sequence; it is the only sequence notation used in the paper.

Assumption 1 (Additive separability)**.** *The instantaneous payoff associated with action d_t in state x_t is additively separable:*

$$U(x_t, d_t, \epsilon_t) \;=\; u(x_t, d_t) + \epsilon_t(d_t), \tag{1}$$

where $u : \mathcal{X} \times \mathcal{D} \to \mathbb{R}$ is the systematic flow payoff and $\epsilon_t = (\epsilon_t(0), \ldots, \epsilon_t(D-1))$ is a vector of idiosyncratic preference shocks.

Assumption 2 (Conditional independence)**.** *The shock vector ϵ_t is drawn i.i.d. across time from a joint distribution $G(\epsilon)$ with continuous support and finite first moments. The shocks are observed by the agent before making the decision at t but are unknown at $t-1$.*

Assumption 3 (Markov transitions)**.** *The state x_t evolves according to a controlled first-order Markov process with transition probability $f(x_{t+1} \mid x_t, d_t)$, independent of ϵ_t.*

Assumption 4 (Discounting)**.** *The agent discounts future payoffs at rate $\beta \in (0,1)$.*

Under these assumptions, the choice-specific value function is

$$v(x, d) \;=\; u(x, d) \;+\; \beta \sum_{x' \in \mathcal{X}} V(x')\, f(x' \mid x, d), \tag{2}$$

where the ex-ante value function $V(x) = \int \max_d \{v(x,d) + \epsilon(d)\}\, dG(\epsilon)$. We write $\mathbf{p}(x) = (p(0 \mid x), \ldots, p(A-1 \mid x))$ for the equilibrium conditional choice probabilities.

2.2 The Arcidiacono–Miller Representation

Finite dependence methods identify structural parameters by expressing the value-function difference $v(x,d) - v(x,d')$ as a finite sum of flow payoffs. Throughout the paper d denotes the *analyzed action* and d' the *reference (baseline) action* against which d is compared; the two are arbitrary distinct elements of $\mathcal{D}$ and the QP formulation in Section 2.3 is symmetric in (d, d'), so the choice of which to label as "reference" is conventional. Identifying $v(x, d) - v(x, d')$ requires finding two future continuation schemes starting from d and d' that produce the same state distribution after ρ periods, canceling the continuation value $V_{t+\rho+1}$.

Lemma 1 (Arcidiacono–Miller representation)**.** *For any state x and any signed weights $\boldsymbol{\omega}(x) = (\omega(0 \mid x), \ldots, \omega(A-1 \mid x))$ satisfying $\sum_{d \in \mathcal{D}} \omega(d \mid x) = 1$, the value function admits*

$$V(x) \;=\; \sum_{d \in \mathcal{D}} \omega(d \mid x) \big[\, v(x, d) + \psi_d(\mathbf{p}(x)) \,\big], \tag{3}$$

where $\psi_d(\mathbf{p}(x)) := V(x) - v(x, d)$. By the Hotz–Miller inversion, ψ_d is a known function of the CCP vector $\mathbf{p}(x)$ and the shock distribution $G(\epsilon)$.[1]

Proof. Substituting $\psi_d(\mathbf{p}(x)) = V(x) - v(x, d)$ into (3) gives $\sum_d \omega(d \mid x)\big[v(x, d) + V(x) - v(x, d)\big] = V(x) \sum_d \omega(d \mid x) = V(x)$, since the weights sum to unity. □

The representation (3) holds for *any* weights summing to unity, including signed weights—the structural opening that generalized finite dependence exploits. We define the *augmented flow payoff*

$$\hat{u}(x, d) \;:=\; u(x, d) + \psi_d(\mathbf{p}(x)). \tag{4}$$

2.3 Flow Parameterization and the Convex QP

Why a flow parameterization? For dependence horizons $\rho \geq 2$, a per-step weighting scheme $\{\omega_\tau\}_{\tau=1}^{\rho}$ enters the value-function representation as a product $\prod_\tau f \cdot \omega_\tau$, and a direct search over the ω_τ's is nonlinear in the decision variables. We resolve this by introducing a single joint flow object that absorbs both transition probabilities and per-step weights into one variable. The resulting feasibility conditions become linear, the matching objective is quadratic, and the problem reduces to a convex quadratic program.

The flow matrix. Starting from (x_0, d_0), define the *flow matrix* $\Phi(x_0, d_0) \in \mathbb{R}^{S^\rho \times D^\rho}$, indexed by future state sequences $x_{1:\rho} \in \mathcal{X}^\rho$ and action sequences $d_{1:\rho} \in \mathcal{D}^\rho$. We write $\phi_{x_{1:\rho}, d_{1:\rho}}(x_0, d_0)$ for its $(x_{1:\rho}, d_{1:\rho})$ entry, the joint flow through the path $(x_{1:\rho}, d_{1:\rho})$ starting from (x_0, d_0). Each ϕ-entry is the (signed) product of the probability of visiting state sequence $x_{1:\rho}$ and the per-step action weights along $d_{1:\rho}$; entries are not restricted to be non-negative.

[1] For instance, when ϵ follows a Type-I extreme-value distribution, $\psi_d(\mathbf{p}(x)) = \gamma_E - \ln p(d \mid x)$, where $\gamma_E \approx 0.5772$ is the Euler–Mascheroni constant.

Linear constraints on Φ. Two types of linear conditions ensure that $\Phi(x_0, d_0)$ represents a valid policy.

(i) Initial flow. For each $x_1 \in \mathcal{X}$, the total flow through histories that begin at state x_1 equals the one-step arrival probability:

$$\sum_{x_{2:\rho}} \sum_{d_{1:\rho}} \phi_{(x_1, x_{2:\rho}), d_{1:\rho}}(x_0, d_0) \;=\; f(x_1 \mid x_0, d_0). \tag{5}$$

(ii) Flow conservation. For each $\tau \in \{1, \dots, \rho - 1\}$, each prefix $(x_{1:\tau}, d_{1:\tau})$, and each next state $x_{\tau+1}$,

$$\sum_{x_{\tau+2:\rho}} \sum_{d_{\tau+1:\rho}} \phi_{x_{1:\rho}, d_{1:\rho}}(x_0, d_0) \;=\; f(x_{\tau+1} \mid x_\tau, d_\tau) \cdot \sum_{x_{\tau+1:\rho}} \sum_{d_{\tau+1:\rho}} \phi_{x_{1:\rho}, d_{1:\rho}}(x_0, d_0), \tag{6}$$

where the left-hand sum holds the prefix $(x_{1:\tau+1}, d_{1:\tau})$ fixed and the right-hand sum holds $(x_{1:\tau}, d_{1:\tau})$ fixed. Summing (5) over x_1 shows that the entries of $\Phi(x_0, d_0)$ sum to one, so $\Phi(x_0, d_0)$ is a signed probability measure on path space.

Per-step decomposition. Constraints (5)–(6) are equivalent to the existence of per-step weights $\omega_\tau(d_\tau \mid x_{0:\tau}, d_{0:\tau-1})$ ($\tau = 1, \dots, \rho$) satisfying $\sum_{d \in \mathcal{D}} \omega_\tau(d \mid x_{0:\tau}, d_{0:\tau-1}) = 1$ at every history with non-zero marginal flow, such that

$$\phi_{x_{1:\rho}, d_{1:\rho}}(x_0, d_0) \;=\; \prod_{\tau=1}^{\rho} f(x_\tau \mid x_{\tau-1}, d_{\tau-1}) \cdot \omega_\tau(d_\tau \mid x_{0:\tau}, d_{0:\tau-1}). \tag{7}$$

The information set conditioning ω_τ contains states $x_0, x_1, \dots, x_\tau$ and actions $d_0, d_1, \dots, d_{\tau-1}$—exactly the history available to a forward-looking agent at the moment d_τ is chosen. The weights need not be non-negative, and they may depend on the full prefix; when they happen to depend only on the contemporaneous state x_τ, (7) reduces to the Markov form imposed by Arcidiacono and Miller (2019). The crucial advantage of working with Φ directly is that the constraints (5)–(6) are linear in Φ, converting the nonlinear search over $\{\omega_\tau\}$ into a convex quadratic program.

Remark 1 (Signed weights and the affine hull)**.** *We do not impose $\Phi(\cdot) \geq 0$. Each deterministic action sequence $d_{1:\rho}$ induces a terminal state distribution $\boldsymbol{\pi}(d_{1:\rho})$. With non-negative weights the set of achievable terminal distributions is the convex hull of $\{\boldsymbol{\pi}(d_{1:\rho})\}_{d_{1:\rho}}$; with signed weights it is the affine hull—the full affine subspace spanned by these distributions. In many settings, affine intersections exist when convex intersections do not. Signed weights expand the feasible set at the cost of potential variance inflation through the effective regressor in the GFD pseudo-likelihood (Section 3.5).*

Terminal distribution and the QP. The terminal-state distribution induced by $\Phi(x_0, d_0)$ is

$$\kappa_{\rho+1}(x' \mid x_0, d_0) \;=\; \sum_{x_{1:\rho}} \sum_{d_{1:\rho}} \phi_{x_{1:\rho},\, d_{1:\rho}}(x_0, d_0) \cdot f(x' \mid x_\rho, d_\rho). \tag{8}$$

The path probability is already absorbed into Φ; no separate factor appears. With analyzed action d and reference action d' from state x_0, the full optimization problem is

$$\begin{aligned} \min_{\Phi(x_0,d),\, \Phi(x_0,d')} \quad & \sum_{x' \in \mathcal{X}} \big(\kappa_{\rho+1}(x' \mid x_0, d) - \kappa_{\rho+1}(x' \mid x_0, d') \big)^2 \\ \text{s.t.} \quad & \text{constraints (5) and (6) for both } \Phi(x_0, d) \text{ and } \Phi(x_0, d'). \end{aligned} \tag{9}$$

The two flow matrices share the same flow-conservation structure but have different initial-flow constraints: $\Phi(x_0, d)$ is constrained by $f(\cdot \mid x_0, d)$ on the right-hand side of (5), while $\Phi(x_0, d')$ is constrained by $f(\cdot \mid x_0, d')$. All constraints are linear and the objective is quadratic in $(\Phi(x_0, d), \Phi(x_0, d'))$, making this a convex QP.

2.4 A Worked Example: $S = 2$, $D = 2$, $\rho = 1$

Time line. Take $\mathcal{X} = \{L, H\}$, $\mathcal{D} = \{0, 1\}$, and $\rho = 1$. At time $t = 0$ the agent observes $x_0 = L$ and chooses an initial action $d_0 \in \mathcal{D}$. At time $t = 1$ the state $x_1 \in \mathcal{X}$ is drawn from $f(\cdot \mid L, d_0)$ and the agent chooses $d_1 \in \mathcal{D}$ according to the per-step weight $\omega_1(\cdot \mid x_0, x_1, d_0)$. The terminal state $x_2 \in \mathcal{X}$ is drawn from $f(\cdot \mid x_1, d_1)$; $\rho = 1$ means we evaluate the matching condition on the distribution of x_2. The QP compares two scenarios that differ *only* in the time-0 action: in scenario A the agent picks $d_0 = d$, in scenario B the agent picks $d_0 = d'$. After time 0 the two scenarios share the same linear constraints, but produce distinct flow matrices because $f(\cdot \mid L, d)$ and $f(\cdot \mid L, d')$ differ.

Flow matrix and initial-flow constraint. With $\rho = 1$ the only path index is (x_1, d_1), and the flow matrix is

$$\Phi(L, d_0) \;=\; \begin{pmatrix} \phi_{L,0}(L, d_0) & \phi_{L,1}(L, d_0) \\ \phi_{H,0}(L, d_0) & \phi_{H,1}(L, d_0) \end{pmatrix}, \qquad d_0 \in \{d, d'\}.$$

Since $\rho = 1$ there is no flow conservation, and the initial-flow constraint (5) reduces to

$$\begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} \phi_{L,0}(L, d_0) \\ \phi_{L,1}(L, d_0) \\ \phi_{H,0}(L, d_0) \\ \phi_{H,1}(L, d_0) \end{pmatrix} \;=\; \begin{pmatrix} f(L \mid L, d_0) \\ f(H \mid L, d_0) \end{pmatrix}.$$

The right-hand side is $\big(f(L \mid L, d),\ f(H \mid L, d)\big)^\top$ when $d_0 = d$, and $\big(f(L \mid L, d'),\ f(H \mid L, d')\big)^\top$ when $d_0 = d'$; the constraint matrix is identical in the two scenarios.

Decomposition into the per-step weight at $\tau = 1$. The decomposition (7) collapses to

$$\phi_{x_1,\,d_1}(L, d_0) \;=\; f(x_1 \mid L, d_0) \cdot \omega_1(d_1 \mid x_0, x_1, d_0), \qquad x_1 \in \{L, H\},\ d_1 \in \{0, 1\}, \tag{10}$$

where $x_0 = L$ and $d_0 \in \{d, d'\}$ is the scenario-specific time-0 action. The row sums of $\Phi(L, d_0)$ recover the one-step arrival probabilities (the initial-flow constraint), and the within-row ratios $\phi_{x_1,d_1}(L, d_0)/\big(\phi_{x_1,0}(L, d_0) + \phi_{x_1,1}(L, d_0)\big) = \omega_1(d_1 \mid x_0, x_1, d_0)$ recover the time-1 per-step weight.

Terminal distribution and the QP. The terminal distribution at $t = 2$ is

$$\begin{pmatrix} \kappa_2(L \mid L, d_0) \\ \kappa_2(H \mid L, d_0) \end{pmatrix} = \begin{pmatrix} f(L \mid L, 0) & f(L \mid L, 1) & f(L \mid H, 0) & f(L \mid H, 1) \\ f(H \mid L, 0) & f(H \mid L, 1) & f(H \mid H, 0) & f(H \mid H, 1) \end{pmatrix} \begin{pmatrix} \phi_{L,0}(L, d_0) \\ \phi_{L,1}(L, d_0) \\ \phi_{H,0}(L, d_0) \\ \phi_{H,1}(L, d_0) \end{pmatrix},$$

and the QP (9) matches $\kappa_2(\cdot \mid L, d)$ against $\kappa_2(\cdot \mid L, d')$. Equation (10) is the simplest instance of the flow decomposition that drives the entire framework: each ϕ-entry is the per-state arrival probability scaled by the time-1 per-step weight, and matching terminal distributions across the two time-0 actions reduces to a finite-dimensional linear-quadratic problem.

2.5 Existence Characterization: Two Theorems

This subsection collects two characterization results that drive the rest of the paper. Theorem 2 shows that the linear constraints (5)–(6) are exactly the constraints needed for Φ to admit the per-step multiplicative form (7); this is the structural justification for the QP reformulation. Theorem 3 then characterizes *pointwise* finite dependence at a single triple (x_0, d, d') as solvability of an explicit linear system $\mathbf{A}\,\mathbf{x} = \mathbf{b}$.

Equivalence between linear constraints and the per-step product form.

Theorem 2 (Constraint–decomposition equivalence)**.** *Fix $(x_0, d_0) \in \mathcal{X} \times \mathcal{D}$ and a horizon $\rho \geq 1$. A flow matrix $\Phi(x_0, d_0) \in \mathbb{R}^{S^\rho \times D^\rho}$ satisfies the initial-flow constraint* (5) *and the flow-conservation (NAC) constraint* (6) *if and only if there exist per-step weights*

$$\omega_\tau\big(\, d_\tau \,\big|\, x_{0:\tau},\, d_{0:\tau-1} \big), \qquad \tau = 1, \ldots, \rho,$$

satisfying $\sum_{d \in \mathcal{D}} \omega_\tau(d \mid x_{0:\tau}, d_{0:\tau-1}) = 1$ at every history with non-zero marginal flow, such that *the multiplicative decomposition* (7) *holds entry-wise:*

$$\phi_{x_{1:\rho},\, d_{1:\rho}}(x_0, d_0) \;=\; \prod_{\tau=1}^{\rho} f(x_\tau \mid x_{\tau-1}, d_{\tau-1}) \cdot \omega_\tau(d_\tau \mid x_{0:\tau}, d_{0:\tau-1}).$$

On every history with non-zero marginal flow, the implied weights ω_τ *are unique.*

Proof sketch. *(⇐).* Substitute the product form into the LHS of (5). Summing over $(x_{2:\rho}, d_{1:\rho})$, the factors $f(x_\tau \mid x_{\tau-1}, d_{\tau-1})$ for $\tau \geq 2$ and the weights ω_τ for $\tau \geq 1$ each sum to one in turn (transitions sum to one over the next state, weights sum to one over the next action), and the only surviving factor is $f(x_1 \mid x_0, d_0)$, matching the RHS. The analogous telescoping verifies (6).

(⇒). For each $\tau = 1, \ldots, \rho$ and each prefix $(x_{1:\tau}, d_{1:\tau-1})$, define the marginal flows

$$M^{\mathrm{s}}_\tau(x_{1:\tau}, d_{1:\tau-1}) \;:=\; \sum_{x_{\tau+1:\rho}} \sum_{d_{\tau:\rho}} \phi_{x_{1:\rho},\, d_{1:\rho}}(x_0, d_0), \quad M^{\mathrm{sa}}_\tau(x_{1:\tau}, d_{1:\tau}) \;:=\; \sum_{x_{\tau+1:\rho}} \sum_{d_{\tau+1:\rho}} \phi_{x_{1:\rho},\, d_{1:\rho}}(x_0, d_0).$$

On histories with $M^{\mathrm{s}}_\tau \neq 0$, set

$$\omega_\tau(d_\tau \mid x_{0:\tau}, d_{0:\tau-1}) \;:=\; M^{\mathrm{sa}}_\tau(x_{1:\tau}, d_{1:\tau}) \,/\, M^{\mathrm{s}}_\tau(x_{1:\tau}, d_{1:\tau-1}).$$

Summing the numerator over d_τ recovers the denominator, so $\sum_d \omega_\tau(d \mid \cdot) = 1$. Flow conservation (6) written as $M^{\mathrm{s}}_{\tau+1}(x_{1:\tau+1}, d_{1:\tau}) = f(x_{\tau+1} \mid x_\tau, d_\tau) \cdot M^{\mathrm{sa}}_\tau(x_{1:\tau}, d_{1:\tau})$, combined with the initial-flow base case $M^{\mathrm{s}}_1(x_1) = f(x_1 \mid x_0, d_0)$, gives the chain identity $\phi = \prod_\tau f \cdot \omega_\tau$ by telescoping. Uniqueness on nonzero-marginal histories is immediate from the definition of ω_τ as a ratio of marginals. □

Remark 2 (Non-anticipativity)**.** *The information set conditioning* ω_τ *in Theorem 2 is exactly* $(x_{0:\tau}, d_{0:\tau-1})$*: states up to time* τ *and actions up to time* $\tau - 1$*. The per-step weight* ω_τ *is therefore measurable with respect to the agent's time-*τ *information and* cannot *look ahead to future state realizations* $x_{\tau+1:\rho}$ *or future actions* $d_{\tau:\rho}$*. The flow-conservation constraint* (6) *is the linear-algebraic encoding of this non-anticipativity property.*

Pointwise finite dependence and linear feasibility.

Definition 1 (Pointwise ρ-period finite dependence)**.** Fix a triple $(x_0, d, d') \in \mathcal{X} \times \mathcal{D}^2$ with $d \neq d'$ (with d the analyzed action and d' the reference action), and a horizon $\rho \geq 1$. We say *pointwise* ρ*-period finite dependence* holds at (x_0, d, d') if there exist flow matrices $\Phi(x_0, d)$ and $\Phi(x_0, d')$ that satisfy initial flow (5) and flow conservation (6) and induce identical terminal distributions:

$$\kappa_{\rho+1}(x' \mid x_0, d) \;=\; \kappa_{\rho+1}(x' \mid x_0, d') \qquad \forall\, x' \in \mathcal{X}. \tag{11}$$

We now construct an explicit linear system whose solvability is equivalent to Definition 1. Let $N := S^\rho$, $M := D^\rho$, and let $\operatorname{vec}\Phi(x_0, d_0) \in \mathbb{R}^{NM}$ denote the column-stack of entries of

$\Phi(x_0, d_0)$, indexed by $(x_{1:\rho}, d_{1:\rho})$.

Define three sparse matrices, each acting on a single $\operatorname{vec}\Phi \in \mathbb{R}^{NM}$.

Initial-flow block $\mathbf{A}_{\text{init}} \in \{0,1\}^{S\times NM}$, with rows indexed by $x_1 \in \mathcal{X}$:

$$(\mathbf{A}_{\text{init}})_{x_1,\,(x_{1:\rho},\,d_{1:\rho})} \;=\; \mathbf{1}\{(\text{first state of } x_{1:\rho}) = x_1\}. \tag{12}$$

Flow-conservation (NAC) block $\mathbf{A}_{\text{nac}}$. Bring (6) to one side: for each $\tau \in \{1, \ldots, \rho - 1\}$, each prefix $(x_{1:\tau}, d_{1:\tau}) \in \mathcal{X}^\tau \times \mathcal{D}^\tau$, and each next state $x_{\tau+1} \in \mathcal{X}$,

$$\underbrace{\sum_{x_{\tau+2:\rho}} \sum_{d_{\tau+1:\rho}} \phi_{x_{1:\rho},\,d_{1:\rho}}(x_0, d_0)}_{\text{flow through } (x_{1:\tau+1},\, d_{1:\tau})} \;-\; f(x_{\tau+1} \mid x_\tau, d_\tau) \underbrace{\sum_{x_{\tau+1:\rho}} \sum_{d_{\tau+1:\rho}} \phi_{x_{1:\rho},\,d_{1:\rho}}(x_0, d_0)}_{\text{flow through } (x_{1:\tau},\, d_{1:\tau})} \;=\; 0. \tag{13}$$

Each instance of (13) is one linear equation in $\operatorname{vec}\Phi$, with coefficients $+1$ on every entry $(x_{1:\rho}, d_{1:\rho})$ extending the prefix $(x_{1:\tau+1}, d_{1:\tau})$ and $-f(x_{\tau+1} \mid x_\tau, d_\tau)$ on every entry extending $(x_{1:\tau}, d_{1:\tau})$; we take $\mathbf{A}_{\text{nac}}$ to be the matrix whose rows are these coefficient vectors, indexed by the quadruple $(\tau, x_{1:\tau+1}, d_{1:\tau})$. At horizon τ there are $S^{\tau+1}D^\tau$ such equations—one per choice of the S^τ state prefixes, D^τ action prefixes, and S next states—so $\mathbf{A}_{\text{nac}} \in \mathbb{R}^{C\times NM}$ with

$$C \;:=\; \text{number of NAC rows} \;=\; \sum_{\tau=1}^{\rho-1} S^{\tau+1} D^\tau. \tag{14}$$

By construction $\mathbf{A}_{\text{nac}} \operatorname{vec}\Phi = \mathbf{0}$ is exactly (6) written across all $(\tau, x_{1:\tau+1}, d_{1:\tau})$.

Terminal-projection block $\mathbf{T}_{\text{last}} \in \mathbb{R}^{S\times NM}$, with rows indexed by $x' \in \mathcal{X}$:

$$(\mathbf{T}_{\text{last}})_{x',\,(x_{1:\rho},\,d_{1:\rho})} \;=\; f(x' \mid x_\rho, d_\rho). \tag{15}$$

By (8), $\mathbf{T}_{\text{last}} \operatorname{vec}\Phi(x_0, d_0) = \kappa_{\rho+1}(\cdot \mid x_0, d_0)$.

Stack the two scenarios into one decision variable

$$\mathbf{x} \;:=\; \begin{pmatrix} \operatorname{vec}\Phi(x_0, d) \\ \operatorname{vec}\Phi(x_0, d') \end{pmatrix} \;\in\; \mathbb{R}^{2NM}.$$

With $\mathbf{f}^{(d_0)} := \big(f(x_1 \mid x_0, d_0)\big)_{x_1 \in \mathcal{X}} \in \mathbb{R}^S$, define

$$\mathbf{A} \;:=\; \begin{pmatrix} \mathbf{A}_{\text{init}} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\text{init}} \\ \mathbf{A}_{\text{nac}} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\text{nac}} \\ \mathbf{T}_{\text{last}} & -\mathbf{T}_{\text{last}} \end{pmatrix}, \qquad \mathbf{b} \;:=\; \begin{pmatrix} \mathbf{f}^{(d)} \\ \mathbf{f}^{(d')} \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix}. \tag{16}$$

The matrix $\mathbf{A} \in \mathbb{R}^{(3S+2C)\times 2NM}$ depends on the transition f, the horizon ρ, and the initial state x_0 (via $\mathbf{A}_{\text{init}}$ rows for x_1); the vector $\mathbf{b}$ additionally depends on the compared actions (d, d').

Theorem 3 (Pointwise FD as linear feasibility)**.** *Fix $(x_0, d, d') \in \mathcal{X} \times \mathcal{D}^2$ with $d \neq d'$, and let $\mathbf{A}, \mathbf{b}$ be the matrix and vector constructed in* (16)*. Then pointwise ρ-period finite dependence at (x_0, d, d') (Definition 1) holds if and only if*

$$\mathbf{b} \in \operatorname{Range}(\mathbf{A}), \qquad \textit{equivalently} \qquad \operatorname{rank}\big([\,\mathbf{A} \mid \mathbf{b}\,]\big) = \operatorname{rank}(\mathbf{A}). \tag{17}$$

Whenever (17) *holds, any solution $\mathbf{x} \in \mathbb{R}^{2NM}$ of $\mathbf{A}\,\mathbf{x} = \mathbf{b}$ yields, by inverse vectorization, flow matrices $\Phi(x_0, d)$ and $\Phi(x_0, d')$ that achieve terminal-distribution matching with horizon ρ.*

Proof. We evaluate $\mathbf{A}\,\mathbf{x}$ block by block and show that the equation $\mathbf{A}\,\mathbf{x} = \mathbf{b}$ is equivalent to the conjunction of the three conditions in Definition 1. Below, $d_0 \in \{d, d'\}$ ranges over the two scenarios; the upper sub-block of each pair corresponds to $d_0 = d$ and the lower to $d_0 = d'$.

Initial-flow blocks. For each $d_0 \in \{d, d'\}$ and each $x_1 \in \mathcal{X}$,

$$\big(\mathbf{A}_{\text{init}} \operatorname{vec} \Phi(x_0, d_0)\big)_{x_1} = \sum_{x_{2:\rho}} \sum_{d_{1:\rho}} \phi_{(x_1, x_{2:\rho}), d_{1:\rho}}(x_0, d_0),$$

which equals $\mathbf{f}_{x_1}^{(d_0)} = f(x_1 \mid x_0, d_0)$ exactly when $\Phi(x_0, d_0)$ satisfies the initial-flow constraint (5).

Flow-conservation (NAC) blocks. The row of $\mathbf{A}_{\text{nac}}$ indexed by $(\tau, x_{1:\tau}, d_{1:\tau}, x_{\tau+1})$, applied to $\operatorname{vec} \Phi(x_0, d_0)$, evaluates to the LHS minus the RHS of (6) at that prefix. The block evaluates to $\mathbf{0}$ for $d_0 = d$ and for $d_0 = d'$ exactly when both $\Phi(x_0, d)$ and $\Phi(x_0, d')$ satisfy flow conservation.

Terminal-matching block. By (8), $\mathbf{T}_{\text{last}} \operatorname{vec} \Phi(x_0, d_0) = \kappa_{\rho+1}(\cdot \mid x_0, d_0)$, so the last block evaluates to $\kappa_{\rho+1}(\cdot \mid x_0, d) - \kappa_{\rho+1}(\cdot \mid x_0, d')$, which equals $\mathbf{0}$ exactly when (11) holds.

Hence $\mathbf{A}\,\mathbf{x} = \mathbf{b}$ is equivalent to Conditions 1–3 of Definition 1, and the two-way characterization (17) follows. The constructive statement (second sentence of the theorem) is immediate from the equivalence. □

Remark 3 (Computational interpretation)**.** *Theorem 3 reduces verification of pointwise finite dependence at (x_0, d, d') to a single sparse linear-algebra computation. The matrix $\mathbf{A}$ has $(3S + 2C) \times 2NM$ entries with row count $3S + 2\sum_{\tau=1}^{\rho-1} S^{\tau+1} D^{\tau}$; both $\mathbf{A}$ and $\mathbf{b}$ are constructed from f and (x_0, d, d') alone, with no appeal to the value function or CCPs. In practice we test $\mathbf{b} \in \operatorname{Range}(\mathbf{A})$ via sparse least squares (LSQR), declaring feasibility when the residual $\|\mathbf{A}\widehat{\mathbf{x}} - \mathbf{b}\|_2$ falls below a fixed tolerance.*

Universal ρ-period finite dependence. We say ρ-period finite dependence holds *universally* if Definition 1 is satisfied at every triple $(x_0, d, d') \in \mathcal{X} \times \mathcal{D}^2$ with $d \neq d'$. By Theorem 3, universal FD is equivalent to pointwise feasibility of $\mathbf{A}^{(x_0,d,d')}\,\mathbf{x} = \mathbf{b}^{(x_0,d,d')}$ at every triple, which can be verified by running the LSQR test of Remark 3 once per triple. The estimator developed in Section 3 requires only this universal-pointwise version; a more compact closed-form characterization on the transition primitives (a controllability-matrix rank test) is possible but technically more involved, and we do not pursue it here.

3 The Generalized Finite Dependence (GFD) Estimator

This section develops the *Generalized Finite Dependence (GFD) estimator*, a two-step CCP estimator of the structural payoff parameter θ that takes a flow-matrix collection $\widehat{\boldsymbol{\Phi}}$ as an explicit input. Treating $\widehat{\boldsymbol{\Phi}}$ as an input is the key conceptual move: it separates *what is being estimated* (θ) from *how the continuation value is canceled* (a researcher-supplied flow-matrix collection that satisfies the FD constraints of Section 2). Different valid choices of $\widehat{\boldsymbol{\Phi}}$ produce different members of the GFD-estimator family, all consistent for θ but generically with different finite-sample variance.

3.1 Setup, Data, and First-Stage Inputs

We observe a panel $\{(x_{it}, d_{it})\}_{i=1,\,t=1}^{N,\,T}$ of N independent agents over T periods, generated by the DDC model of Section 2 at an unknown structural parameter $\theta_0 \in \Theta \subseteq \mathbb{R}^{d_\theta}$. Throughout this section we maintain Assumptions 1–4 together with the following two:

Assumption 5 (Linear-in-parameter payoff)**.** *The systematic flow payoff is linear in θ:*

$$u(x, d; \theta) \;=\; z(x, d)^\top \theta, \qquad d \in \mathcal{D}, \tag{18}$$

where $z : \mathcal{X} \times \mathcal{D} \to \mathbb{R}^{d_\theta}$ is a known regressor, normalized so that the reference action $d = 0$ has zero payoff: $z(x, 0) = \mathbf{0}$ for all x.

Assumption 6 (Type-I extreme-value shocks)**.** *The shocks $\epsilon_t(d)$ are i.i.d. Type-I extreme-value across (t, d), so that $\psi_d(\mathbf{p}(x)) = \gamma_E - \ln p(d \mid x)$ and the equilibrium CCP is the multinomial logit $p(d \mid x) = \exp(v(x, d))/\sum_{d''} \exp(v(x, d''))$.*

First-stage estimates. From the panel, the researcher constructs $\sqrt{N}$-consistent nonparametric estimates $\widehat{f}(x' \mid x, d)$ and $\widehat{p}(d \mid x)$ (cell-frequency, sieve, or kernel). The plug-in Hotz–Miller correction is $\psi_d(\widehat{\mathbf{p}}(x)) = \gamma_E - \ln \widehat{p}(d \mid x)$.

3.2 The Flow-Matrix Input $\widehat{\boldsymbol{\Phi}}$

The defining input of the GFD estimator is a *flow-matrix collection*

$$\widehat{\boldsymbol{\Phi}} \;=\; \left\{ \widehat{\Phi}(x, d_0) \,:\, x \in \mathcal{X},\; d_0 \in \mathcal{D} \right\} \tag{19}$$

indexed by every (x, d_0) pair, where each $\widehat{\Phi}(x, d_0) \in \mathbb{R}^{S^\rho \times D^\rho}$ is a flow matrix in the sense of Section 2.3. The researcher supplies the entire collection as input.

Definition 2 (FD-feasible flow input)**.** The collection $\widehat{\boldsymbol{\Phi}}$ is *FD-feasible at horizon ρ under transitions $\widehat{f}$* if, for every $x \in \mathcal{X}$,

(a) each $\widehat{\Phi}(x, d_0)$ satisfies the initial-flow constraint (5) and the flow-conservation (NAC) constraint (6), both with the transition $\widehat{f}$ in place of f; and

(b) the induced terminal distribution $\widehat{\kappa}_{\rho+1}(\cdot \mid x, d_0)$ from (8) is the *same for every $d_0 \in \mathcal{D}$*:

$$\widehat{\kappa}_{\rho+1}(x' \mid x, d) \;=\; \widehat{\kappa}_{\rho+1}(x' \mid x, d_0), \qquad \forall\, d, d_0 \in \mathcal{D},\; \forall\, x' \in \mathcal{X}. \tag{20}$$

Condition (a) is the per-action structural feasibility from Theorem 2; condition (b) is the FD matching condition extended to all D actions simultaneously, and is what makes the continuation value cancel from every value-difference $v(x, d; \theta) - v(x, 0; \theta)$ for $d \in \mathcal{D}$.

How to obtain $\widehat{\boldsymbol{\Phi}}$. Common choices include:

- *Moore–Penrose minimum-ℓ_2-norm solution.* For each x, solve the joint feasibility QP that enforces (a) and (b) and pick the minimum-ℓ_2-norm solution. This is a single sparse linear solve per x; we call this the *canonical input.*
- *Variance-optimal weights.* Choose $\widehat{\boldsymbol{\Phi}}$ to minimize a plug-in estimate of the asymptotic variance of $\widehat{\theta}^{\mathrm{GFD}}$; Section 4 treats this in detail.
- *Structure-exploiting closed form.* For shift-register transitions, $\widehat{\boldsymbol{\Phi}}$ admits an explicit construction (Proposition 10) that bypasses the QP entirely.

The estimator below is defined for any FD-feasible $\widehat{\boldsymbol{\Phi}}$; consistency holds for any feasible choice, and the variance depends on the specific input.

3.3 The GFD Value-Difference Identity

The GFD estimator builds on a closed-form representation of the value-function difference $v(x, d; \theta) - v(x, 0; \theta)$ in terms of the flow-matrix input.

Theorem 4 (GFD value-difference identity)**.** *Suppose Assumptions 1–6 hold and that $\boldsymbol{\Phi} = \{\Phi(x, d_0)\}_{x, d_0}$ is FD-feasible at horizon ρ under the true transitions f (Definition 2 with f in place of $\widehat{f}$). Then for every $(x, d) \in \mathcal{X} \times \mathcal{D}$ with $d \neq 0$,*

$$v(x, d; \theta) \,-\, v(x, 0; \theta) \;=\; z(x, d)^\top \theta \,+\, \sum_{\tau=1}^{\rho} \beta^\tau \sum_{x_{1:\rho}, d_{1:\rho}} \left[\phi_{x_{1:\rho}, d_{1:\rho}}(x, d) - \phi_{x_{1:\rho}, d_{1:\rho}}(x, 0)\right] \hat{u}(x_\tau, d_\tau; \theta, \mathbf{p}), \tag{21}$$

where $\hat{u}(x, d; \theta, \mathbf{p}) = z(x, d)^\top \theta + \psi_d(\mathbf{p}(x))$ *is the augmented flow payoff* (4). *The right-hand side of* (21) *depends only on the flow input* $\mathbf{\Phi}$*, the parameter* θ*, and the population CCPs* $\mathbf{p}$*; the continuation value* $V_{t+\rho+1}$ *does not appear.*

Proof sketch. Iterate the AM representation (3) of Lemma 1 ρ times along the path measure encoded by $\mathbf{\Phi}$. After ρ substitutions, the choice-specific value function $v(x, d_0; \theta)$ becomes

$$v(x, d_0; \theta) \;=\; z(x, d_0)^\top \theta + \sum_{\tau=1}^{\rho} \beta^\tau \sum_{x_{1:\rho}, d_{1:\rho}} \phi_{x_{1:\rho}, d_{1:\rho}}(x, d_0) \cdot \hat{u}(x_\tau, d_\tau; \theta, \mathbf{p}) + \beta^{\rho+1} \sum_{x'} \kappa_{\rho+1}(x' \mid x, d_0)\, V(x'),$$

with $\hat{u}(x_\tau, d_\tau; \theta, \mathbf{p}) = z(x_\tau, d_\tau)^\top \theta + \psi_{d_\tau}(\mathbf{p}(x_\tau))$. Taking the difference of this identity at $d_0 = d$ and $d_0 = 0$ and invoking the common-terminal condition (20) (which makes the V-term cancel) yields (21). □

3.4 Definition of the GFD Estimator

Define the *GFD pseudo value-difference* at parameter θ and inputs $(\widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}})$ as the sample analogue of the right-hand side of (21):

$$\widetilde{v}^{\rho\text{-FD}}(x, d;\, \theta, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) \;:=\; z(x, d)^\top \theta \;+\; \sum_{\tau=1}^{\rho} \beta^\tau \sum_{x_{1:\rho}, d_{1:\rho}} \left[\widehat{\phi}_{x_{1:\rho}, d_{1:\rho}}(x, d) - \widehat{\phi}_{x_{1:\rho}, d_{1:\rho}}(x, 0)\right] \hat{u}(x_\tau, d_\tau; \theta, \widehat{\mathbf{p}}), \tag{22}$$

with the normalization $\widetilde{v}^{\rho\text{-FD}}(x, 0; \cdot) = 0$ for every x. By Assumption 6, the FD-implied conditional choice probability is the multinomial logit

$$\Lambda_d(x;\, \theta, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) \;:=\; \frac{\exp\big(\widetilde{v}^{\rho\text{-FD}}(x, d;\, \theta, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}})\big)}{\sum_{d'' \in \mathcal{D}} \exp\big(\widetilde{v}^{\rho\text{-FD}}(x, d'';\, \theta, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}})\big)}. \tag{23}$$

Definition 3 (GFD estimator at input $\widehat{\mathbf{\Phi}}$)**.** Given first-stage estimates $(\widehat{f}, \widehat{p})$ and an FD-feasible flow input $\widehat{\mathbf{\Phi}}$ (under $\widehat{f}$), the *Generalized Finite Dependence estimator at input* $\widehat{\mathbf{\Phi}}$ is

$$\widehat{\theta}^{\text{GFD}}\big(\widehat{\mathbf{\Phi}}\big) \;:=\; \arg\max_{\theta \in \Theta} \; \frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} \log \Lambda_{d_{it}}\big(x_{it};\, \theta, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}\big). \tag{24}$$

What $\widehat{\theta}^{\text{GFD}}$ estimates. The estimator targets the structural payoff parameter θ_0 governing $u(x, d; \theta_0) = z(x, d)^\top \theta_0$. Under exact ρ-period FD with $\mathbf{\Phi}$ feasible under the true f, Theorem 4 gives $\widetilde{v}^{\rho\text{-FD}}(x, d;\, \theta_0, \mathbf{\Phi}, \mathbf{p}) = v(x, d; \theta_0) - v(x, 0; \theta_0)$, so $\Lambda_d(x;\, \theta_0, \mathbf{\Phi}, \mathbf{p}) = p_0(d \mid x)$ at the true parameter and the population GFD log-likelihood (24) is the correctly specified conditional multinomial log-likelihood. The GFD estimator inherits the standard MLE consistency: $\widehat{\theta}^{\text{GFD}}(\widehat{\mathbf{\Phi}}) \xrightarrow{p} \theta_0$ for any FD-feasible $\widehat{\mathbf{\Phi}}$ and any $\sqrt{N}$-consistent first-stage $(\widehat{f}, \widehat{p})$.

What $\widehat{\mathbf{\Phi}}$ does in the estimator. The flow input $\widehat{\mathbf{\Phi}}$ enters (24) through the GFD pseudo value-difference (22), which packs the ρ-period continuation cancelation into a finite weighted sum over paths. *The estimator does not solve the Bellman equation at any θ.* Different feasible choices of $\widehat{\mathbf{\Phi}}$ produce different $\widehat{\theta}^{\text{GFD}}(\widehat{\mathbf{\Phi}})$, all consistent for θ_0 but with different asymptotic variances; choosing $\widehat{\mathbf{\Phi}}$ to minimize the asymptotic-variance trace is the subject of Section 4.

3.5 Linear-in-Parameter Form: Effective Regressors and Offsets

Under Assumption 5 the FD pseudo value-difference is linear in θ. Substituting $z(x_\tau, d_\tau)^\top\theta + \psi_{d_\tau}$ into (22) and grouping the θ-dependent and θ-independent terms:

$$\widetilde{v}^{\rho\text{-FD}}(x, d;\, \theta, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) \;=\; \widetilde{H}(x, d;\, \widehat{\mathbf{\Phi}})^\top\theta \;+\; \widetilde{h}(x, d;\, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}), \tag{25}$$

where the *effective regressor* $\widetilde{H}(x, d;\, \widehat{\mathbf{\Phi}}) \in \mathbb{R}^{d_\theta}$ and the *effective offset* $\widetilde{h}(x, d;\, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) \in \mathbb{R}$ are

$$\widetilde{H}(x, d;\, \widehat{\mathbf{\Phi}}) \;:=\; z(x, d) \;+\; \sum_{\tau=1}^{\rho} \beta^\tau \sum_{x_{1:\rho}, d_{1:\rho}} \big[\, \widehat{\phi}_{x_{1:\rho}, d_{1:\rho}}(x, d) - \widehat{\phi}_{x_{1:\rho}, d_{1:\rho}}(x, 0) \,\big]\, z(x_\tau, d_\tau), \tag{26}$$

$$\widetilde{h}(x, d;\, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) \;:=\; \sum_{\tau=1}^{\rho} \beta^\tau \sum_{x_{1:\rho}, d_{1:\rho}} \big[\, \widehat{\phi}_{x_{1:\rho}, d_{1:\rho}}(x, d) - \widehat{\phi}_{x_{1:\rho}, d_{1:\rho}}(x, 0) \,\big]\, \widehat{\psi}_{d_\tau}(x_\tau). \tag{27}$$

With this decomposition the GFD log-likelihood (24) becomes a standard logit log-likelihood with generated regressor $\widetilde{H}(x, d;\, \widehat{\mathbf{\Phi}})$ and generated offset $\widetilde{h}(x, d;\, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}})$:

$$\widehat{\theta}^{\text{GFD}}\big(\widehat{\mathbf{\Phi}}\big) \;=\; \arg\max_{\theta\in\Theta}\; \frac{1}{NT}\sum_{i,t} \log \frac{\exp\big(\widetilde{H}(x_{it}, d_{it})^\top\theta + \widetilde{h}(x_{it}, d_{it})\big)}{\sum_{d''} \exp\big(\widetilde{H}(x_{it}, d'')^\top\theta + \widetilde{h}(x_{it}, d'')\big)}, \tag{28}$$

which is concave in θ and admits standard Newton–Raphson or BFGS optimization.

3.6 The Two-Step GFD Procedure

Putting the pieces together, the full GFD procedure is:

Step 1 (*First stage*). From the panel $\{(x_{it}, d_{it})\}_{i,t}$, compute nonparametric estimates $\widehat{f}(x' \mid x, d)$ and $\widehat{p}(d \mid x)$ together with the plug-in Hotz–Miller correction $\psi_d(\widehat{\mathbf{p}}(x))$.

Step 2 (*Solve for the flow input $\widehat{\mathbf{\Phi}}$*). For each $x \in \mathcal{X}$, solve the sparse joint feasibility problem

$$\min_{\{\Phi(x, d_0)\}_{d_0\in\mathcal{D}}} \;\sum_{d=1}^{D-1} \sum_{x'\in\mathcal{X}} \Big(\widehat{\kappa}_{\rho+1}(x' \mid x, d) - \widehat{\kappa}_{\rho+1}(x' \mid x, 0)\Big)^2 \tag{29}$$

subject to (5)–(6) (with $\widehat{f}$) for each $\Phi(x, d_0)$. By Theorem 3 (applied pairwise against the reference $d = 0$), the minimum is zero whenever pointwise FD holds at each triple $(x, d, 0)$.

The default *canonical input* is the Moore–Penrose minimum-ℓ_2-norm solution.

Step 3 (*Construct effective regressors and offsets*). Using $\widehat{\mathbf{\Phi}}$ from Step 2 and the plug-in $\psi_d(\widehat{\mathbf{p}}(x))$ from Step 1, compute $\widetilde{H}(x, d; \widehat{\mathbf{\Phi}})$ and $\widetilde{h}(x, d; \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}})$ via (26)–(27) for every (x, d) in the data.

Step 4 (*Pseudo-likelihood maximization*). Solve the concave logit optimization (28) for $\widehat{\theta}^{\mathrm{GFD}}(\widehat{\mathbf{\Phi}})$.

Computational cost. Step 1 is $O(NT)$ in the sample size. Step 2 is dominated by $|\mathcal{X}|$ sparse linear solves of dimension $D \cdot S^\rho D^\rho$ each, scaling roughly as $|\mathcal{X}| \cdot S^\rho D^\rho$ when LSQR is used and exploited sparsity is taken into account; this cost is paid once and does not repeat across parameter evaluations. Step 3 is one weighted sum of length $S^\rho D^\rho$ per (x, d) pair. Step 4 is a low-dimensional (d_θ-variable) concave optimization. Crucially, none of Steps 2–4 solves the Bellman fixed point at any θ, in contrast to nested fixed-point methods that re-solve the Bellman equation at every parameter update during optimization.

Choice of input matters for variance, not consistency. Definition 3 produces a different estimator for each FD-feasible $\widehat{\mathbf{\Phi}}$: the canonical Moore–Penrose input gives one $\widehat{\theta}^{\mathrm{GFD}}$, a variance-optimized input (Section 4) gives another, and a non-negative AM-style input (when feasible) gives a third. All are consistent for θ_0, but the asymptotic variance $\mathbf{\Sigma}(\widehat{\mathbf{\Phi}})$ depends explicitly on the input through $\widetilde{H}$. The *variance–input* mapping $\widehat{\mathbf{\Phi}} \mapsto \mathbf{\Sigma}(\widehat{\mathbf{\Phi}})$ is the central object of Section 4.

3.7 Counterfactual Computation via Finite Dependence

The GFD identity (21) was developed for estimation, but it delivers a second payoff: when finite dependence holds, the counterfactual conditional choice probabilities under a payoff-only policy experiment can be computed by a fixed-point iteration that does not solve the Bellman equation at any point.

Definition 4 (Payoff-only counterfactual)**.** A counterfactual policy experiment is *payoff-only* if it modifies the structural payoff parameter from θ_0 to $\theta_1 \neq \theta_0$ while holding the transition kernel $f(x' \mid x, d)$ fixed. The counterfactual object of interest is the equilibrium CCP vector $\mathbf{p}_1(x) = \mathbf{p}(x; \theta_1, f)$ under the new payoff.

Payoff-only counterfactuals cover policy evaluations that change entry costs, subsidies, taxes, or other payoff shifters without altering the physical state transition. They exclude merger simulations or market-structure changes that modify f itself; for those, the linear-representation approach of Kalouptsidi et al. (2021) remains the operative tool.

Theorem 5 (Bridge: payoff-only counterfactual CCPs from FD)**.** *Suppose Assumptions 1–6 hold and that $\mathbf{\Phi}$ is FD-feasible at horizon ρ under the true transitions f (Definition 2). Consider a payoff-only counterfactual that changes $\theta_0 \mapsto \theta_1$. Then:*

(a) Fixed-point representation. *The counterfactual CCP vector* $\mathbf{p}_1(\cdot) = \mathbf{p}(\cdot;\, \theta_1, f)$ *satisfies*

$$p_1(d \mid x) \;=\; \Lambda_d\big(x;\, \theta_1, \mathbf{\Phi}, \mathbf{p}_1\big), \qquad \forall\, (x, d) \in \mathcal{X} \times \mathcal{D}, \tag{30}$$

where Λ_d *is the GFD multinomial-logit map* (23) *evaluated at the counterfactual payoff* θ_1 *and counterfactual CCPs* $\mathbf{p}_1$*. No object on the right-hand side requires evaluation of the Bellman equation at* θ_1*.*

(b) Single-agent uniqueness and local convergence. *The Bellman operator at* θ_1 *is a* β*-contraction in the sup-norm, so the counterfactual value function* V_1 *is unique and—by the Hotz–Miller bijection—the Bellman-solution CCP* $\mathbf{p}_1$ *is unique and is a fixed point of* (30)*. The CCP-space iteration* $\mathbf{p}^{(m+1)}(\cdot) = \Lambda(\widetilde{v}^{\rho\text{-}FD}(\cdot;\, \theta_1, \mathbf{\Phi}, \mathbf{p}^{(m)}))$ *converges locally to* $\mathbf{p}_1$ *when the spectral radius of the Jacobian of the composite map at* $\mathbf{p}_1$ *is below unity. Theorem 5 does* not *claim that* $\mathbf{p}_1$ *is the only fixed point of* (30)*: extraneous fixed points of the FD map that do not correspond to any Bellman solution are not ruled out by the contraction argument alone. The practical guardrail is the restart/spectral-radius diagnostic of Remark 4.*

(b′) Multi-player games. *In a Markov-perfect-equilibrium (MPE) setting with player-specific perceived transitions* $f_i^{\mathbf{p}_{-i}^{\mathrm{MPE}}}$, *every MPE profile* $\mathbf{p}_1^{\mathrm{MPE}}$ *at counterfactual payoffs* $\boldsymbol{\theta}_1 = (\theta_{1,i})_i$ *induces a fixed point of* (30) *stacked across players, provided* ρ*-period finite dependence holds at each perceived transition* $f_i^{\mathbf{p}_{-i}^{\mathrm{MPE}}}$ *(automatic in shared-action Kronecker games via the action-invariant cancellation of Lemma 11). MPE uniqueness is not delivered by the bridge; verifying that a given fixed point of* (30) *is an MPE requires a one-shot Bellman best-response check at the candidate profile—a diagnostic constructed from the primitives* (u, f, β)*, not delegated to an external equilibrium-existence theorem.*

(c) Identification of $\mathbf{p}_1$ from observables. *For the canonical Moore–Penrose flow input* $\mathbf{\Phi}^{\mathrm{MP}}(f)$*, the right-hand side of* (30) *is a function of* $(\theta_1, f, \mathbf{p}_1)$ *only. For the variance-optimal flow input* $\mathbf{\Phi}^{\mathrm{opt}}(\theta_0, \mathbf{p}_0, f)$ *of Section 4.4, an additional dependence on the estimable DGP triple* $(\theta_0, \mathbf{p}_0)$ *enters and must be plugged in via a* $\sqrt{N}$*-consistent first-stage estimate.*

Proof sketch. Part (a): substitute (21) (evaluated at the counterfactual payoff θ_1 and CCPs $\mathbf{p}_1$) into the Hotz–Miller logit-CCP map at θ_1. Part (b): standard Bellman contraction plus Hotz–Miller bijection deliver uniqueness; Theorem 4 then guarantees that the unique Bellman solution satisfies (30); local convergence follows from the spectral-radius condition on the FD-map Jacobian. Part (b′) applies Theorem 4 player by player at the equilibrium perceived transitions; the equilibrium continuation value is exactly the object the FD-weighted sum reproduces under exact FD at $f_i^{\mathbf{p}_{-i}^{\mathrm{MPE}}}$. Part (c) follows by tracing which population objects $\mathbf{\Phi}$ depends on under each selector rule. □

Remark 4 (Convergence and validation diagnostics)**.** *Two practical diagnostics screen out extraneous fixed points and confirm that the FD iteration has converged to the Bellman-solution CCP.*

(i) Restart from multiple initializations. *Run* (30) *from at least five starts—the baseline observational CCP, plus a spread of constant vectors—and verify that all converge to the same limit; convergence to disjoint limits flags a non-Bellman fixed point. (ii)* Spectral radius at the converged point. *Compute the operator-norm Jacobian of the composite map* $\mathbf{p} \mapsto \Lambda(\widetilde{v}^{\rho\text{-}FD}(\cdot;\, \theta_1, \mathbf{\Phi}, \mathbf{p}))$ *at the converged* $\mathbf{p}_1$*; values above* 0.9 *recommend increasing* ρ*. Section 6.2 implements both diagnostics on the 3-player entry/exit DGP at a counterfactual entry-cost increase of* 50%*: starting from five initializations (the baseline observational CCP plus four constant profiles spanning the simplex from* $p(\textit{enter}) = 0.05$ *to* $p(\textit{enter}) = 0.95$*), all five iterations converge to the same FD fixed point* $\mathbf{p}_1^{\mathrm{FD}}$ *in* 12–14 *iterations with maximum pairwise difference* 8.7×10^{-7} *across the five limits. The converged FD CCP agrees with the NFXP MPE to maximum absolute error* 8.2×10^{-3}*, which reflects the finite-precision residual of the rank-deficient flow QP at the counterfactual perceived transition (LSQR terminal mismatch* $\|\mathbf{C}\widehat{\mathbf{\Phi}}\| \in [1.5, 6.6] \times 10^{-4}$ *across the three players); a higher-precision flow QP (direct KKT or larger* ρ*) is what closes the residual. Reproducible from* `python3 code/test_bridge_counterfactual.py`*.*

Relation to identification/computation literature. Arcidiacono and Miller (2020) show that counterfactual CCPs are generically not identified from short panels unless finite dependence holds. Kalouptsidi et al. (2021) (KSSR) provide a general null-space compatibility condition for counterfactual identification under arbitrary linear payoff transformations, including transition-changing counterfactuals. Theorem 5 is complementary to both: it gives a checkable computational route within the payoff-only class, combining the FD identity with restart-based validation. Uniqueness of $\mathbf{p}_1$ itself is inherited from the Bellman contraction; the contribution of the bridge is the iteration plus its verification scaffold, not a new uniqueness theorem.

3.8 Extension to Unobserved Heterogeneity

Many applied DDC models feature unobserved heterogeneity, with agents of unobserved type $\ell \in \{1, \ldots, L\}$ differing in either preferences $\theta^{(\ell)}$ or transitions $f^{(\ell)}(x' \mid x, d)$, or both. Following Arcidiacono and Miller (2011) and the nonparametric identification results of Kasahara and Shimotsu (2009), GFD integrates into an Expectation–Maximization framework with a substantial computational saving over AM-style EM.

The observed CCP is the type mixture $p(d \mid x) = \sum_{\ell=1}^{L} \pi_\ell \, p_\ell(d \mid x)$, where π_ℓ is the type prevalence and $p_\ell(\cdot \mid x)$ is the type-ℓ optimal CCP. The econometrician does not observe ℓ directly but infers it from the choice history.

Proposition 6 (Type-specific finite dependence)**.** *Suppose agents are heterogeneous with type-specific transitions* $F_d^{(\ell)}$*.*

(a) Common transitions. *If all types share the transition structure* $F_d^{(\ell)} = F_d$ *for every* ℓ *and differ only in preferences* $\theta^{(\ell)}$*, then a single FD-feasible flow input* $\mathbf{\Phi}(f)$ *works for every* *type. The pointwise feasibility test of Theorem 3 and the minimal horizon $\rho^\star$ are identical across types.*

(b) Type-specific transitions. *If types differ in both preferences and transitions, the feasibility test must be verified type by type. Let $\rho^\star_\ell$ denote the minimal FD horizon for type ℓ. The model admits FD estimation at horizon $\rho = \max_\ell \rho^\star_\ell$ using type-specific flow inputs $\mathbf{\Phi}^{(\ell)}(f^{(\ell)})$.*

Proof. Part (a): the linear feasibility system $\mathbf{Ax} = \mathbf{b}$ of Theorem 3 depends only on f, so a single $\mathbf{\Phi}$ satisfies the matching condition for every ℓ; the augmented payoff $\widehat{u}(x, d; \theta^{(\ell)}, \mathbf{p}_\ell) = z(x, d)^\top \theta^{(\ell)} + \psi_d(\mathbf{p}_\ell(x))$ varies across types but enters only the value of the GFD identity, not its validity. Part (b): when $F^{(\ell)} \neq F^{(\ell')}$, the constraint matrix differs, so feasibility is verified per type; the pointwise rank condition is monotone in ρ along the recursive structure of (29). $\square$

Part (a) covers the common applied specification in which unobserved types index preference heterogeneity (e.g. taste for work, discount-rate heterogeneity) while physical state transitions are type-invariant. In that case, GFD applies with no additional computational overhead per type: the flow input $\widehat{\mathbf{\Phi}}$ is solved once, the effective regressor $\widetilde{H}^{(\ell)}$ requires only the type-specific payoff regressor (which is shared across types up to $\theta^{(\ell)}$ when payoffs are linear in θ), and the only per-type, per-iteration update is the Hotz–Miller offset $\widetilde{h}^{(\ell),(m)}(d, x)$.

GFD-EM algorithm. Adapting the AM EM scheme of Arcidiacono and Miller (2011) to the GFD identity yields the following inner loop. *Initialize* type probabilities $\widehat{\pi}^{(0)}_\ell$ and type-specific CCPs $\widehat{p}^{(0)}_\ell(d \mid x)$ from a finite-mixture model on the observed choice frequencies. *Pre-loop* (computed once, outside the EM iteration): solve the flow QP for $\widehat{\mathbf{\Phi}}^{(\ell)}$ under $\widehat{f}^{(\ell)}$ (or once across types under common transitions), and form the effective regressors $\widehat{H}^{(\ell)}$. *Iterate* for $m = 1, 2, \ldots$ until convergence:

(E) *Posterior types.* For each i compute $q^{(m)}_{i\ell} = \widehat{\pi}^{(m-1)}_\ell \prod_t \widehat{p}^{(m-1)}_\ell(d_{it} \mid x_{it}) / \sum_{\ell'} \widehat{\pi}^{(m-1)}_{\ell'} \prod_t \widehat{p}^{(m-1)}_{\ell'}(d_{it} \mid x_{it})$, and refresh the type-specific Hotz–Miller offsets $\widehat{h}^{(\ell),(m)}$ from $\widehat{p}^{(m-1)}_\ell$.

(M) *Pseudo-MLE.* Maximize the type-weighted GFD pseudo-log-likelihood

$$\widehat{\theta}^{(m)} \;=\; \arg\max_\theta \sum_{i,t} \sum_{\ell=1}^{L} q^{(m)}_{i\ell} \log \Lambda_{d_{it}}\big(x_{it};\, \theta^{(\ell)}, \widehat{\mathbf{\Phi}}^{(\ell)}, \widehat{\mathbf{p}}^{(m-1)}_\ell\big).$$

(U) *Update mixture and CCPs.* Set $\widehat{\pi}^{(m)}_\ell = N^{-1} \sum_i q^{(m)}_{i\ell}$ and $\widehat{p}^{(m)}_\ell(d \mid x) = \Lambda_d(x;\, \widehat{\theta}^{(m)}, \widehat{\mathbf{\Phi}}^{(\ell)}, \widehat{\mathbf{p}}^{(m-1)}_\ell)$.

Computational advantage over AM-EM. The structural saving is in step (E)/(M): the regressors $\widehat{H}^{(\ell)}$ depend on $(\widehat{f}^{(\ell)}, \widehat{\mathbf{\Phi}}^{(\ell)})$ but not on the iterating CCPs, so they are constructed once before the EM loop and reused across iterations. Only the offsets $\widehat{h}^{(\ell),(m)}$—an $O(|\mathcal{X}| \cdot \rho)$

Hotz–Miller pass per type—are refreshed each iteration. AM-EM, by contrast, requires the matrix inversion $(I - \beta F_{\widehat{\mathbf{p}}})^{-1}$ at every iteration because the CCP-weighted transition $F_{\widehat{\mathbf{p}}} = \sum_d \widehat{p}(d \mid \cdot)\, F_d$ depends on the iterating CCPs; this is an $O(|\mathcal{X}|^3)$ operation that compounds across iterations. When $L \cdot |\mathcal{X}|$ is large, GFD-EM's per-iteration cost is asymptotically negligible compared with AM-EM's. Convergence diagnostics (observed-data log-likelihood monotonicity across EM iterations, and multiple initializations) are recommended in the standard EM fashion when the per-step weights implied by $\widehat{\mathbf{\Phi}}^{(\ell)}$ are signed: monotonicity is a clean Jensen-based property only in the non-negative-weight, fixed-weight regime.

4 Variance–Input Mapping and Efficient GFD

The GFD estimator of Section 3 is defined for any FD-feasible flow input $\widehat{\mathbf{\Phi}}$, and is consistent for θ_0 regardless of which feasible input is used. The QP (29), however, is generically under-determined: at horizons $\rho \geq 2$ or in models with Kronecker-separable structure, the feasible set is a non-trivial affine variety, so multiple distinct flow inputs all satisfy Definition 2. This section asks: among the multiple feasible solutions, which one yields the asymptotically most *efficient* GFD estimator?

4.1 Non-uniqueness of the FD-feasible Flow Input

Fix $x \in \mathcal{X}$ and consider the joint constraint system enforced by the QP (29). The joint decision variable stacks the D flow matrices $\{\Phi(x, d_0)\}_{d_0 \in \mathcal{D}}$:

$$\operatorname{vec} \mathbf{\Phi}(x) \;:=\; \big(\operatorname{vec} \Phi(x,0)^\top, \ldots, \operatorname{vec} \Phi(x, A-1)^\top\big)^\top \;\in\; \mathbb{R}^{D \cdot S^\rho D^\rho}.$$

The constraints (5)–(6) for each $\Phi(x, d_0)$ together with the common-terminal condition (20) form a single block linear equation in $\operatorname{vec} \mathbf{\Phi}(x)$. The total row count is $(2D-1)\,S + D\,C$ with $C = \sum_{\tau=1}^{\rho-1} S^{\tau+1} D^\tau$, versus a column count of $D \cdot S^\rho D^\rho$ that is typically much larger.

Definition 5 (Feasible variety)**.** The *feasible variety* at x under f, denoted $\mathcal{F}(x; f)$, is the set of $\mathbf{\Phi}(x)$ satisfying the constraint system above. When $\mathcal{F}(x; f) \neq \emptyset$—which holds whenever pointwise FD of Theorem 3 holds at every triple $(x, d, 0)$, $d \neq 0$—it is an affine subspace whose dimension equals the column count minus the rank of the constraint operator.

The dimension of $\mathcal{F}(x; f)$ is typically substantial at shift-register and Kronecker models with moderate ρ. Every point of $\mathcal{F}(x; f)$ gives a valid input for the GFD estimator (Definition 3); the family $\{\widehat{\theta}^{\mathrm{GFD}}(\widehat{\mathbf{\Phi}}) : \widehat{\mathbf{\Phi}} \in \mathcal{F}(\cdot; \widehat{f})\}$ is therefore a continuum of consistent estimators of θ_0, parameterized by the input.

Null-space coordinates. Let $\mathbf{\Phi}^{\mathrm{MP}}(x; f) \in \mathcal{F}(x; f)$ denote the Moore–Penrose minimum-ℓ_2-norm element (the *canonical input*), and let $\mathbf{N}(x; f)$ be a basis for the null space of the constraint operator. Every feasible flow input has the representation

$$\mathbf{\Phi}(x;\mathbf{q}) \;=\; \mathbf{\Phi}^{\mathrm{MP}}(x;f) + \mathbf{N}(x;f)\,\mathbf{q}, \tag{31}$$

where the null-space coordinate $\mathbf{q}$ ranges over a Euclidean space whose dimension equals that of $\mathcal{F}(x;f)$. We treat $(\mathbf{\Phi}^{\mathrm{MP}},\mathbf{N})$ as deterministic functions of f (computed once and held fixed), and let the user-tunable parameter be the vector $\mathbf{q}$.

4.2 Asymptotic Distribution of the GFD Estimator

Before turning to efficient choice of $\mathbf{q}$, we record the asymptotic distribution of $\widehat{\theta}^{\mathrm{GFD}}(\widehat{\mathbf{\Phi}})$ for an arbitrary FD-feasible input. The proof is a standard two-step M-estimator argument and is deferred to Appendix B.

Theorem 7 (Asymptotic normality of the GFD estimator)**.** *Suppose Assumptions 1–6 hold; the first-stage estimates $\widehat{f}$ and $\widehat{p}$ are $\sqrt{N}$-consistent with asymptotically linear influence-function representations (Appendix B Assumption B.2); $\widehat{\mathbf{\Phi}}$ is FD-feasible under $\widehat{f}$ for every N and $\widehat{\mathbf{\Phi}} \xrightarrow{p} \mathbf{\Phi}^*$ for some $\mathbf{\Phi}^*$ FD-feasible under f; and the GFD information matrix $\mathbf{J}(\mathbf{\Phi}) := -\mathbb{E}[\nabla_\theta^2 \log \Lambda_{d_t}(x_t;\theta_0,\mathbf{\Phi},\mathbf{p})]$ is non-singular at $\mathbf{\Phi}^*$. Then*

$$\sqrt{N}\big(\widehat{\theta}^{\mathrm{GFD}}(\widehat{\mathbf{\Phi}}) - \theta_0\big) \;\xrightarrow{d}\; \mathcal{N}\big(\mathbf{0},\,\mathbf{\Sigma}(\mathbf{\Phi}^*)\big), \tag{32}$$

where $\mathbf{\Sigma}(\mathbf{\Phi}) := \mathbf{J}(\mathbf{\Phi})^{-1}\,\mathbf{V}(\mathbf{\Phi})\,\mathbf{J}(\mathbf{\Phi})^{-1}$ is the standard two-step sandwich variance, with $\mathbf{V}(\mathbf{\Phi})$ the variance of the GFD score augmented by the first-stage plug-in correction (full formula in Appendix B).

Since $\widetilde{H}(x,d;\mathbf{\Phi})$ is affine in $\mathbf{\Phi}$ by (25), $\mathbf{\Sigma}(\mathbf{\Phi})$ is a smooth (rational) function of $\mathbf{\Phi}$.

4.3 The Variance–Input Mapping

Composing the parameterization (31) with the asymptotic-variance map $\mathbf{\Sigma}(\cdot)$ from Theorem 7 defines the *variance–input mapping*

$$\mathbf{q} \;\longmapsto\; \mathbf{\Sigma}\big(\mathbf{\Phi}(\mathbf{q})\big), \tag{33}$$

a smooth (rational) function on $\mathbb{R}^{d_q}$. The key identification of this section is that asymptotic efficiency among GFD estimators is the problem of *minimizing this mapping over* $\mathbf{q}$. The natural scalar criterion is the trace

$$\mathbf{q} \;\longmapsto\; \mathrm{tr}\big(\mathbf{\Sigma}(\mathbf{\Phi}(\mathbf{q}))\big), \tag{34}$$

the sum of asymptotic variances of the individual coordinates of $\widehat{\theta}^{\mathrm{GFD}}$.

Remark 5 (Non-convexity and multi-start)**.** *Because $\widetilde{H}$ is affine in $\mathbf{q}$, the GFD information matrix and score variance underlying $\boldsymbol{\Sigma}$ are quadratic in $\mathbf{q}$, and the trace map* (34) *is rational. Convexity is* not *guaranteed: at high-dimensional null spaces $\mathbb{R}^{d_q}$ the trace can have multiple stationary points. In practice we recommend solving the inner optimization* (35) *below from several random starts (the Moore–Penrose anchor $\mathbf{q} = \mathbf{0}$ being one) and selecting the smallest-trace optimum. The continuity of the trace on the regular sub-region $\mathcal{Q}$ ensures the global minimum is attained (Theorem 8); finding it with certainty is the standard issue for non-convex smooth optimization.*

4.4 The Variance-Optimal GFD Estimator

Definition 6 (Variance-optimal flow input)**.** The *variance-optimal flow input* (under the trace criterion) is

$$\boldsymbol{\Phi}^{\mathrm{opt}} \ := \ \boldsymbol{\Phi}^{\mathrm{MP}} + \mathbf{N}\,\mathbf{q}^{\mathrm{opt}}, \qquad \mathbf{q}^{\mathrm{opt}} \ := \ \arg\min_{\mathbf{q}\in\mathcal{Q}} \operatorname{tr}\big(\boldsymbol{\Sigma}(\boldsymbol{\Phi}(\mathbf{q}))\big), \tag{35}$$

where $\mathcal{Q} \subseteq \mathbb{R}^{d_q}$ is a compact regular sub-region on which $\mathbf{J}(\boldsymbol{\Phi}(\mathbf{q}))$ is uniformly bounded away from singularity (Assumption B.1 in Appendix B).

Theorem 8 (Existence and FD-class trace optimality of $\boldsymbol{\Phi}^{\mathrm{opt}}$)**.** *Suppose the conditions of Theorem 7 hold and the compact regular sub-region $\mathcal{Q} \subseteq \mathbb{R}^{d_q}$ is non-empty (Appendix B Assumption B.1). Then:*

- *(i)* Existence. *The minimum in* (35) *is attained at some $\mathbf{q}^{\mathrm{opt}} \in \mathcal{Q}$, and the corresponding $\boldsymbol{\Phi}^{\mathrm{opt}} = \boldsymbol{\Phi}^{\mathrm{MP}} + \mathbf{N}\,\mathbf{q}^{\mathrm{opt}} \in \mathcal{F}(\cdot;\, f) \cap \mathcal{Q}$ is a variance-optimal flow input.*
- *(ii)* First-order condition. *Every interior minimizer $\mathbf{q}^{\mathrm{opt}} \in \operatorname{int}(\mathcal{Q})$ satisfies $\nabla_{\mathbf{q}} \operatorname{tr}\big(\boldsymbol{\Sigma}(\boldsymbol{\Phi}(\mathbf{q}^{\mathrm{opt}}))\big) = \mathbf{0}$.*
- *(iii)* FD-class trace optimality. *The GFD estimator at input $\boldsymbol{\Phi}^{\mathrm{opt}}$ attains the trace-minimum asymptotic variance over the class of GFD estimators at horizon ρ:*

$$\operatorname{tr}\big(\boldsymbol{\Sigma}(\boldsymbol{\Phi}^{\mathrm{opt}})\big) \ = \ \min_{\boldsymbol{\Phi}\in\mathcal{F}(\cdot;\, f)\cap\mathcal{Q}} \operatorname{tr}\big(\boldsymbol{\Sigma}(\boldsymbol{\Phi})\big). \tag{36}$$

Proof. (i) The trace map $\mathbf{q} \mapsto \operatorname{tr}(\boldsymbol{\Sigma}(\boldsymbol{\Phi}(\mathbf{q})))$ is continuous on the compact set $\mathcal{Q}$ (rational with no singularities on $\mathcal{Q}$ by Assumption B.1, which keeps $\mathbf{J}$ uniformly non-singular). The extreme value theorem gives existence of a minimizer. (ii) Standard FOC at an interior optimum of a smooth objective. (iii) Restatement of Definition 6: the minimum is attained at $\boldsymbol{\Phi}^{\mathrm{opt}}$ by part (i), so its trace is the infimum. □

Theorem 8 is the paper's strongest *unconditional* efficiency result for the GFD class: no further controllability, rank, or instrument-Jacobian assumptions are imposed—only the FD-feasibility of the input and the regularity of the sub-region $\mathcal{Q}$. The non-convexity caveat of

Remark 5 applies to *finding* the minimum numerically; existence and the value of the minimum are unconditional.

Sample-feasible procedure. Definition 6 is a population concept; the sample-feasible variant replaces f with $\widehat{f}$ and the population variance with a plug-in estimator $\widehat{\boldsymbol{\Sigma}}$, evaluated at a preliminary $\widehat{\theta}^{(0)}$ (the GFD estimator using $\widehat{\boldsymbol{\Phi}}^{\mathrm{MP}}$). The full algorithm is:

2a. Compute $\widehat{\boldsymbol{\Phi}}^{\mathrm{MP}}$ from (29) (canonical input).

2b. Compute the preliminary estimate $\widehat{\theta}^{(0)} := \widehat{\theta}^{\mathrm{GFD}}(\widehat{\boldsymbol{\Phi}}^{\mathrm{MP}})$ via the four-step procedure of Section 3.6.

2c. Form the plug-in sandwich variance $\widehat{\boldsymbol{\Sigma}}(\boldsymbol{\Phi}(\mathbf{q}); \widehat{\theta}^{(0)})$ using sample analogues of the population sandwich pieces in Theorem 7 (Appendix B for explicit formulas).

2d. Solve $\widehat{\mathbf{q}}^{\mathrm{opt}} := \arg\min_{\mathbf{q}} \operatorname{tr}\big(\widehat{\boldsymbol{\Sigma}}(\boldsymbol{\Phi}(\mathbf{q}); \widehat{\theta}^{(0)})\big)$ by quasi-Newton in $\mathbf{q}$.

2e. Set $\widehat{\boldsymbol{\Phi}}^{\mathrm{opt}} := \widehat{\boldsymbol{\Phi}}^{\mathrm{MP}} + \mathbf{N}\,\widehat{\mathbf{q}}^{\mathrm{opt}}$ and re-run Steps 3–4 of Section 3.6 to obtain the final $\widehat{\theta}^{\mathrm{GFD}}(\widehat{\boldsymbol{\Phi}}^{\mathrm{opt}})$.

The replacement of Steps 2–4 by Steps 2a–2e–3–4 adds one inner optimization over $\mathbf{q}$ and one re-evaluation of the logit MLE; both are standard low-dimensional smooth optimizations and add a modest constant factor to total computational cost.

4.5 Conditional Extension: ρ-Horizon Semiparametric Efficiency

Theorem 8 states optimality *within* the GFD class. A natural follow-up question is whether $\boldsymbol{\Sigma}(\boldsymbol{\Phi}^{\mathrm{opt}})$ also matches the broader semiparametric efficiency bound for the class of *all* ρ-horizon CCP-based estimators.

The broader class. Define $\mathcal{E}_\rho$ as the class of regular two-step estimators of θ_0 that use only the joint distribution of the data window $(x_{it}, d_{it}, x_{i,t+1}, \ldots, x_{i,t+\rho})$, with first-stage nonparametric estimation of f and p. Every GFD estimator at any feasible $\widehat{\boldsymbol{\Phi}}$ belongs to $\mathcal{E}_\rho$, but $\mathcal{E}_\rho$ also includes other ρ-horizon CCP estimators (e.g., truncated Hotz–Miller, sieve-based pseudo-likelihoods). The Chamberlain (1987) semiparametric efficiency bound for $\mathcal{E}_\rho$ is the smallest asymptotic variance achievable in this class, denoted $\boldsymbol{\Sigma}_\rho^{\mathrm{eff}}$.

The conditional efficiency claim. Whether $\boldsymbol{\Sigma}(\boldsymbol{\Phi}^{\mathrm{opt}}) = \boldsymbol{\Sigma}_\rho^{\mathrm{eff}}$ depends on whether the family of FD inputs $\mathcal{F}(\cdot; f)$ is rich enough to span the efficient ρ-horizon influence function and on properties of the first-stage variance correction. We isolate this in four sufficient conditions:

(G1) *Full controllability.* For every triple (x_0, d, d') with $d \neq d'$, the pair-specific controllability matrix $\widetilde{\mathcal{C}}_\rho^{(x_0,d,d')}$ has full rank $S-1$, i.e., the achievable terminal-distribution perturbations span the entire $(S-1)$-dim simplex tangent. This is the standard linear-system rank condition; it holds for renewal, shift-register, and a class of Kronecker-separable models, and fails generically only when transitions are degenerate (e.g., rank-one).

(G2) *Instrument-Jacobian rank.* The mapping $\mathbf{q} \mapsto \widetilde{H}(\cdot; \boldsymbol{\Phi}(\mathbf{q}))$ has Jacobian of full row rank d_θ at $\mathbf{q}^{\text{opt}}$ (equivalently, the directional derivatives $\partial\widetilde{H}/\partial q_k$ jointly span a d_θ-dim subspace of $L^2(p_0)$).

(G3) *FD-coincidence.* There exists $\boldsymbol{\Phi}^\dagger \in \mathcal{F}(\cdot; f)$ at which the GFD-induced asymptotic variance equals the broader ρ-horizon efficiency bound: $\boldsymbol{\Sigma}(\boldsymbol{\Phi}^\dagger) = \boldsymbol{\Sigma}_\rho^{\text{eff}}$. This is the load-bearing alignment condition: it asserts that some feasible flow input realizes the efficient ρ-horizon influence function. It is *not* implied by (G1)+(G2): spanning richness of the FD instrument family does not automatically imply that some member realizes the efficient information.

(G4) *Oracle equivalence.* The first-stage plug-in correction in (32) does not strictly inflate the GFD variance relative to its oracle (known-first-stage) counterpart. For cell-frequency first-stage CCP estimators this follows from the multinomial-distribution efficiency identity of Chamberlain (1987, Lemma 2 + Theorem 3); for sieve or kernel first-stage estimators it requires a separate verification.

Corollary 9 (Conditional ρ-horizon semiparametric efficiency)**.** *Under the conditions of Theorem 7 together with gates (G1)–(G4) at the data-generating process,*

$$\boldsymbol{\Sigma}(\boldsymbol{\Phi}^{\text{opt}}) \;=\; \boldsymbol{\Sigma}_\rho^{\text{eff}}. \tag{37}$$

Status of the four gates. (G1) is verifiable from transition primitives alone (a linear-algebraic rank check). (G2) is generic in the sense that the rank-deficiency set in the parameter space has Lebesgue measure zero, and is numerically verifiable at the estimated $\widehat{\mathbf{q}}^{\text{opt}}$. (G3) is the load-bearing condition; it must be checked on a per-DGP basis (numerically or, for special structures, algebraically). (G4) is automatic for cell-frequency first-stage CCPs at the population point under exact FD; for other first-stage schemes it requires a separate identity check.

Practical reading. Theorem 8 (FD-class trace optimality) is the unconditional result and the default benchmark: at any DGP where the GFD framework applies, the variance-optimal input $\boldsymbol{\Phi}^{\text{opt}}$ delivers the smallest-trace GFD estimator. Corollary 9 promotes this to ρ-horizon semiparametric efficiency *contingent on* (G1)–(G4), which are non-trivial conditions verified per DGP. The interpretation is deliberately layered: trace optimality within the GFD class is universal; matching the broader Chamberlain bound is a stronger conditional claim.

5 Applications

This section verifies the GFD framework on the suite of ten canonical DDC models implemented in the accompanying code library. For each model we identify the transition structure, predict the minimal FD horizon $\rho^\star$ from primitives where possible, and confirm the prediction by *two* independent computations: (i) a direct Bellman solve for the true value-difference, used as a reference (`solve_mdp` in `HistoryDependentFDSolver.py`), and (ii) the pointwise linear-feasibility test of Theorem 3 (`fd_holds` in `fd_linear_system.py`). Agreement between the two computations across all ten models confirms both the rank test and the GFD value-difference identity (Theorem 4) to machine precision.

5.1 A Taxonomy of FD-Amenable Transition Structures

We distinguish four transition structures that organize the ten canonical models in Section 5.2.

Type R (Renewal). One reference action sets the next-period state distribution independently of the current state—that is, $f(\cdot \mid x, 0) = f^r(\cdot)$ for all x, so the transition matrix $\mathbf{F}_0$ has rank one. After a single application of the renewal action, the system forgets the initial state. Pointwise FD at $(x_0, d, 0)$ holds at $\rho^\star = 1$ for any analyzed action d, and the QP (29) admits the classical AM non-negative solution.

Type S (Shift register, p-lag). The state factorizes into a p-component lag register of past actions $(d_{t-p}, d_{t-p+1}, \ldots, d_{t-1})$ together with an exogenous component evolving under an action-independent kernel. The current action d_t enters the next state *only* through the lag register, so any two initial actions d, d' become indistinguishable after p shifts have flushed them out.

Proposition 10 (FD horizon for shift-register models). *For Type-S models with p-lag register, pointwise FD holds at every triple (x_0, d, d') at $\rho^\star = p$. The QP solution is the trivial flush: at $\rho = p$, the flow matrices $\Phi(x_0, d)$ and $\Phi(x_0, d')$ that replay the same future action sequence $(d_1, \ldots, d_p)$ produce identical terminal registers $(d_1, \ldots, d_p)$, independent of the initial action.*

Type K (Kronecker with action-invariant exogenous factor). The state factorizes as $x_t = (x_t^{\text{endo}}, x_t^{\text{exo}})$ with Kronecker transition $\mathbf{F}_d = \mathbf{F}_d^{\text{endo}} \otimes \mathbf{F}^{\text{exo}}$, where the exogenous factor is *action-invariant* ($\mathbf{F}^{\text{exo}}$ does not depend on d). Capital accumulation with exogenous productivity, or inventory with exogenous demand, are the canonical examples.

Lemma 11 (Action-invariant Kronecker cancellation). *Suppose $\mathbf{F}_d = \mathbf{F}_d^{\text{endo}} \otimes \mathbf{F}^{\text{exo}}$ with $\mathbf{F}^{\text{exo}}$ action-invariant. Then for every triple (x_0, d, d'), pointwise FD on the joint state holds at horizon $\rho^\star$ if and only if pointwise FD on the x^{endo}-projected problem holds at $\rho^\star$. The joint flow input factorizes as $\Phi(x_0, d_0) = \Phi^{\text{endo}}(x_0^{\text{endo}}, d_0) \otimes \Phi^{\text{exo}}(x_0^{\text{exo}})$, where Φ^{exo} is the action-invariant evolution of x^{exo}.*

When the action-dependent endogenous block has *strong control* (every action moves x^{endo} by a one-step shift on a grid), the projected problem is FD-feasible at $\rho^\star = 1$, and the joint problem inherits the same horizon despite arbitrary persistence in the exogenous factor. Lemma 11 is the structural mechanism behind the $\rho^\star = 1$ verifications for Investment and Inventory in Section 5.2.

Type D (Triple Connector / Strategic Ripple). Multi-component states whose components are *all* action-dependent (e.g., entry/exit with action-dependent productivity), or shared-action games where opponents' equilibrium strategies enter the per-player transition. No single component is renewing, no factor is action-invariant, and the state is not a lag register. Type-D models generically require signed flow inputs (Remark 1); $\rho^\star$ must be determined by direct verification with Theorem 3.

Remark 6 (Pointwise vs. universal FD in Type-D models)**.** *For Type-D* single-agent *models with endogenous productivity (Entry/Exit, Endogenous), pointwise FD can hold at lower ρ than the universal-FD/controllability-rank prediction would suggest, because the propagated distributional gaps lie in a low-dimensional subspace of the controllability column space even when that column space does not span the full $(S-1)$-dim simplex tangent. The Type-D* multi-player *games below (the entry/exit game and the dynamic oligopoly) are different: any pointwise $\rho = 1$ certificate obtained from the per-player projection is an artifact of projecting out opponents' equilibrium states and does not reflect the true equilibrium-induced strategic ripple, which acts at horizon two through opponents' best-response feedback into the player's own continuation. The operative horizon for any multi-player game in this class is therefore $\rho^\star = 2$, regardless of what the projected pointwise test reports at $\rho = 1$.*

5.2 Verification on Ten Canonical Models

We verify the framework on ten canonical DDC models spanning the four transition structures of Section 5.1. The models are:

1. *Engine Replacement* (Rust, 1987): state $x \in \{0, 1, \ldots, 9\}$ records discrete mileage bins; actions $d \in \{0, 1\}$ are continue and replace, with replacement deterministically resetting to bin 0 ($\mathbf{F}_1$ rank-one).

2. *Job Search*: state x is the offered wage in a discrete grid; actions are accept and reject, with rejection drawing a fresh i.i.d. wage offer next period.

3. *Female Labor* ($p = 1$): state $x = (\ell_{t-1}) \in \{0, 1\}$ is the lagged labor-supply decision; transitions are deterministic shift-register updates.

4. *Investment*: state $x = (k, z)$ pairs capital k with exogenous productivity z on a Tauchen grid (Tauchen, 1986); three actions (depreciate, maintain, invest) shift capital by at most one grid point while z evolves as an action-invariant AR(1).

5. *Inventory Control*: state $x = (s, \delta)$ pairs stock s with exogenous demand δ; three actions (idle, order-low, order-high) shift stock while demand follows an action-invariant AR(1).

6. *Entry/Exit with endogenous productivity*: state $x = (m, \omega, z)$ combines market status, productivity, and an exogenous shock; entry and exit shift both market status and the productivity distribution (triple connector).

7. *Altug–Miller* ($p = 3$) (Altug and Miller, 1998): state $x = (\ell_{t-3}, \ell_{t-2}, \ell_{t-1}) \in \{0, 1\}^3$ is a $p = 3$ shift register over labor history.

8. *Education* ($p = 4$): state $x = (d_{t-4}, \ldots, d_{t-1})$ is a $p = 4$ shift register over schooling decisions, in the spirit of Keane and Wolpin (1997).

9. *Entry/Exit Game*: 3-player Markov-perfect dynamic entry/exit game with state $x = (z, y_1, y_2, y_3)$ pairing market size z with each player's incumbency status y_i; per-player projection has $|\mathcal{X}_i| = 24$.

10. *Dynamic Oligopoly* (Ericson and Pakes, 1995): state $x = (q, m)$ pairs firm quality q with market structure m; the action shifts quality while m evolves under an action-invariant Markov process.

Table 1 reports the structural type, the predicted minimal FD horizon (from the primitives or the Type-S/K propositions above), the smallest ρ at which pointwise feasibility is certified by LSQR (`fd_holds_universal` in `code/fd_linear_system.py`), and the worst-case LSQR residual at the certifying horizon. The agreement between predicted and verified $\rho^\star$ across the eight tested models, together with residuals at machine precision, jointly confirms the rank test and the underlying value-difference identity.

Reading the table. Three patterns emerge. First, the four Type-R and Type-K models (rows 1–2, 4–5) achieve $\rho^\star = 1$ with structurally non-negative flow inputs and recover the classical AM non-negative weighting class. Second, the three Type-S models (rows 3, 7, 8) achieve $\rho^\star = p$ by the trivial shift-flush solution; Female Labor at $p = 1$, Altug–Miller at $p = 3$, and Education at $p = 4$ all match Proposition 10. Third, the three Type-D models (rows 6, 9, 10) display the signed-weight phenomenon: Entry/Exit with endogenous productivity requires $\rho^\star = 2$ with signed flow input; the entry/exit game and dynamic oligopoly are multi-player Markov-perfect models in which the equilibrium-induced strategic ripple acts at horizon two through opponents' best-response feedback, so the operative FD horizon is $\rho^\star = 2$ regardless of what a per-player projected feasibility check reports (Remark 6).

5.3 Three Walkthroughs

Engine Replacement (Type R, $\rho^\star = 1$). The state $x_t \in \{0, 1, \ldots, 9\}$ records mileage bins; actions $d_t \in \{0, 1\}$ are continue and replace. Under $d_t = 1$, mileage resets to bin 0

Model	Source	Type	S	D	$\rho^\star$ pred.	$\rho^\star$ verif.	LSQR resid.
Engine Replacement	Rust (1987)	R	10	2	1	1	8.7×10^{-14}
Job Search	—	R	10	2	1	1	1.3×10^{-16}
Female Labor (p=1)	Altug and Miller (1998)	S	2	2	1	1	3.1×10^{-16}
Investment	—	K	20	3	1	1	7.9×10^{-14}
Inventory Control	—	K	20	3	1	1	7.9×10^{-14}
Entry/Exit (Endo)	—	D	16	2	2	2	4.8×10^{-13}
Altug–Miller (p=3)	Altug and Miller (1998)	S	8	2	3	3	1.9×10^{-13}
Education (p=4)	Keane and Wolpin (1997)	S	16	2	4	—†	—
Entry/Exit Game	Ericson and Pakes (1995)	D	24	2	2	$2^\ddagger$	9.7×10^{-17}
Dynamic Oligopoly	Ericson and Pakes (1995)	D	16	2	2	$2^\ddagger$	1.9×10^{-13}

Table 1: Verification of the ten canonical models. Type codes: R = renewal; S = shift register; K = Kronecker with action-invariant exogenous factor; D = triple connector / strategic ripple. Predicted $\rho^\star$ comes from the structure (Proposition 10, Lemma 11, or direct verification at $\rho = 1$ followed by $\rho = 2$); verified $\rho^\star$ is the smallest horizon at which $\mathbf{A}^{(x_0,d,d')}\mathbf{x} = \mathbf{b}^{(x_0,d,d')}$ is feasible at every triple, computed by sparse LSQR. † Education ($p = 4$) is predicted by Proposition 10 at $\rho^\star = 4$; the pointwise LSQR test exceeds the per-model variable budget ($2 \cdot S^\rho D^\rho > 8000$) and is not run in the default configuration. ‡ The entry/exit game and the dynamic oligopoly are multi-player Markov-perfect models in which the equilibrium-induced strategic ripple acts at horizon two through opponents' best-response feedback; the operative horizon is therefore $\rho^\star = 2$. A per-player projected feasibility check can produce an apparent $\rho = 1$ certificate; we treat this as an artifact of the projection rather than a property of the equilibrium-induced game, and adopt $\rho = 2$ throughout (Remark 6). Reproducible from `python3 code/fd_linear_system.py`.

deterministically, so $\mathbf{F}_1$ has every column equal to $\mathbf{e}_0$ (rank one). Pointwise FD at $(x_0, 0, 1)$ at $\rho = 1$ uses

$$\Phi(x_0, 1) \;=\; \mathbf{e}_0\,\mathbf{e}_{(\cdot,d_1)}^\top \quad \text{(degenerate at } x_1 = 0), \qquad \Phi(x_0, 0) \;=\; \sum_{x'} f(x' \mid x_0, 0)\,\mathbf{e}_{x'}\,\mathbf{e}_{(\cdot,d_1)}^\top \cdot \omega_1(d_1 \mid x_0, x', 0),$$

where the per-step weight ω_1 on the $\Phi(x_0, 0)$ side is chosen by the QP to push the terminal distribution to $\mathbf{e}_0$ (matching $\Phi(x_0, 1)$). The solution is the AM (2011) renewal-action weight; LSQR residual is 8.7×10^{-14}. Both the direct Bellman value-difference and the FD-implied value-difference agree to the same precision in the code library (`example_engine_replacement.py`).

Altug–Miller (Type S, $p = 3$, $\rho^\star = 3$). The state is the labor-history register $(\ell_{t-3}, \ell_{t-2}, \ell_{t-1}) \in \{0, 1\}^3$; the transition deterministically shifts the register and appends $\ell_t = d_t$. After $\rho = 3$ periods of any future action sequence (d_1, d_2, d_3), the register becomes (d_1, d_2, d_3), independent of the initial action d_0. Pointwise feasibility at $\rho = 3$ is therefore trivial: both $\Phi(x_0, d)$ and $\Phi(x_0, d')$ assign full mass to the same future action sequence, with the deterministic shift register matching the terminal state automatically. At $\rho < 3$ the initial action remains in the register and pointwise FD fails; LSQR returns positive residuals at $\rho = 1, 2$ and machine zero (1.9×10^{-13}) at $\rho = 3$, matching Proposition 10 exactly (`example_altug_miller.py`).

Investment (Type K, $\rho^\star = 1$). The state factorizes as (capital $k_t \in \{1, \ldots, 5\}$, productivity $o_t \in \{1, \ldots, 4\}$), with three actions (depreciate, maintain, invest). The transition is $\mathbf{F}_d = \mathbf{F}_d^k \otimes \mathbf{F}^o$ where the capital block $\mathbf{F}_d^k$ is action-dependent (deterministic shift by -1, 0, $+1$ on the capital grid) and the productivity block $\mathbf{F}^o$ is action-invariant (AR(1) Tauchen discretization). By Lemma 11 the joint FD problem reduces to the capital-only FD problem; with three actions and a five-point grid, the capital-only problem has full controllability at $\rho = 1$ via the shift-by-± 1 structure, so $\rho^\star = 1$ for the joint state. LSQR confirms feasibility at $\rho = 1$ with residual 7.9×10^{-14}, despite the productivity factor having near-unit-root persistence (`example_investment.py`).

5.4 Pointwise vs. Universal FD: A Cautionary Observation

A practical caution about pointwise feasibility certificates in multi-player Markov-perfect games. Running the pointwise test of Theorem 3 at $\rho = 1$ on the *per-player projected perceived transition* $f_i^{\widehat{\mathbf{P}}_{-i}}$—obtained after integrating out opponents' first-stage CCPs—can return a feasibility certificate at every projected triple. We observe this for both the entry/exit game and the dynamic oligopoly in Table 1: the controllability subspace of the projected per-player problem at $\rho = 1$ is strictly smaller than $S - 1$ (3 of 23 for the game; 12 of 15 for the oligopoly), but the propagated baseline gaps of the projected problem still fall inside it.

This certificate is, however, an artifact of the projection rather than a property of the equilibrium-induced game. The strategic ripple that distinguishes a multi-player game from a single-agent decision problem operates at horizon two: the player's continuation-value difference depends on opponents' best-response feedback to the player's own action, which is itself a $\rho = 2$ object once the equilibrium fixed point is taken into account. The operative FD horizon for any multi-player Markov-perfect game in this class is therefore $\rho^\star = 2$, and we use $\rho = 2$ in the Monte Carlo of Section 6.2 regardless of the projected $\rho = 1$ certificate. For Type-D *single-agent* models with endogenous productivity (Entry/Exit, Endogenous), the pointwise $\rho <$ universal-rank gap is a real feature rather than a projection artifact, and the practical recommendation "run the pointwise test at the smallest ρ first" applies as stated.

5.5 Summary

The four-type taxonomy organizes the ten canonical models tested in code: two Type-R (renewal), three Type-S (shift register at lag depths $p = 1, 3, 4$), two Type-K (Kronecker with action-invariant exogenous factor), and three Type-D (triple connector and strategic ripple, all at $\rho^\star = 2$). For Types R, S, and K the minimal FD horizon $\rho^\star$ is read off the primitives without numerical computation; for Type D single-agent models the pointwise QP delivers $\rho^\star$ directly with possibly signed flow input, while for Type D multi-player games the operative horizon is $\rho^\star = 2$ regardless of what the per-player projected feasibility certificate reports (Remark 6). Every model is verified two ways in the accompanying code library—once by direct Bellman solve as a reference value-difference, and once by the LSQR pointwise rank test of Theorem 3—with agreement at machine precision across all configurations tested. The Monte Carlo evidence in the next section quantifies the finite-sample performance of the GFD estimator on representative members of each type.

6 Monte Carlo Evidence

This section evaluates the finite-sample performance of the GFD estimator on two of the canonical models from Section 5: the single-agent investment model with capital accumulation (Type K, $\rho^\star = 1$) and the three-player dynamic entry/exit game (Type D shared-action, $\rho^\star = 2$). For each DGP we compare the GFD estimator at the canonical (Moore–Penrose) input $\widehat{\mathbf{\Phi}}^{\text{MP}}$ against a nested fixed-point/pseudo-likelihood baseline. The Monte Carlo scripts are `code/monte_carlo_investment.py` and `code/monte_carlo_game.py`; results in this section are reproducible from those scripts.

6.1 Investment Model (Type K, $\rho^\star = 1$)

DGP. The state factorizes as (capital $k_t \in \{0, 1, 2, 3, 4\}$, productivity o_t with 4-state Tauchen discretization), giving $S = 20$. Three actions $d_t \in \{0, 1, 2\}$ correspond to depreciate, maintain, and invest, each shifting capital by -1, 0, or $+1$ deterministically; productivity evolves under an AR(1) Tauchen kernel that is action-invariant. By Lemma 11 the FD problem reduces to the capital component, with $\rho^\star = 1$ via the shift-by-± 1 structure. Discount $\beta = 0.95$. The structural parameter is $\theta_0 = (\theta^{\text{rev}}, \theta^{\text{cost}}, \theta^{\text{adj}})^\top = (2.5,\ 1.2,\ 0.8)^\top$, governing the linear-in-θ flow payoff $u(x_t, d_t; \theta) = \theta^{\text{rev}} \cdot o_t\sqrt{k_t} - \theta^{\text{cost}} d_t - \theta^{\text{adj}} d_t^2$.

Estimators. We compare three estimators:

- *GFD($\rho = 1$, MLE)*: the multinomial-logit MLE form (28) of Definition 3 at the canonical input $\widehat{\mathbf{\Phi}}^{\text{MP}}$. The implementation factors the augmented FD constraint matrix once (it is shared across initial states), batches all $|\mathcal{X}| \cdot (D-1)$ flow-QP right-hand sides into a single matrix multiplication, builds the cached effective regressors $(\widetilde{H}, \widetilde{h})$ once, and then runs a Newton iteration on the strictly concave logit pseudo-likelihood that converges in ~ 5 steps.

- *NFPM (Nested Fixed-Point Maximum Likelihood)*: solves the Bellman fixed point at every parameter evaluation; the canonical asymptotically efficient benchmark.

- *CCP-2step (Hotz–Miller)*: two-step estimator using the infinite-horizon Hotz–Miller representation $V = (I - \beta F_p)^{-1}\bar{u}$.

Results. Table 2 reports bias, RMSE, and mean estimation time per replication across 30 Monte Carlo replications at $N = 1000$, $T = 15$ (15,000 observations).

Method	$\theta^{\text{rev}} = 2.5$ Bias	RMSE	$\theta^{\text{cost}} = 1.2$ Bias	RMSE	$\theta^{\text{adj}} = 0.8$ Bias	RMSE	Time(s)
GFD(ρ=1)	0.069	0.293	−0.058	0.158	0.054	0.156	0.005
NFPM	−0.008	0.228	−0.005	0.088	0.006	0.076	5.13
CCP-2step	−0.006	0.229	−0.005	0.088	0.005	0.076	1.84

Table 2: Investment model Monte Carlo results, 30 replications, $N = 1000$, $T = 15$. Bias and RMSE in θ-units; Time is mean wall-clock seconds per replication. The GFD column reports the MLE form of Definition 3: a one-shot QP solve for the flow input $\widehat{\mathbf{\Phi}}^{\text{MP}}$ followed by Newton iteration on the multinomial-logit pseudo-likelihood (28) on cached effective regressors $(\widetilde{H}, \widetilde{h})$. Reproducible from `python3 code/monte_carlo_investment.py --quick` (which sets `MC_REPLICATIONS = 30`, `SAMPLE_SIZES = [(1000, 15)]`).

All three estimators are consistent: biases are within sampling error on every component. GFD-MLE's RMSE is roughly 1.3–1.8× larger than NFPM and CCP-2step across the three components, with the gap concentrated on θ^{cost} and θ^{adj}. This is the variance–input cost of GFD at the canonical (Moore–Penrose) flow input on a small-state calibration: the FD identity uses signed flow weights whose variance contribution is captured by the trace term in Theorem 8; tuning the null-space coordinate toward the variance-optimal flow narrows the gap (we revisit this trade-off in Section 6.3).

Timing. GFD completes a replication in 0.005 seconds—roughly 370× faster than CCP-2step and 1,000× faster than NFPM on this calibration. Step 1 (CCP and transition estimation) is shared across all three methods; Step 2 (a single batched flow-QP solve that returns the FD weights for all $|\mathcal{X}| \cdot (D-1)$ initial-state / action pairs at once) and Step 3 (cache effective regressors $(\widetilde{H}, \widetilde{h})$ via einsum) together cost a few milliseconds; Step 4 (Newton iteration on the strictly concave logit log-likelihood) converges in ~ 5 iterations. NFPM pays for the inner Bellman fixed point at every parameter trial; CCP-2step performs a single Hotz–Miller matrix inversion plus a non-linear search.

Scaling with state-space size. The structural advantage of GFD is that the FD weight solve $\widehat{\mathbf{\Phi}}$ is amortized over the parameter optimization: a single batched solve at the start of the replication yields effective regressors that are reused by every Newton step. NFXP, by contrast, re-solves the Bellman fixed point at every parameter trial. To document how the advantage evolves with the state-space size, we run the same investment DGP with capital and productivity grids tuned so that the joint state space ranges from $|\mathcal{X}| = 20$ to $|\mathcal{X}| = 2{,}000$, while holding all other DGP primitives fixed.

The advantage of GFD over NFXP *grows* with $|\mathcal{X}|$ once the action-invariant Kronecker structure of the transition is exploited. The empirical pattern: a stable $\sim$ 13–19× across $|\mathcal{X}| \in [20, 600]$, jumping to 112× at $|\mathcal{X}| = 2{,}000$ and 136× at $|\mathcal{X}| = 5{,}000$. The mechanism is asymmetric scaling at large $|\mathcal{X}|$: NFXP's Bellman matrix solve is $O(|\mathcal{X}|^2)$ per parameter

$\lvert\mathcal{X}\rvert$	GFD time (s)	NFXP time (s)	NFXP / GFD	GFD RMSE	NFXP RMSE
20	0.017	0.30	17×	0.68	0.56
60	0.025	0.32	13×	0.33	0.19
200	0.054	0.89	17×	0.32	0.18
600	0.11	2.10	19×	0.20	0.15
2,000	0.31	34.6	112×	0.05	0.12
5,000	1.39	189.2	136×	0.14	0.11

Table 3: Scaling of GFD vs. NFXP wall time as the joint state-space size grows. RMSE columns report mean RMSE across the three structural parameters. Same DGP as Table 2 (deterministic capital + Tauchen-discretized AR(1) productivity, three actions, $\rho = 1$, $N = 1000$, $T = 15$); $|\mathcal{X}|$ is varied by changing the capital grid size and the number of productivity nodes (5×4, 10×6, 20×10, 30×20, 50×40, 100×50). Replications: 5 per row except $|\mathcal{X}| = 5{,}000$, where each estimator is run for 3 replications due to NFXP wall-time cost. The GFD row uses (i) a smoothed first-stage CCP estimator (multinomial logit on the basis $(1, k, o, k\cdot o)$, $\sim$10 free parameters fitted by Newton iteration on the joint pseudo-likelihood; see Appendix B), in place of the empirical-frequency $\widehat{p}(a \mid s) = \text{counts}_{s,a}/n_s$ that has variance $O(|\mathcal{X}|/(NT))$; and (ii) the *structured* (Kronecker-decomposed) flow-QP solver of Appendix A, exploiting the action-invariant Kronecker factorisation $f(s' \mid s, a) = T_{\text{endo}}(k' \mid k, a)\cdot T_{\text{exo}}(o' \mid o)$ guaranteed by Lemma 11. The structured solver collapses the FD problem from the full state space ($|\mathcal{X}| = M\cdot Z$ states) to the endogenous component alone (M states), so each per-initial-state QP is $\sim Z^2$ times smaller and the number of distinct QPs solved drops from $|\mathcal{X}|\cdot(D-1)$ to $M\cdot(D-1)$. The full-state effective regressors $(\widetilde{H}, \widetilde{h})$ are then assembled in factored Kronecker form—never materialising the $(|\mathcal{X}|\cdot D)^2$ flow tensor—which keeps memory $O(M^2D^2)$ even at $|\mathcal{X}| = 5{,}000$. Reproducible from `python3 code/scalability_benchmark.py`.

trial $\times O(M)$ outer trials and inherits no benefit from the Kronecker structure (the Bellman fixed-point operator on the full state space is the same matrix-vector product whether or not f factors), while GFD's structured flow-QP scales as $O(M^2\cdot Z)$ for the cached regressors and $O(M^3D^3)$ for the per-state QPs (with M = endogenous-state size, Z = exogenous-state size, $|\mathcal{X}| = M\cdot Z$). Doubling Z at fixed M leaves the GFD QP cost unchanged but multiplies NFXP's Bellman work by 4.

The structural mechanism is Lemma 11 (action-invariant Kronecker cancellation): when one component of the state evolves independently of actions, it factors out of the FD problem entirely, and the whole flow-QP machinery operates on the endogenous component alone. The investment DGP has this structure (capital is endogenous, productivity is action-invariant Tauchen AR(1)), and so does any DDC where the choice variables interact only with a subset of state variables — a generic property of applied dynamic discrete-choice models with separable taste shocks, idiosyncratic demand shifters, or aggregate shocks.

Three further structural margins. Beyond the Kronecker exploit, GFD's relative advantage widens with:

(i) Multi-player Markov-perfect games. NFXP's inner loop becomes an iterated equilibrium

fixed point: at every candidate θ, a Gauss–Seidel best-response over K players runs ~ 100–200 per-player Bellman iterations to convergence. GFD computes one flow QP per player per *replication*, with no equilibrium re-solve at any candidate θ. The speedup ratio scales with the number of players, the slowness of MPE contraction, and the tightness of the outer convergence tolerance. The three-player entry/exit DGP of Section 6.2 delivers $\sim 4\times$ at $|\mathcal{X}| = 24$ per player; combining with Kronecker-exploited large $|\mathcal{X}|$ would multiply this further.

(ii) *High-dimensional structural parameter.* GFD-MLE has an analytic Hessian on cached effective regressors (Section 3.4); Newton converges in 5–10 steps regardless of $\dim\theta$. NFXP outer searches scale poorly with $\dim\theta$: Nelder–Mead at $O(d_\theta^2)$ trials, quasi-Newton at $O(d_\theta)$ but with each trial paying the full Bellman cost. As $\dim\theta$ grows, the GFD/NFXP gap widens independently of $|\mathcal{X}|$.

(iii) *Counterfactual computation.* The bridge theorem of Section 3.7 lets a single Φ solve be reused across many counterfactual scenarios, while NFXP re-solves the MPE for each one. The relative advantage scales with the number of counterfactuals.

The headline empirical fact: GFD wins on raw wall time across every state-space size we test, with a speedup that *grows super-linearly* from $|\mathcal{X}| = 600$ onward, reaching $\sim 140\times$ at $|\mathcal{X}| =$ 5,000. The accompanying RMSE column sits within a factor of $\sim 1.5\times$ of NFXP's at every row once the smoothed first-stage CCP of Appendix B is used.

The first-stage CCP estimator matters. Table 3's GFD RMSE is comparable to NFXP across the entire sweep, with the absolute level falling from 0.68 at $|\mathcal{X}| = 20$ to 0.05–0.14 at $|\mathcal{X}| \in \{2{,}000, 5{,}000\}$ (lower-rep counts at the largest cells contribute single-rep noise to the comparison). This stability is *not* automatic: it relies on the smoothed first-stage CCP estimator $\widehat{\boldsymbol{p}}^{\text{logit}}$ described in the table note. The naive empirical-frequency estimator $\widehat{p}^{\text{emp}}(a \mid x) = \text{counts}_{x,a}/n_x$ has pointwise variance $O(|\mathcal{X}|/(NT))$—with the panel size held fixed at $NT = 15{,}000$, the per-state observation count shrinks from 750 at $|\mathcal{X}| = 20$ to 7.5 at $|\mathcal{X}| = 2{,}000$, and this noise enters the GFD pseudo-likelihood linearly through the cached offset $\widetilde{h}_x$. The smoothed estimator collapses the per-state CCP variance to $O(p_{\text{dim}}/(NT))$ where $p_{\text{dim}} \approx 10$ is the number of basis coefficients in the multinomial logit, removing the $|\mathcal{X}|$ dependence almost entirely. NFXP avoids this issue by construction: it never plugs in a first-stage $\widehat{\boldsymbol{p}}$ and re-solves the Bellman fixed point at each candidate θ, so its variance is set by the Cramér–Rao bound on the full NT-observation likelihood.

Decomposing the residual gap: oracle CCPs. The smoothed-CCP GFD RMSE still sits roughly 1.5–1.7× above NFXP across the table; that residual gap is what the trace term in Theorem 8 predicts when the canonical MP flow input is used in place of the variance-optimal one. To verify this attribution, we re-run a parallel sweep with two alternative GFD first stages alongside NFXP:

$\lvert\mathcal{X}\rvert$	GFD (emp. $\widehat{p}$)	GFD (smooth. $\widehat{p}^{\text{logit}}$)	GFD (oracle $\boldsymbol{p}(\theta_0)$)	NFXP
20	0.76	0.68	0.68	0.56
60	0.37	0.33	0.22	0.19
200	0.51	0.32	0.22	0.18
600	0.60	0.20	0.148	0.146

Two readings: (i) smoothed CCPs absorb most of the empirical-CCP variance penalty (e.g. at $\lvert\mathcal{X}\rvert = 600$, the empirical column 0.60 falls to 0.20 under smoothing), and (ii) the remaining smoothed-vs-oracle gap then collapses fully under the oracle, with 0.148 matching NFXP's 0.146 to within 2%. So both the first-stage misspecification and the MP flow-norm *can* be closed; the practical levers are (a) a smoother / sieve / model-based first-stage CCP estimate, and (b) the variance-optimal flow input of Section 4.4, which targets the trace term directly. We revisit lever (b) on the game DGP in Tables 5 and 6; we have not yet ported the variance-optimal solver to the $\lvert\mathcal{X}\rvert = 2{,}000$ single-agent benchmark, and the smoothed-CCP version of Table 3 is the headline result for this section.

A complementary structural speedup arises in multi-player Markov-perfect games, where NFXP's inner loop becomes an iterated equilibrium fixed point rather than a single Bellman solve and the relative gap depends on the contraction rate of the MPE iteration; we document that case in Section 6.2. The single-agent and game speedups are not substitutes: the structured Kronecker FD solver of Appendix A (A.5) and the NFXP Gauss–Seidel inner loop are independent sources of GFD's relative advantage, and a model with both Kronecker structure and multi-player MPE would compound them.

6.2 Dynamic Entry/Exit Game (Type D, $\rho^\star = 2$)

DGP. A three-player dynamic entry/exit game with market sizes $\{2, 6, 10\}$ (3 levels, AR transition) and binary actions (exit/enter) per player. The joint state has $\lvert\mathcal{S}\rvert = 3 \cdot 2^3 = 24$ elements per player after the equilibrium-induced per-player projection (compute_ftran_game in `utils.py`). Discount $\beta = 0.9$; structural parameter for player 0 is $\theta_0 = (\theta_1, \theta_2, \theta_3, \theta_4)^\top = (1, 1, 1, 1)^\top$, with payoffs linear in market size, incumbent status, and competitor configuration. The minimal FD horizon is $\rho^\star = 2$: at $\rho = 1$ a per-player feasibility check using the projected perceived transition can appear to satisfy the pointwise test (cf. Remark 6), but this artifact reflects the projection rather than the equilibrium-induced strategic ripple, which acts at horizon two through the opponents' best-response feedback into the player's own continuation. We therefore use $\rho = 2$, the horizon required by any multi-player dynamic game in this class.

Estimators.

- *GFD($\rho = 2$)*: canonical input $\widehat{\mathbf{\Phi}}^{\mathrm{MP}}$, GMM moments from (21), applied per-player after Lemma 11 collapses opponents' equilibrium states. Following the same one-shot pattern as the single-agent case, the player-i flow input $\widehat{\mathbf{\Phi}}_i^{\mathrm{MP}}$ is solved *once per replication* on the perceived transition $\widehat{f}_i^{\widehat{\mathbf{p}}_{-i}}$ assembled from the first-stage cell-frequency CCP estimate $\widehat{\mathbf{p}}_{-i}$ rather than from the unobserved MPE profile $\mathbf{p}_{-i}^{\mathrm{MPE}}$. The resulting plug-in bias is asymptotically negligible at $\sqrt{N}$ rate provided the first-stage CCPs are $\sqrt{N}$-consistent (which they are under cell frequencies once each (x, d) cell is visited $O(N)$ times) and is left in the reported numbers without a further outer-loop correction.

- *NFXP (Nested Fixed-Point)*: at every candidate θ, re-solve the full Markov-perfect equilibrium via the Gauss–Seidel best-response iteration of `solve_game_equilibrium` (typically 100–200 iterations to convergence), then evaluate the per-player log-likelihood at the resulting equilibrium CCP profile. This is the natural Bellman-fixed-point benchmark: every parameter trial pays the cost of an MPE solve. The Aguirregabiria and Mira (2007) NPL alternative, which iterates between CCP updates and pseudo-likelihood maximization without solving the equilibrium fixed point exactly, is a faster but more fragile contemporaneous benchmark; we discuss it briefly at the end of this subsection.

Results. Table 4 reports bias, RMSE, and mean per-replication wall time across 100 Monte Carlo replications at $N = 2000$, $T = 20$ (40,000 observations) for player 0. Bias and RMSE are vectors over the four payoff components $(\theta_1, \theta_2, \theta_3, \theta_4)$.

	Bias				RMSE				Time(s)
Method	θ_1	θ_2	θ_3	θ_4	θ_1	θ_2	θ_3	θ_4	
GFD(ρ=2)	0.001	−0.008	0.012	−0.001	0.052	0.161	0.056	0.022	1.22
NFXP	−0.005	−0.017	0.004	−0.001	0.044	0.146	0.051	0.021	4.51

Table 4: Three-player entry/exit game Monte Carlo results, 100 replications, $N = 2000$, $T = 20$, player 0. True $\boldsymbol{\theta}_0 = (1, 1, 1, 1)^\top$. NFXP solves the full Markov-perfect equilibrium fixed point at every parameter trial; GFD solves the joint flow QP *once* per replication and then runs a Newton iteration on the multinomial-logit GFD pseudo-likelihood on cached effective regressors. Reproducible from `python3 code/monte_carlo_game.py --nfxp` (with internal defaults `n_reps = 100`, `sample_sizes = [(2000, 20)]`).

Both estimators are consistent: biases are within sampling error on every component. RMSE is essentially equal: mean RMSE 0.073 for GFD versus 0.066 for NFXP—an absolute gap of 0.007, of the same order as the Monte Carlo standard error at 100 replications. The third payoff component θ_3 has the largest dispersion under both methods (RMSE roughly 0.15–0.16), reflecting relatively weak identification of that component even at the 40,000-observation sample.

Timing: a structural advantage that scales with problem size. At this calibration ($S = 24$ per player, $D = 2$, $\rho = 2$), GFD is roughly 3.7× faster than NFXP (1.22s vs 4.51s per replication). The asymmetry is structural rather than incidental, and explains why the timing comparison flips between Table 2 and Table 4. GFD's cost is *front-loaded* into Step 2 of the procedure (Section 3.6): one joint feasibility QP per state, solved once per replication. After Step 2 the effective regressors $(\widetilde{H}, \widetilde{h})$ are cached and the parameter solve in Step 3 reduces to a low-dimensional MLE that touches no further FD machinery. NFXP, by contrast, must re-solve the full Markov-perfect equilibrium fixed point inside the inner parameter optimization, repeating the MPE iteration at every parameter trial.

The implication is that the relative cost of GFD vs. NFXP/NFPM is governed by the ratio

$$\frac{\text{GFD QP solve}}{\text{number of parameter trials} \times \text{MPE/Bellman solve}}.$$

On the small-state investment calibration of Table 2 ($S = 20$, single agent), GFD-MLE wins by 370× over CCP-2step and 1,000× over NFPM because the QP solve is small, the FD constraint matrix is shared across initial states (factored once and applied to all $S \cdot (D-1)$ right-hand sides simultaneously), and Steps 3–4 reduce to a single Newton iteration on cached regressors. On the larger-state, multi-player Game DGP of Table 4 (effective $S = 24$ per player, three players, $\rho = 2$), the inner MPE solve inside NFXP is genuinely expensive, and the front-loaded one-shot QP solve of GFD pays off 3.7×. The advantage widens further in regimes where NFXP's inner work grows faster than GFD's flow-QP solve does. Four margins are relevant: (i) the number of players in MPE games, where NFXP's inner Gauss–Seidel loop adds a per-player Bellman recursion that GFD avoids entirely; (ii) the dimension d_θ of the structural parameter, where NFXP's outer search scales as $O(d_\theta)$–$O(d_\theta^2)$ trials but GFD-MLE runs Newton in 5–10 steps regardless; (iii) the FD horizon ρ, where NFXP/NFPM backward induction depth grows but the GFD QP cost grows only polynomially in the per-state dimension; and (iv) the size of the state space $|\mathcal{X}| = M \cdot Z$ *when one component of the state is action-invariant*, since GFD's structured solver collapses the FD problem to the endogenous component M alone via the Kronecker cancellation of Lemma 11, while NFXP's Bellman matrix solve inherits no such reduction. Table 3 documents margin (iv) on the single-agent investment DGP, where the GFD/NFXP speedup grows from 13× at $|\mathcal{X}| = 60$ to 136× at $|\mathcal{X}| = 5{,}000$.

Note on NPL. The Aguirregabiria–Mira NPL alternative (Aguirregabiria and Mira, 2007) runs in ~ 0.13s on this calibration—faster than both GFD and NFXP—because the iterative-update form replaces the full MPE solve with a sequence of cheap CCP-update steps. NPL is the natural "upper bound" on what a fast contemporaneous CCP-based estimator can deliver on a small, well-behaved game DGP. We do not display NPL alongside NFXP in Table 4 because the comparison would mislead in a way the table cannot fix: NPL's per-iteration cost is not what applied researchers worry about; the operative concerns are multiplicity of fixed points (NPL can converge to a non-MPE solution), slow convergence in larger games, and the fact that NPL's runtime advantage on small DGPs vanishes as the per-iteration MPE update becomes expensive at scale. GFD avoids these issues *by construction*: the joint flow QP is convex with a unique minimum, the variance characterization of Theorem 8 gives a closed efficiency target, and the estimator never iterates over θ.

6.3 Variance Reduction by Null-Space Tuning (Theorem 6 demonstration)

The two preceding subsections use the canonical Moore–Penrose flow input $\widehat{\mathbf{\Phi}}^{\mathrm{MP}}$ for both DGPs. By Theorem 8, this is generically not the asymptotic variance-optimal input: tuning the null-space coordinate $\mathbf{q}$ in the descent direction of $\mathrm{tr}(\mathbf{\Sigma}(\mathbf{\Phi}(\mathbf{q})))$ at $\mathbf{q} = \mathbf{0}$ strictly reduces the asymptotic variance trace. We verify this empirically with two scripts.

Asymptotic trace reduction (`test_variance_optimal.py`). On a five-state subset of the game DGP ($x_0 \in \{3, 9, 12, 17, 21\}$), the script computes the canonical input $\widehat{\mathbf{\Phi}}^{\mathrm{MP}}(x_0)$ via the joint feasibility QP, the null-space basis $\mathbf{N}(x_0)$ (per-state $d_q = 2352$), the GFD information matrix

$$\mathbf{J}(\mathbf{\Phi}) \; = \; \sum_{x_0} w_{x_0}\, p_0(0 \mid x_0)\, p_0(1 \mid x_0)\, \widetilde{H}(x_0; \mathbf{\Phi})\, \widetilde{H}(x_0; \mathbf{\Phi})^{\top}$$

at the true θ_0, and the analytic gradient $\nabla_{\mathbf{q}} \mathrm{tr}(\mathbf{\Sigma}(\mathbf{\Phi}(\mathbf{q})))$ at $\mathbf{q} = \mathbf{0}$. A single descent step at $\eta = 10^{-3}$ reduces $\mathrm{tr}(\mathbf{\Sigma})$ from 3.80×10^3 (canonical) to 2.0×10^1 (a 99.5% reduction). The asymptotic trace can in principle be driven to zero by scaling $\mathbf{q}$ along the descent direction, because the FD-feasible variety $\mathcal{F}$ is unbounded; the compact regular sub-region $\mathcal{Q}$ of Theorem 8 restricts to a finite minimizer by absorbing the implicit bound on $\|\mathbf{q}\|$ that arises from second-order finite-sample corrections.

Finite-sample MC comparison (`monte_carlo_game_variance_opt.py`, `monte_carlo_game_variance_opt_mle` We run the same Monte Carlo as Section 6.2 ($N = 1000$, $T = 15$, 50 replications, $\rho = 2$) with two GFD inputs:

- *GFD-MP*: canonical Moore–Penrose $\widehat{\mathbf{\Phi}}^{\mathrm{MP}}$ at every x_0.
- *GFD-Opt*: $\widehat{\mathbf{\Phi}}^{\mathrm{MP}}$ updated by one step of the global analytic gradient descent on $\mathrm{tr}(\mathbf{\Sigma})$, computed across all 24 states of the per-player projection.

We sweep the descent step size η to map out the asymptotic-vs-finite trade-off, and we report results under *both* estimator forms of Definition 3: the squared-residual MD form (Table 5) and the multinomial-logit MLE form (Table 6).

Reading the tables: a finite-sample fragility, not an implementation artefact. Asymptotic trace reduction is real and substantial: at $\eta = 10^{-4}$, $\mathrm{tr}(\mathbf{\Sigma})$ falls by 96%. *Yet finite-sample RMSE does not improve under either estimator form.* Tables 5 and 6 show the same monotone pattern: as η grows the asymptotic trace falls but the empirical RMSE *rises*, with the MLE

η	$\Delta\,\mathrm{tr}(\mathbf{\Sigma})$	Bias θ_1	Bias θ_2	Bias θ_3	Bias θ_4	RMSE θ_1	RMSE θ_2	RMSE θ_3	RMSE θ_4
GFD-MP	(reference)	0.001	0.036	−0.017	−0.009	0.085	0.304	0.081	0.040
10^{-6}	$\sim 0\%$	−0.004	0.016	−0.013	−0.010	0.086	0.307	0.081	0.040
10^{-5}	69%	0.008	−0.022	0.045	−0.045	0.196	0.738	0.178	0.077
10^{-4}	96%	−0.063	−0.144	−0.017	−0.089	0.126	0.434	0.100	0.097

Table 5: GFD-Opt vs. GFD-MP at three descent step sizes η, *MD form* (squared-residual GMM). $\Delta\,\mathrm{tr}(\mathbf{\Sigma})$ is the asymptotic trace reduction relative to the canonical anchor. Bias and RMSE are over 50 Monte Carlo replications, $N = 1000$, $T = 15$, player 0 of the 3-player game. The $\eta = 10^{-5}$ row uses the smaller-sample configuration ($N = 500$, $T = 10$, 10 replications) and is included for trend visibility. Reproducible from `python3 code/monte_carlo_game_variance_opt.py --eta <value>`.

η	$\Delta\,\mathrm{tr}(\mathbf{\Sigma})$	Bias θ_1	Bias θ_2	Bias θ_3	Bias θ_4	RMSE θ_1	RMSE θ_2	RMSE θ_3	RMSE θ_4
GFD-MP-MLE	(reference)	−0.003	0.029	−0.024	−0.008	0.084	0.287	0.083	0.037
10^{-6}	$\sim 0\%$	−0.008	0.010	−0.021	−0.008	0.085	0.288	0.083	0.037
10^{-5}	69%	−0.032	−0.078	−0.007	−0.016	0.095	0.335	0.090	0.040
10^{-4}	96%	−0.065	−0.142	−0.027	−0.088	0.126	0.430	0.106	0.095

Table 6: GFD-Opt vs. GFD-MP at three descent step sizes η, *MLE form* (multinomial logit pseudo-likelihood, Newton-iterated). All rows use $N = 1000$, $T = 15$, 50 replications; the linear index in the binary logit is the GFD value difference $\widetilde{h}(x_0; \mathbf{\Phi})^\top \theta + \text{const}$. Reproducible from `python3 code/monte_carlo_game_variance_opt_mle.py --eta <value>`.

form behaving qualitatively the same as the MD form. Migrating to the concave logit pseudo-likelihood (which removes the squared-residual non-convexity, the actions $\neq d_{\mathrm{ref}}$ subsetting, and the Nelder–Mead inner solve, replacing it with one Newton solve in ~ 5 iterations) does *not* close the gap.

The diagnostic in `diagnose_phi_opt.py` identifies the mechanism. The variance-optimal direction in the FD null space pushes each per-state flow input $\mathbf{\Phi}^{\mathrm{opt}}(x_0)$ to a solution with much larger flow coefficients than the canonical anchor. At $\eta = 10^{-4}$ the per-state flow magnitudes increase by 6.5× in Frobenius norm. The GFD value-difference identity contains a Hotz–Miller correction that propagates the first-stage CCP estimate $\widehat{\mathbf{p}}$ into the linear index through the $\mathbf{\Phi}$-weighted sum of $\psi_d(\widehat{\mathbf{p}}(x))$ terms; with 6.5× larger flow weights, the resulting RMSE of this finite-sample correction term inflates by ≈20× relative to the $\widehat{\mathbf{\Phi}}^{\mathrm{MP}}$ baseline. This finite-sample CCP-noise channel is invisible to the asymptotic formula—which conditions on $\widehat{\mathbf{p}} = \mathbf{p}_0$ and collapses Theorem 8 to a one-step ahead information geometry—but it dominates the realized variance at $N = 1000$.

Sensitivity to sample size. Because the dominant channel through which $\mathbf{\Phi}^{\mathrm{opt}}$ damages finite-sample performance is the amplification of CCP estimation noise at rate $\|\mathbf{\Phi}^{\mathrm{opt}}\|/\|\mathbf{\Phi}^{\mathrm{MP}}\|$, that channel must shrink as N grows. We test this by re-running the MLE comparison at

$N \in \{10000, 50000\}$ at two step sizes (Table 7). Two patterns emerge clearly. *First*, with a small step ($\eta = 10^{-6}$), GFD-Opt and GFD-MP are essentially tied at $N = 10000$ and GFD-Opt strictly dominates GFD-MP at $N = 50000$ on every coefficient (RMSE reductions of 8%–20%, with bias on θ_1, θ_2 also halved). *Second*, the aggressive step $\eta = 10^{-4}$ never wins: even at $N = 50000$ the bias on θ_4 remains pinned near -0.08 and does not shrink between $N = 10000$ and $N = 50000$. This indicates that the descent step at $\eta = 10^{-4}$ is overshooting the local quadratic regime in which the gradient direction is informative, leaving a non-vanishing bias even at large N.

The combined message is a bias–variance trade-off in η: at any finite N there exists an optimal step $\eta^*(N)$ that is small enough to keep $\|\mathbf{\Phi}\|$ from inflating CCP-correction noise but large enough to extract some asymptotic gain. Empirically $\eta^*(N)$ shrinks more slowly than the $1/N$ rate suggested by a naive bias-variance balance, presumably because of the second-order overshoot identified above.

		Bias				RMSE			
N	method	θ_1	θ_2	θ_3	θ_4	θ_1	θ_2	θ_3	θ_4
10000	GFD-MP-MLE	0.013	0.053	−0.006	0.000	0.029	0.108	0.026	0.010
10000	GFD-Opt-MLE ($\eta = 10^{-6}$)	0.008	0.036	−0.003	−0.001	0.028	0.101	0.025	0.011
10000	GFD-Opt-MLE ($\eta = 10^{-4}$)	−0.044	−0.095	−0.013	−0.081	0.055	0.148	0.029	0.082
50000	GFD-MP-MLE	0.009	0.037	−0.003	0.000	0.014	0.052	0.009	0.006
50000	GFD-Opt-MLE ($\eta = 10^{-6}$)	0.005	0.021	0.000	−0.001	0.012	0.042	0.008	0.006
50000	GFD-Opt-MLE ($\eta = 10^{-5}$)	−0.015	−0.052	0.010	−0.009	0.018	0.063	0.013	0.011

Table 7: Sensitivity of the GFD-Opt vs. GFD-MP comparison to sample size, MLE form, $T = 15$, 50 replications. At $N = 10000$ the small step is essentially tied with the canonical anchor; at $N = 50000$ the small step strictly dominates on every coefficient (8%–20% RMSE reductions). The aggressive step $\eta = 10^{-4}$ remains dominated at every N tested; in particular the bias on θ_4 does not shrink between $N = 10000$ and $N = 50000$, consistent with a non-vanishing overshoot of the local linear regime in which the analytic gradient is informative. Reproducible from `python3 code/monte_carlo_game_variance_opt_mle.py --eta <value> --N <value> --T 15 --reps 50`.

What is established by the tables. Tables 5–7 together demonstrate empirically:

1. Theorem 8 holds: the canonical anchor is *not* asymptotically variance-optimal, and the asymptotic trace falls monotonically as one descends along the analytic gradient direction in the FD null space (three η values, 0% → 69% → 96% reduction).

2. For *large samples* ($N = 50000$ on the 24-state game DGP), the asymptotic gain does translate into a finite-sample RMSE improvement, provided the descent step η is small enough to prevent inflation of the $\mathbf{\Phi}$-weighted first-stage CCP correction. At $\eta = 10^{-6}$ on $N = 50000$, GFD-Opt strictly dominates GFD-MP on every coefficient.

3. For *small samples* ($N \leq 10000$ at any η, or any N at $\eta \geq 10^{-5}$), the unconstrained variance-optimal $\mathbf{\Phi}^{\text{opt}}$ is dominated by the canonical anchor: the CCP-noise inflation channel of the previous paragraph overwhelms the asymptotic information gain. A variance-stabilized $\mathbf{\Phi}^{\text{opt}}$ that penalizes flow magnitude explicitly—or, equivalently, a two-step procedure that updates the first-stage CCPs at the optimized flow before re-evaluating the moment—is the natural extension and the principled finite-sample correction; it is beyond the scope of this paper.

6.4 Discussion

Three observations from the Monte Carlo evidence.

Consistency. The GFD estimator (in both the MD and the game implementations) is consistent on both DGPs, with biases at the order of sampling error in the game and modest bias on the harder MD investment calibration. This confirms the population consistency discussed after Definition 3 for any FD-feasible input.

Implementation matters: MLE vs. MD. The investment and game MCs in Sections 6.1–6.2 implement the GFD estimator as a *minimum-distance* (MD) GMM on the squared FD value-difference moments. The variance-optimal MC of Section 6.3 additionally reports the multinomial-logit MLE form (28) of Definition 3, in which the GFD value difference serves as the linear index of a binary logit and the inner solve reduces to one Newton iteration on a strictly concave likelihood ($\sim$5 steps to convergence vs. Nelder–Mead). At the canonical anchor $\widehat{\mathbf{\Phi}}^{\text{MP}}$ the two forms deliver near-identical RMSE (Tables 5–6, top rows). The investment MC is roughly $2\times$ noisier than NFPM/CCP-2step (Table 2); migrating it to the MLE form is straightforward but does not affect the qualitative ordering. The variance-optimal selection delivers the asymptotic gain predicted by Theorem 8 but does not translate it to finite-sample RMSE under *either* estimator form, for the finite-sample CCP-noise reason discussed in Section 6.3.

Variance versus the baseline at the canonical input. The Moore–Penrose anchor is the simplest—but not the variance-optimal—input. The variance-optimal procedure of Section 4.4 tunes the null-space coordinate $\mathbf{q}$ to minimize trace variance, generically producing a strict RMSE improvement at the GFD-class trace optimum (Theorem 8). The Monte Carlo in this section uses the canonical input throughout; a variance-optimized GFD would compress the GFD row of both tables further, at the cost of one inner optimization in $\mathbf{q}$ per state.

Computational profile. GFD's structural advantage—solving the Bellman fixed point *zero* times during θ-optimization—becomes substantive only when (i) the state space is large and the Bellman fixed point is slow to compute, or (ii) the model has many parameters and the parameter-by-Bellman-iteration cost dominates total estimation time. For the small calibrations in this section ($S = 20$ and $S = 24$ per player), both NFPM/NPL and the MLE form of GFD complete a single replication in under two seconds, and constant-factor overhead in the GFD QP solve dominates the timing comparison. The GFD speedup over NFPM/NPL is therefore quantitatively understated by these calibrations and should be evaluated empirically on each application's actual state space.

7 Conclusion

This paper develops a computable framework for finite dependence in dynamic discrete choice models, organized around four results. Theorem 3 reduces existence of ρ-period finite dependence to a pointwise linear feasibility test $\mathbf{A}\mathbf{x} = \mathbf{b}$ on the structural transitions, giving practitioners a pre-estimation diagnostic at every (x_0, d, d') triple. Definition 3 packages the resulting GFD identity into an estimator that takes the FD-feasible flow input $\widehat{\boldsymbol{\Phi}}$ as an explicit argument, separating instrument selection from the rest of the procedure. Theorem 5 shows that the same identity yields a Bellman-free fixed-point system for payoff-only counterfactual CCPs, with restart-based validation against extraneous fixed points. Theorem 8 characterizes the asymptotically variance-optimal flow input as the trace-minimizing solution over a compact regular sub-region of the FD-feasible variety; the conditional extension of Section 4.5 relates this FD-class optimum to the ρ-horizon semiparametric efficiency bound under explicit gates.

Three limitations bound the present scope. First, the framework is developed for discrete state spaces with time-invariant transitions; extension to continuous or non-stationary states via sieve approximation remains open. Second, the bridge theorem is restricted to the payoff-only counterfactual class; transition-changing experiments require the broader compatibility machinery of Kalouptsidi et al. (2021). Third, the worst-case computational cost of the joint feasibility QP grows as $D \cdot S^{\rho} \cdot D^{\rho}$; the reachability pruning and iterative-solver techniques of Appendix A keep this manageable for the ten canonical models verified in Section 5.2, but pushing ρ above five in dense-transition models would require either a sparser reformulation or stronger structural restrictions. Each of these directions is a natural target for future work.

An invitation to applied work. The framework is designed to slot into existing structural DDC pipelines with minimal modification. Three concrete invitations follow. *First*, before committing to a nested fixed-point estimator, run the pointwise feasibility test of Theorem 3 on the estimated transitions; a positive verdict at $\rho \leq 2$—which holds for eight of the ten canonical models verified in Section 5.2, including ones with multi-component Kronecker-separable state spaces and shared-action multi-player games—makes GFD an immediate drop-in alternative with the same observable inputs and a substantial speed gain over Bellman-based benchmarks: ~ 370–$1{,}000\times$ faster than CCP-2step and NFPM on the small-state single-agent investment calibration of Table 2, $\sim 4\times$ faster than NFXP on the multi-player Markov-perfect entry/exit game of Table 4, and—when one component of the state is action-invariant, so that the structured Kronecker FD solver of Appendix A applies—up to $\sim 136\times$ faster than NFXP on the single-agent investment DGP at $|\mathcal{X}| = 5{,}000$, with a speedup that grows super-linearly in $|\mathcal{X}|$ (Table 3). *Second*, when conducting the payoff-only counterfactual policy analysis that dominates the applied DDC literature—changes in subsidies, fees, taxes, entry costs, or other monetary instruments—use the bridge of Theorem 5 to compute counterfactual conditional choice probabilities without re-solving the dynamic program at the counterfactual payoff, with the restart and spectral-radius diagnostics of Section 3.7 as the validation scaffold against extraneous fixed points. *Third*, when finite dependence holds with multiple feasible flow inputs—the generic case at $\rho \geq 2$—the variance-optimal selection of Theorem 8 costs one inner optimization in the null-space coordinate per initial state and recovers the FD-class trace minimum at no additional first-stage cost. A replication package accompanying this paper provides production-ready Python implementations of all three procedures together with the ten-canonical-model verification of Section 5.2 and the Monte Carlo benchmarks of Section 6.

Outlook. The combination of a checkable existence diagnostic, a Bellman-free estimator, a Bellman-free counterfactual computation, and a closed efficiency theorem within the FD class—packaged with reproducible code—is, to our knowledge, the first such bundle for the dynamic-discrete-choice setting. We expect it to broaden the empirical reach of structural dynamic work to settings in which the computational and functional-form burdens of full-solution methods have been a binding constraint on what applied researchers can credibly ask: large state spaces in industrial-organization entry/exit and dynamic oligopoly, multi-period labor histories with action-specific human capital, multi-component asset and demographic state vectors in household problems, and multi-player Markov-perfect games beyond the two- or three-player calibrations that have dominated the recent literature. The pre-estimation feasibility test of Theorem 3 makes the binary "does-finite-dependence-hold-here?" question a routine first step in applied work; we hope it will become as standard a diagnostic as a first-stage F-statistic in instrumental-variables practice.

A Scalable Computation for Large State Spaces

The joint feasibility QP (29) has decision variable $\operatorname{vec}\mathbf{\Phi}(x) \in \mathbb{R}^{D\cdot S^{\rho}\cdot D^{\rho}}$ per state x, and the associated history tree has $\sum_{\tau=0}^{\rho-1} S\cdot(D\cdot S)^{\tau}$ nodes. For $S = 20$, $D = 3$, and $\rho = 3$ this is 73,220 nodes per state and 219,660 flow variables: prohibitive in dense memory but tractable once two structural features of applied DDC models are exploited.

A.1 Reachability pruning

Many applied transitions are sparse or deterministic: capital shifts by at most one grid point under investment, labor histories evolve deterministically given the current choice, Tauchen-discretized AR(1) shocks have localized support. Define the *reachable successor set*

$$\mathcal{R}(d,x) \;:=\; \{x' \in \mathcal{X} : f(x' \mid x, d) > 0\}, \tag{38}$$

and let $b := \max_{d,x} |\mathcal{R}(d,x)|$ denote the maximum branching factor. For deterministic transitions $b = 1$; for localized stochastic transitions $b \ll S$. The pruned history tree replaces the full child enumeration at each step with $\{(\mathbf{h}, d, x') : x' \in \mathcal{R}(d, x_\tau)\}$; flow conservation, terminal-distribution matching, and the QP objective are all evaluated only over the pruned node set. The pruned tree has at most $\sum_{\tau=0}^{\rho-1} S \cdot (D \cdot b)^\tau$ nodes, a factor-$(S/b)^{\rho-1}$ reduction that grows fast with ρ. For the investment model ($S = 20$, $D = 3$, $\rho = 3$, with deterministic capital transitions and one Tauchen-stochastic productivity factor) the node count drops from 73,220 to 3,140, a 96% reduction.

The pruning is activated automatically once the estimated full tree size exceeds a threshold (we use 50,000 nodes); below threshold the solver builds the full tree, preserving exact behavior on small models.

A.2 Iterative KKT solving via LSQR

For small problems the QP (29) is solved by forming the KKT system

$$\begin{pmatrix} 2\,\mathbf{C}^\top\mathbf{C} + \varepsilon\,\mathbf{I} & \mathbf{A}_{\text{eq}}^\top \\ \mathbf{A}_{\text{eq}} & -\delta\,\mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{w} \\ \boldsymbol{\lambda} \end{pmatrix} = \begin{pmatrix} \mathbf{0} \\ \mathbf{b}_{\text{eq}} \end{pmatrix}, \tag{39}$$

where $\mathbf{C}$ maps flow variables to the terminal-distribution difference $\kappa_{\rho+1}(\cdot \mid x_0, d) - \kappa_{\rho+1}(\cdot \mid x_0, d')$, ε is a Tikhonov regularizer, and $\delta \approx 0$ ensures non-singularity. This requires forming $\mathbf{C}^\top\mathbf{C}$ explicitly, which is dimension $n_{\text{tot}} \times n_{\text{tot}}$ and becomes near-dense when $S \ll n_{\text{tot}}$.

For larger problems we replace the direct solve with a penalized least-squares formulation

$$\min_{\mathbf{w}} \left\| \begin{pmatrix} \sqrt{2}\,\mathbf{C} \\ \sqrt{\varepsilon}\,\mathbf{I} \\ \mu\,\mathbf{A}_{\text{eq}} \end{pmatrix} \mathbf{w} - \begin{pmatrix} \mathbf{0} \\ \mathbf{0} \\ \mu\,\mathbf{b}_{\text{eq}} \end{pmatrix} \right\|_2^2, \tag{40}$$

with $\mu \gg 1$ enforcing the equality constraints. Solving (40) via LSQR (Paige and Saunders, 1982) requires only matrix–vector products with the (sparse) stacked operator and never forms the normal equations, so memory usage is $O(\text{nnz})$ rather than $O(n_{\text{tot}}^2)$. The solver switches automatically between the direct KKT system (39) (used for $n_{\text{tot}} \le 5{,}000$) and the iterative formulation (40) otherwise.

A.3 Empirical performance

Table 8 reports the combined effect of pruning and iterative solving on the investment model ($S = 20$, $D = 3$, $\beta = 0.95$). At $\rho = 3$ the unpruned full tree exceeds available memory on a standard workstation; with pruning and LSQR the same problem solves in under one second to machine-precision matching accuracy.

ρ	Nodes (full)	Nodes (pruned)	Reduction	Solver	Time (s)
1	20	20	—	Direct KKT	0.01
2	1,220	1,220	—	LSQR	0.18
3	73,220	3,140	95.7%	LSQR	0.50

Table 8: Combined effect of reachability pruning and iterative LSQR solving on the investment model ($S = 20$, $D = 3$, $\beta = 0.95$). Without pruning, $\rho = 3$ runs out of memory; with pruning and LSQR it completes in 0.5 seconds with terminal-distribution matching error $\|\kappa_{\rho+1}(\cdot \mid x_0, d) - \kappa_{\rho+1}(\cdot \mid x_0, d')\|_1 < 10^{-15}$.

A.4 Batched flow-QP solve at $\rho = 1$: dense pinv vs. KKT-LU

The pruning-LSQR machinery of A.1–A.2 attacks *deep-tree* growth (the $b^{\rho-1}$ explosion in node count as ρ increases). At $\rho = 1$ the tree is one level deep and this machinery is moot; the cost is instead dominated by the *wide-state* dimension $|\mathcal{X}| \cdot D$ of the per-state flow vector, and the binding consideration is how to reuse a factorisation across the $|\mathcal{X}| \cdot (D-1)$ right-hand-side vectors $b^{(x_0,k)}$ that share a common constraint matrix A_{eq} (only the initial-flow block depends on x_0). Two paths are implemented in `code/scalability_benchmark.py`:

Dense SVD pseudoinverse (default for $|\mathcal{X}| \cdot D \leq 6{,}000$). Form the augmented matrix $\mathbf{A}_{\text{aug}} = [\sqrt{w}\,\mathbf{A}_{\text{eq}}; \mathbf{C}]$ explicitly, compute its Moore–Penrose pseudoinverse via LAPACK's `dgesvd`, and apply it to the $|\mathcal{X}| \cdot (D-1)$ right-hand sides as a single matrix multiplication. LAPACK's BLAS implementation makes this unbeatable up to a few thousand columns on modern hardware; cost is $O((|\mathcal{X}| \cdot D)^3)$ for the SVD, which dominates as $|\mathcal{X}|$ grows.

Sparse symmetric-indefinite KKT factorisation (default for $|\mathcal{X}| \cdot D > 6{,}000$). The constrained QP $\min_{\mathbf{x}} \|\mathbf{C}\mathbf{x}\|^2$ s.t. $\mathbf{A}_{\text{eq}}\mathbf{x} = \mathbf{b}$ has KKT system

$$\begin{pmatrix} 2\,\mathbf{C}^\top\mathbf{C} + \varepsilon\mathbf{I} & \mathbf{A}_{\text{eq}}^\top \\ \mathbf{A}_{\text{eq}} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \boldsymbol{\lambda} \end{pmatrix} = \begin{pmatrix} \mathbf{0} \\ \mathbf{b} \end{pmatrix}. \tag{41}$$

Both $\mathbf{C}$ and $\mathbf{A}_{\text{eq}}$ are sparse for any sparse-transition DGP (Tauchen-discretised AR(1), deterministic shifts, Kronecker-separable factor models), so $\mathbf{C}^\top\mathbf{C} + \varepsilon\mathbf{I}$ remains sparse and the KKT block is symmetric-indefinite-but-sparse. We factor it once via `scipy.sparse.linalg.splu` (using the default UMFPACK backend) and then triangular-solve for each right-hand side $(\mathbf{0}; \mathbf{b}^{(x_0,k)})$. Empirical fill-in for the investment DGP is moderate, with factorisation cost that scales sub-cubically in $|\mathcal{X}|$ and triangular-solve cost that is essentially linear in $|\mathcal{X}|$ per RHS; end-to-end the path scales considerably better than the dense pseudoinverse for general (non-Kronecker-separable) DGPs. For DGPs that admit the Kronecker factorisation (42) of A.5 below—the case in Table 3—the structured solver dominates this KKT-LU path by another factor of $\sim Z^c$.

Why not normal equations. The standard alternative, $\mathbf{x} = (\mathbf{A}_{\text{aug}}^\top \mathbf{A}_{\text{aug}})^{-1}\mathbf{A}_{\text{aug}}^\top \mathbf{b}$ via sparse splu, squares the condition number of $\mathbf{A}_{\text{aug}}$. With the penalty weight $w = 10^6$ already needed to enforce the constraints to acceptable precision, $\text{cond}(\mathbf{A}_{\text{aug}}) \sim 10^6$ becomes $\text{cond}(\mathbf{A}_{\text{aug}}^\top \mathbf{A}_{\text{aug}}) \sim 10^{12}$, which a small Tikhonov ridge cannot regularise without introducing significant bias; in our trials this path produced estimates that diverged numerically at $|\mathcal{X}| = 200$. The KKT formulation (41) avoids the squaring entirely and is the recommended sparse path.

A.5 Structured Kronecker solver: when the state factors as endo $\times$ exo

The dominant scaling exploit, used to produce Table 3, is to recognise that many applied DDC transitions factor as a Kronecker product

$$f(s' \mid s, a) \;=\; T_{\text{endo}}(k' \mid k, a) \,\cdot\, T_{\text{exo}}(o' \mid o), \tag{42}$$

with $s = (k, o)$, where k is an endogenous state (responding to actions) and o is an action-invariant exogenous state. The investment model has $k =$ capital and $o =$ Tauchen-discretised productivity; the labour-supply model has $k =$ labour history and $o =$ wage shock; multi-agent games have player-i state as k and opponent-state-and-environment as o at the equilibrium-induced perceived transition. Lemma 11 shows that under (42), the FD problem on the full state space reduces to the FD problem on the endogenous component k alone: the exogenous factor T_{exo} cancels identically between the two compared actions. Concretely, the optimal flow weight at the full state factors as

$$\mathbf{\Phi}^{\text{full}}\big((k, o),\, a;\, (k', o')\big) \;=\; \mathbf{\Phi}^{\text{endo}}(k,\, a;\, k') \,\cdot\, T_{\text{exo}}(o' \mid o), \tag{43}$$

so the structured solver computes $\mathbf{\Phi}^{\text{endo}}$ on the M-state endogenous problem ($M := |\{k\}|$) and never materialises the $|\mathcal{X}|^2$ full-state weight tensor.

Effective regressors via factored Kronecker form. With $\mathbf{\Phi}^{\text{endo}}$ in hand, the GFD-MLE effective regressor $\widetilde{H}[s, a, r]$ at full state $s = (k, o)$ decomposes by basis component r via the conditional expectation operator $E_{o'|o}[\cdot] := \sum_{o'} T_{\text{exo}}(o' \mid o)\,\cdot$. For the investment basis

$z[(k,o),a,\cdot] = (o\cdot k,\ -a,\ -a^2)$:

$$\widetilde{H}[s,a,0] = \beta\, E_{o'|o}[o'] \cdot \sum_{k',a'} \Delta\mathbf{\Phi}^{\text{endo}}(k,a;k',a')\cdot k' \qquad \text{(revenue, action-invariant base term)}$$
$$\widetilde{H}[s,a,1] = -(a-a_{\text{ref}}) \ - \ \beta \sum_{k',a'} \Delta\mathbf{\Phi}^{\text{endo}}(k,a;k',a')\cdot a' \qquad \text{(cost)}$$
$$\widetilde{H}[s,a,2] = -(a^2-a_{\text{ref}}^2) \ - \ \beta \sum_{k',a'} \Delta\mathbf{\Phi}^{\text{endo}}(k,a;k',a')\cdot a'^2 \qquad \text{(adjustment)}$$

The Hotz–Miller offset $\widetilde{h}$ is similarly factored:

$$\widetilde{h}[s,a] \ = \ \beta \sum_{k',a'} \Delta\mathbf{\Phi}^{\text{endo}}(k,a;k',a') \ \cdot \ E_{o'|o}\big[\widehat{e}(a',(k',o'))\big],$$

where $\widehat{e}(a,s) = -\log\widehat{p}(a \mid s) + \gamma_E$ is the Hotz–Miller correction. All conditional expectations $E_{o'|o}[\cdot]$ are computed once as $Z\times Z$ matrix-vector products against T_{exo}, never as full-state sums.

Cost reduction. At fixed D and fixed ρ, the structured solver does $M\cdot(D-1)$ flow-QP solves of size $M\cdot D$, instead of $|\mathcal{X}|\cdot(D-1) = M\cdot Z\cdot(D-1)$ solves of size $|\mathcal{X}|\cdot D = M\cdot Z\cdot D$ in the unstructured path. The work ratio is

$$\frac{\text{unstructured cost}}{\text{structured cost}} \ = \ Z \ \cdot \ \left(\frac{|\mathcal{X}|}{M}\right)^c \ = \ Z^{1+c},$$

where c is the per-QP-size scaling exponent ($c=3$ for dense pseudoinverse, $c\approx 1.5$–2 for sparse KKT-LU). On the investment DGP at $|\mathcal{X}| = 5{,}000 = 100\cdot 50$, this predicts a $\sim Z^4 = 6{,}250{,}000\times$ reduction relative to dense pinv on the full state and $\sim Z^3 = 125{,}000\times$ reduction relative to KKT-LU on the full state. The empirical reduction is more modest because of constant-factor overhead (matrix construction, basis precomputation), but Table 3 shows that the structured solver brings GFD wall time at $|\mathcal{X}| = 5{,}000$ from 111s (KKT-LU on full state) down to 1.39s, an $80\times$ reduction in this single benchmark.

B Smoothed First-Stage CCP Estimator

The two-step GFD estimator of Section 3.4 requires a first-stage estimate $\widehat{\boldsymbol{p}}$ of the conditional choice probabilities $p_0(d \mid x)$, which enters the value-difference identity through the Hotz–Miller correction $\widehat{\psi}_d(x) = -\log\widehat{p}(d \mid x) + \gamma_E$. Any consistent estimator suffices for $\sqrt{N}$-consistency of $\widehat{\theta}$, but the *finite-sample* variance of $\widehat{\theta}$ inherits the variance of $\widehat{\boldsymbol{p}}$ entry-by-entry through the cached offset $\widetilde{h}_x$ of (28). With the empirical-frequency estimator

$$\widehat{p}^{\text{emp}}(d \mid x) \ = \ \frac{n_{x,d}}{n_x}, \qquad n_{x,d} := \sum_{i,t} \mathbf{1}\{x_{it}=x,\ d_{it}=d\}, \tag{44}$$

the per-state variance is $\mathrm{Var}[\widehat{p}^{\mathrm{emp}}(d \mid x)] = p_0(d \mid x)(1-p_0(d \mid x))/n_x$, and the average occupancy count $n_x = NT/|\mathcal{X}|$ shrinks with the state space. Concretely, with $NT = 15{,}000$ observations the average n_x falls from 750 at $|\mathcal{X}| = 20$ to 7.5 at $|\mathcal{X}| = 2{,}000$, and the empirical-CCP first stage drives the GFD-MP RMSE divergence reported in the central column of the diagnostic panel of Section 6.1.

B.1 Multinomial-logit smoother on a low-dimensional basis

We replace (44) with a parametric multinomial logit fit on a low-dimensional basis $\mathbf{z}(x) \in \mathbb{R}^{p_{\mathrm{dim}}}$ of the state. For the investment DGP of Section 6.1 we use

$$\mathbf{z}(x) \;=\; \big(1,\, k(x),\, o(x),\, k(x) \cdot o(x)\big)^\top, \qquad p_{\mathrm{dim}} = 4, \tag{45}$$

where $k(x)$ is the capital index and $o(x)$ the productivity grid value. The smoother fits coefficients $\boldsymbol{\gamma}_a \in \mathbb{R}^{p_{\mathrm{dim}}}$ for each non-reference action $a \in \{1, \ldots, D-1\}$ (with action 0 as the reference) by maximising the multinomial-logit log-likelihood

$$\ell(\boldsymbol{\gamma}) \;=\; \sum_{i,t} \log \frac{\exp\!\big(\mathbf{z}(x_{it})^\top \boldsymbol{\gamma}_{d_{it}}\big)}{\sum_{a'=0}^{D-1} \exp\!\big(\mathbf{z}(x_{it})^\top \boldsymbol{\gamma}_{a'}\big)}, \qquad \boldsymbol{\gamma}_0 \equiv \mathbf{0}, \tag{46}$$

by Newton iteration on the joint Hessian. Because $\ell(\boldsymbol{\gamma})$ is strictly concave, Newton converges in $\sim$5 iterations from any reasonable start. The fitted CCPs are then evaluated at every state of the model:

$$\widehat{p}^{\mathrm{logit}}(d \mid x) \;=\; \frac{\exp\!\big(\mathbf{z}(x)^\top \widehat{\boldsymbol{\gamma}}_d\big)}{\sum_{a'=0}^{D-1} \exp\!\big(\mathbf{z}(x)^\top \widehat{\boldsymbol{\gamma}}_{a'}\big)}, \qquad x \in \mathcal{X}. \tag{47}$$

Implementation in `code/scalability_benchmark.py` (function `_smoothed_ccp_logit`); the entire smoother runs in $O(NT \cdot p_{\mathrm{dim}}^2)$ time per Newton iteration, dominated by the assembly of the $p_{\mathrm{dim}} \cdot (D-1) \times p_{\mathrm{dim}} \cdot (D-1)$ Hessian.

B.2 Variance properties

The pointwise variance of $\widehat{p}^{\mathrm{logit}}(d \mid x)$ scales as $p_{\mathrm{dim}}/(NT)$ rather than as $|\mathcal{X}|/(NT)$, because all NT observations contribute to the estimation of every $\widehat{\boldsymbol{\gamma}}_d$. For the investment calibration of Section 6.1, this gives a per-state variance reduction factor of approximately $|\mathcal{X}|/p_{\mathrm{dim}} = 500$ at $|\mathcal{X}| = 2{,}000$, which is what drives the RMSE collapse from 1.08 (empirical) to 0.19 (smoothed) reported in Table 3.

B.3 When parametric smoothing is appropriate

The logit smoother is consistent for the true CCP function whenever $p_0(d \mid x) = \Lambda_d(\mathbf{z}(x)^\top \boldsymbol{\gamma}_d)$ for some $\boldsymbol{\gamma}$, i.e. when the basis $\mathbf{z}(\cdot)$ is rich enough to span the log-odds surface $\log[p_0(d \mid x)/p_0(0 \mid x)]$.

In our DGPs this is satisfied exactly because the equilibrium CCP is itself a multinomial logit in the basis $\mathbf{z}$ (the structural payoff is linear in the same $(k, o, k \cdot o)$ regressors). When this exact-parametric assumption fails, three robust alternatives preserve the variance gain:

1. *Sieve logit.* Enrich $\mathbf{z}(x)$ with polynomial, spline, or tensor-product basis terms; choose the dimension by cross-validation. Asymptotic results in Newey and McFadden (1994) apply provided the sieve dimension grows at the standard nonparametric rate.

2. *Kernel smoothing.* Replace (44) by a Nadaraya–Watson estimator over a metric on $\mathcal{X}$: $\widehat{p}^{\text{kernel}}(d \mid x) = \sum_{i,t} K_h(x - x_{it})\mathbf{1}\{d_{it} = d\} / \sum_{i,t} K_h(x - x_{it})$, with bandwidth h chosen by cross-validation. This is the original Hotz and Miller (1993) recommendation.

3. *Frequency estimator with Laplace smoothing.* $\widehat{p}^{\text{lap}}(d \mid x) = (n_{x,d} + \alpha)/(n_x + \alpha D)$ with $\alpha \in (0, 1]$. Cheap; partly mitigates the zero-cell problem but does *not* reduce the $|\mathcal{X}|/(NT)$ variance scaling.

The choice between (1)–(3) is the standard CCP first-stage selection problem in the two-step DDC literature; the GFD second stage is invariant to this choice in large-sample efficiency, but finite-sample RMSE is highly sensitive when $|\mathcal{X}|$ grows relative to NT.

References

Abadie, A., A. Diamond, and J. Hainmueller (2019). Synthetic control methods with signed donor weights and a recursive computation. *Working paper*.

Adusumilli, K. and D. Eckardt (2025). Temporal-difference estimation of dynamic discrete choice models. *Review of Economic Studies*. Forthcoming.

Aguirregabiria, V. and M. Marcoux (2021). Imposing equilibrium restrictions in the estimation of dynamic discrete games. *Quantitative Economics 12*(4), 1223–1271.

Aguirregabiria, V. and P. Mira (2007). Sequential estimation of dynamic discrete games. *Econometrica 75*(1), 1–53.

Altug, S. and R. A. Miller (1998). The effect of work experience on female wages and labour supply. *Review of Economic Studies 65*(1), 45–85.

Arcidiacono, P. and R. A. Miller (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica 79*(6), 1823–1867.

Arcidiacono, P. and R. A. Miller (2019). Nonstationary dynamic models with finite dependence. *Quantitative Economics 10*(3), 853–890.

Arcidiacono, P. and R. A. Miller (2020). Identifying dynamic discrete choice models off short panels. *Journal of Econometrics 215*(1), 48–68.

Bajari, P., C. L. Benkard, and J. Levin (2007). Estimating dynamic models of imperfect competition. *Econometrica 75*(5), 1331–1370.

Berry, S. T. and G. Compiani (2023). An instrumental variable approach to dynamic models. *Review of Economic Studies 90*(4), 1724–1758.

Blevins, J. R. (2025). Identification and estimation of continuous-time dynamic discrete choice games. *Quantitative Economics*. Forthcoming.

Borusyak, K. and P. Hull (2024). Negative weights are no concern in design-based specifications. *AEA Papers and Proceedings 114*, 597–600.

Bunting, J. and T. Ura (2025). Faster estimation of dynamic discrete choice models using index invertibility. *Journal of Econometrics 250*, 106004.

Chamberlain, G. (1987). Asymptotic efficiency in estimation with conditional moment restrictions. *Journal of Econometrics 34*(3), 305–334.

Chen, E. (2025). A model-adaptive approach to the estimation of dynamic discrete choice models with large state spaces. *arXiv:2501.18746*.

de Chaisemartin, C. and X. D'Haultfœuille (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review 110*(9), 2964–2996.

Dix-Carneiro, R. (2014). Trade liberalization and labor market dynamics. *Econometrica 82*(3), 825–885.

Ericson, R. and A. Pakes (1995). Markov-perfect industry dynamics: A framework for empirical work. *Review of Economic Studies 62*(1), 53–82.

Gayle, W.-R. (2021). CCP estimation of dynamic discrete/continuous choice models with generalized finite dependence. Working Paper.

Hao, Y. and H. Kasahara (2024). Conditional choice probability estimation of dynamic discrete choice models with 2-period finite dependence. arXiv Working Paper.

Hotz, V. J. and R. A. Miller (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies 60*(3), 497–529.

Kalouptsidi, M., P. T. Scott, and E. Souza-Rodrigues (2021). Identification of counterfactuals in dynamic discrete choice models. *Quantitative Economics 12*(2), 351–403.

Kasahara, H. and K. Shimotsu (2009). Nonparametric identification of finite mixture models of dynamic discrete choices. *Econometrica 77*(1), 135–175.

Keane, M. P. and K. I. Wolpin (1997). The career decisions of young men. *Journal of Political Economy 105*(3), 473–522.

Liberzon, D. (2003). *Switching in Systems and Control*. Boston: Birkhäuser.

Magnac, T. and D. Thesmar (2002). Identifying dynamic discrete decision processes. *Econometrica 70*(2), 801–816.

Murphy, K. M. and R. H. Topel (1985). Estimation and inference in two-step econometric models. *Journal of Business & Economic Statistics 3*(4), 370–379.

Newey, W. K. (1990). Efficient instrumental variables estimation of nonlinear models. *Econometrica 58*(4), 809–837.

Newey, W. K. and D. McFadden (1994). Large sample estimation and hypothesis testing. In R. F. Engle and D. McFadden (Eds.), *Handbook of Econometrics*, Volume 4, pp. 2111–2245. Elsevier.

Paige, C. C. and M. A. Saunders (1982). LSQR: An algorithm for sparse linear equations and sparse least squares. *ACM Transactions on Mathematical Software 8*(1), 43–71.

Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica 55*(5), 999–1033.

Ryan, S. P. (2012). The costs of environmental regulation in a concentrated industry. *Econometrica 80*(3), 1019–1061.

Sun, Z. and S. S. Ge (2005). *Switched Linear Systems: Control and Design*. London: Springer.

Tauchen, G. (1986). Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters 20*(2), 177–181.