A Generalized Finite Dependence Framework for Dynamic Discrete Choice Models

Yu Hao*
Faculty of Business and Economics
The University of Hong Kong
haoyu@hku.hk

Hiroyuki Kasahara
Vancouver School of Economics
University of British Columbia
hkasahar@mail.ubc.ca

Katsumi Shimotsu
Faculty of Economics
The University of Tokyo
shimotsu@e.u-tokyo.ac.jp

June 7, 2026

Abstract

Finite dependence (FD) makes dynamic discrete choice models estimable without solving the Bellman equation, but establishing it has relied on case-by-case construction, fragile in games where feasibility hinges on equilibrium play. We turn existence into a single checkable test: pointwise ρ-step FD holds at every state-action triple if and only if $\mathbf{b} \in \text{Range}(\mathbf{A})$, a sparse linear-feasibility condition on the kernel alone. Searching over ρ returns the first certified horizon ρ^* (Theorem 3). Signed, history-conditional weights strictly enlarge the feasible set, reaching strategic kernels beyond prior non-negative weights. The generalized FD (GFD) estimator attains the minimum asymptotic-variance trace within the FD-feasible class (Theorem 6) and extends the same representation to payoff-only counterfactual conditional choice probabilities (CCPs) in non-stationary settings (Theorem 8). The replication package verifies the test on eight dynamic discrete choice (DDC) examples. GFD estimates a non-stationary ten-carrier U.S. airline entry game (2011–2023) with one regression per carrier. It recovers dollar structure comparable to a structural benchmark and rejects carrier-year entry-value stationarity for eight of ten carriers.

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

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

*We are grateful to Robert Miller, Victor Aguirregabiria, and participants at the DSE-HKU 2025 conference for valuable comments on this paper. Yu Hao also thanks his UBC PhD committee (Vadim Marmer, Kevin Song, and Florian Hoffmann) for their guidance during the development of the ideas in this paper. Hao gratefully acknowledges financial support from the Hong Kong Research Grants Council Early Career Scheme (ECS, Project No. 27603224, 2024 Exercise) under the project "Estimating Dynamic Discrete-Continuous Choice Models using a Generalized Euler Equation Approach".

1 Introduction

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

Finite dependence (FD) circumvents the Bellman fixed point. If two current actions induce future state distributions that match after a short horizon ρ, the continuation values cancel from the value-function difference. The difference then reduces to a finite sum of flow utilities and conditional choice probabilities (CCPs; Hotz and Miller, 1993). Arcidiacono and Miller (2011) introduced finite dependence. Arcidiacono and Miller (2019) extended it to non-stationary models and dynamic games. Gayle (2021) generalized it to discrete–continuous choice. An early instance of the cancellation appears in Altug and Miller (1998). In each case, the analyst constructs weights for an amenable structure and then verifies the minimal horizon case by case. Games make this step more onerous because feasibility hinges on equilibrium play. No general characterization tells us, for an arbitrary transition kernel, whether ρ-step FD admits a feasible representation.

We close this gap. The main result (Theorem 3) reduces pointwise ρ-step FD at every triple (x_t, d_t, d'_t) to a single sparse linear-feasibility condition $\mathbf{b} \in \text{Range}(\mathbf{A})$. The condition is constructed from the transition kernel f alone. It does not use the value function or CCPs. The condition is necessary and sufficient, and analysts can run it before any estimation step. It rests on allowing signed, history-conditional weights, which strictly enlarges the set of kernels at which FD holds. We recover three recurring transition classes in closed form: renewal, shift-register, and Kronecker with an action-invariant exogenous factor. The rank test reaches the fourth class, the strategic "triple-connector" class of dynamic games, at $\rho^* = 2$. For this class, the non-negative contemporaneous-state weights of prior FD provably fail. We verify the existence test on eight benchmark and stylized DDC examples with numerical residuals reported in Section 4.1. After the test establishes existence, the generalized finite dependence (GFD) estimator of Section 3 takes any FD-feasible flow input as an explicit argument. Within the FD-feasible class it attains the minimum trace of the asymptotic variance over a compact regular sub-region (Theorem 6). For payoff-only counterfactuals, the same FD representation gives a fixed-point system for counterfactual CCPs in non-stationary models.

The efficiency claim is deliberately scoped. GFD is a class of CCP estimators. It does not generally dominate a full-solution maximum-likelihood estimator (MLE) that exploits all Bellman restrictions or Markov-perfect-equilibrium restrictions. Instead, GFD chooses the variance-minimizing flow input among FD-feasible estimators that use a ρ-period window. Under explicit gates, this input attains the ρ-horizon semiparametric efficiency bound (Corollary 7). When those gates fail, a full-solution estimator can use restrictions outside the GFD class. In the strategic game simulation below, the trace of the GFD covariance is 1.24 times the covariance of the nested fixed point (NFXP) estimator. We therefore do not claim that GFD is generally more efficient than full-solution methods. We show that GFD makes the finite-horizon information set operational and measures the cost when the full-information bound is out of reach.

A key advantage of GFD is its natural accommodation of non-stationary environments. In many labor economics applications the environment is inherently non-stationary. Life-cycle models span several domains. Examples include occupational choice (James, 2014), education (Arcidiacono et al., 2016), and female labor supply and fertility (Altug and Miller, 1998; Gayle and Golan, 2012; Gayle et al., 2018). These settings feature finite horizons, aging, and aggregate shocks. Traditional full-solution methods cannot rely on a stationary Bellman contraction in these settings. Researchers must often make strong assumptions about the terminal value function or the evolution of states beyond the sample period. Finite dependence instead matches state distributions over a short horizon ρ. It bypasses the continuation value V over that horizon. As emphasized by Arcidiacono and Miller (2019), this allows estimation of non-stationary models without imposing a terminal value-function path in settings where finite dependence holds.

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

We make four contributions. First, we give an existence and rank-test characterization. The test rests on a constraint–decomposition equivalence (Theorem 2) and recovers the four FD-amenable transition classes of Section 4.1. Second, we define GFD as a class of two-step CCP estimators. The class is consistent for any FD-feasible input (Theorem 5). Third, we prove FD-class trace optimality at $\widehat{\mathbf{\Phi}}^{\text{opt}}$ (Theorem 6) and state when this optimum attains the ρ-horizon semiparametric efficiency bound. Fourth, we record the payoff-only counterfactual implication (Theorem 8). Its role here is not a new counterfactual method. It shows that the FD representation extends the range of non-stationary models in which payoff-only counterfactual CCPs can be computed without a value-function path. This is the identity used in the non-stationary airline application. Appendix B collects the numerical infrastructure: reachability pruning, iterative least-squares QR (LSQR), Karush–Kuhn–Tucker (KKT)-LU batched solving, and the structured Kronecker solver. Appendix E gives a short latent-type extension.

The empirical application illustrates the non-stationary payoff channel without imposing a terminal value-function path. We estimate a dynamic entry game of ten U.S. carriers over 2011–2023. A nested-logit demand system and Nash–Bertrand pricing recover dollar variable profits. The GFD value-difference identity then turns dynamic entry estimation into one logit regression per carrier. The renewal offset uses next-year CCPs. Thus 2023 supplies the boundary CCP input for the 2022 estimating year, while the estimation sample ends in 2022. A Wald test rejects stationarity of the carrier-year entry value for eight of the ten carriers. After we control for dollar profitability, the economic cost of stationarity is modest. The method therefore detects non-stationarity without loading the Coronavirus Disease 2019 (COVID-19) demand collapse onto a structural entry-value intercept.

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. Arcidiacono and Ellickson (2011) survey the broader

CCP/finite-dependence toolkit. Existing applications verify finite dependence through model-specific arguments, such as renewal actions or terminal paths. This paper replaces that step with a rank condition on the transition kernel. Many empirical DDC models instead use full-solution methods, including 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 with action-specific human-capital depreciation (Dix-Carneiro, 2014). We cite these as model-class contrasts. Arcidiacono and Miller (2020) show that counterfactual CCPs are generally not identified off short panels unless finite dependence holds, linking our existence test to identification as well as computation.

The term "generalized finite dependence" is due to Gayle (2021), who introduce it for dynamic discrete and continuous choice models. We adopt the terminology in the same spirit. Relative to that work, we give the arbitrary-ρ existence and rank characterization, the convex-QP feasibility test for signed weights, and the FD-class trace-optimality result. The closest methodological comparison is Kalouptsidi, Scott, and Souza-Rodrigues (KSSR; Kalouptsidi et al., 2021). They provide a linear-representation approach to counterfactual analysis in DDC models. Our framework is complementary. Theorem 8 gives a computational route for payoff-only counterfactuals after the analyst verifies finite dependence. Magnac and Thesmar (2002) establish that DDC models are generically underidentified absent normalization restrictions. Our results work within the usual conditional-independence setting and add finite dependence as an operational source of value-function cancellation. Berry and Compiani (2023) study dynamic models with serially correlated unobservables, which are outside the present scope.

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. In dynamic games, Aguirregabiria and Marcoux (2021) emphasize equilibrium restrictions. Our approach uses discrete-time transition structure. The rank test applies to arbitrary finite transition kernels, and the flow parameterization turns weight-finding into a convex QP. Adusumilli and Eckardt (2025) use temporal-difference learning to approximate value functions in DDC models with continuous or high-dimensional state spaces. Chen (2025) develops model-adaptive sieve approximations to the Bellman equation. GFD takes a different route: when finite dependence holds, it cancels the continuation value rather than approximating it. The closest Bellman-free alternative is the Euler-equation approach of Aguirregabiria and Magesan (2018). That approach replaces the Bellman fixed point with local first-order conditions in CCP space. GFD instead cancels the continuation value over a finite horizon via signed transition weights.

The signed weights connect the paper to econometric work on signed measures. Borusyak and Hull (2024) discuss negative weights in design-based specifications. de Chaisemartin and D'Haultfœuille (2020) show how negative weights can matter in two-way fixed-effects estimators with heterogeneous effects. Abadie et al. (2019) allow signed donor weights in synthetic control. Our weights apply to structural transition primitives, not to heterogeneous treatment effects. Exact finite dependence therefore preserves the payoff-difference estimand regardless of signs. Under approximate finite dependence, signed weights create a variance and bias-amplification concern through their norm.

The asymptotic theory builds on standard two-step M-estimation (Murphy and Topel, 1985; Newey and McFadden, 1994). The variance-input mapping connects to Newey (1990)'s effi-

cient instrumental-variables framework. The GFD QP/least-squares flow-weight machinery is a structurally constrained instance of the asymptotic-least-squares class (Pesendorfer and Schmidt-Dengler, 2008). The short latent-type extension in Appendix E uses the nonparametric mixture-identification condition of Kasahara and Shimotsu (2009). The controllability recursion is related to the discrete-time free-switching reachability literature (Liberzon, 2003; Sun and Ge, 2005).

1.1 Discount-Factor Identification

A companion paper, Hao et al. (2026), studies the identification of the discount factor and payoff function in stationary DDC models. Its object is different from ours. It does not construct a GFD estimator. It asks when β and flow payoffs are identified without imposing the common normalization that one action's payoff is zero. The paper shows that, in single-agent models with an exclusion restriction, β is identified up to a finite set. The cardinality is bounded by the state-space size. It also shows how nonparametric restrictions on payoffs can identify β. The restrictions include homogeneity and monotonicity. They also include concavity, zero cross-derivatives, and complementarity. Once β is identified, the payoff function is identified as well.

Finite dependence enters that problem through the order of the identifying equations. Without FD, eliminating the value function typically leaves terms such as $(I - \beta Q_d)(I - \beta Q_K)^{-1}$. The determinant and adjugate of the state-space inverse generate polynomials whose degree scales with the number of states used in the reduced system. Under ρ-period FD, the same inverse expansion truncates after ρ steps. In the one-period case,

$$(Q_d - Q_K)Q_K = 0 \quad \Rightarrow \quad (I - \beta Q_d)(I - \beta Q_K)^{-1} = I + \beta(Q_K - Q_d),$$

so the relevant entries are linear in β. More generally, FD reduces the polynomial order of the discount-factor restrictions to the finite dependence horizon and the number of rank-increasing restrictions, rather than the full dynamic-program state dimension. This is the connection to the present paper: our existence test tells the analyst when such a finite-horizon reduction is available, while the companion paper gives the discount-factor identification argument.

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 and the payoff-only counterfactual GFD identity (Section 3.8). Subsection 3.7 of Section 3 establishes asymptotic normality, the variance–input mapping, and optimal weight selection. Section 4 presents the four-type taxonomy with rank-test verification on eight benchmark and stylized DDC examples. It then reports Monte Carlo evidence on a single-agent investment model and a three-player Markov-perfect entry/exit game. Section 5 takes the framework to data, estimating a non-stationary entry game of ten U.S. carriers over 2011–2023 on a dollar demand-and-pricing structure. Section 6 concludes. Appendix D applies the GFD framework to Altug and Miller (1998)'s 14-parameter female labour-supply model on Panel Study of Income Dynamics (PSID) data. Appendix B collects the scalable computation machinery.

2 Existence of Finite Dependence

This section sets up the dynamic discrete choice framework and characterizes when a finite-dependence representation of value differences exists. The main result reduces existence at a given horizon to a single linear-feasibility test on the transition kernel.

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. This is the only sequence notation used throughout. We impose the standard regularity condition that payoffs are uniformly bounded across t. More general summability, $\sum_{s\geq 0} \beta^s \sup_{x,d} |u_{t+s}(x,d)| < \infty$, also suffices. The value functions V_t are then well-defined and finite even in the non-stationary case.

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

$$U_t(x_t, d_t, \epsilon_t) \;=\; u_t(x_t, d_t) + \epsilon_t(d_t), \tag{1}$$

where $u_t : \mathcal{X} \times \mathcal{D} \to \mathbb{R}$ 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 independently and identically distributed (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 $f_t(x_{t+1} \mid x_t, d_t)$, 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

$$v_t(x_t, d_t) \;=\; u_t(x_t, d_t) \;+\; \beta \sum_{x_{t+1} \in \mathcal{X}} V_{t+1}(x_{t+1})\, f_t(x_{t+1} \mid x_t, d_t), \tag{2}$$

where the ex-ante value function $V_t(x_t) = \int \max_{d_t} \{v_t(x_t, d_t) + \epsilon_t(d_t)\}\, dG(\epsilon_t)$. We write $\mathbf{p}_t(x_t) = (p_t(0 \mid x_t), \ldots, p_t(D-1 \mid x_t))$ for the equilibrium conditional choice probabilities.

2.2 The Arcidiacono–Miller Representation

Finite dependence methods identify structural parameters by expressing the value-function difference $v_t(x_t, d_t) - v_t(x_t, d'_t)$ as a finite sum of flow payoffs. Throughout the paper d_t denotes the analyzed action and d'_t the reference (baseline) action against which d_t is compared. The two are arbitrary distinct elements of $\mathcal{D}$. The QP formulation in Section 2.4 is symmetric in (d_t, d'_t), so the choice of which to label as "reference" is conventional. Identifying $v_t(x_t, d_t) - v_t(x_t, d'_t)$

requires finding two future continuation schemes starting from d_t and d'_t 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_t and any signed weights $\boldsymbol{\omega}(x_t) = (\omega(0 \mid x_t), \ldots, \omega(D-1 \mid x_t))$ satisfying $\sum_{d_t \in \mathcal{D}} \omega(d_t \mid x_t) = 1$, the value function admits the representation of Arcidiacono and Miller (2019, Theorem 1 and eq. (2.4), pp. 857–858):*

$$V_t(x_t) \;=\; \sum_{d_t \in \mathcal{D}} \omega(d_t \mid x_t) \big[\, v_t(x_t, d_t) + \psi_{d_t}(\mathbf{p}_t(x_t)) \,\big], \tag{3}$$

where $\psi_{d_t}(\mathbf{p}_t(x_t)) := V_t(x_t) - v_t(x_t, d_t)$. By the Hotz–Miller inversion, ψ_{d_t} is a known function of the CCP vector $\mathbf{p}_t(x_t)$ and the shock distribution $G(\epsilon)$.[12]

The representation (3) holds for any weights summing to unity, including signed weights: the structural opening that generalized finite dependence exploits. The ρ-period flow Φ of Section 2.4 is the multi-period analogue of the AM per-step weight ω. We define the augmented flow payoff

$$\mu_t(x_t, d_t) \;:=\; u_t(x_t, d_t) + \psi_{d_t}(\mathbf{p}_t(x_t)). \tag{4}$$

2.3 Finite Dependence and History-Conditional Weights

Finite dependence and the AM weight matrix. Arcidiacono and Miller (2019) define ρ-period finite dependence at a triple (x_t, d_t, d'_t) through Markov per-step weights $\omega_{t+\tau}(d_{t+\tau} \mid x_{t+\tau})$ for $\tau = 1, \ldots, \rho$. Each step satisfies the unit-sum condition $\sum_{d_{t+\tau}} \omega_{t+\tau}(d_{t+\tau} \mid x_{t+\tau}) = 1$. Finite dependence holds when the two paths starting from (x_t, d_t) and (x_t, d'_t) reach the same terminal state distribution after ρ steps.

The ρ-period expansion of the choice-specific value function $v_t(x_t, d_t)$ shows how the weight enters the value-function difference. Take a sequence of Markov per-step weights $\omega_{t+1}, \ldots, \omega_{t+\rho}$. These induce a length-$\rho$ path measure $\mathbb{Q}_\omega$. The measure is a genuine joint probability when the ω's are non-negative, and a signed path measure when they are not. It admits the explicit product form

$$\mathbb{Q}_\omega(x_{t+1:t+\rho},\, d_{t+1:t+\rho} \mid x_t, d_t) \;=\; \prod_{\tau=1}^{\rho} f_{t+\tau-1}(x_{t+\tau} \mid x_{t+\tau-1}, d_{t+\tau-1})\, \omega_{t+\tau}(d_{t+\tau} \mid x_{t+\tau}), \tag{5}$$

and the terminal-state distribution at time $t+\rho+1$ is obtained by integrating one additional transition:

$$\kappa_{t+\rho+1}(x_{t+\rho+1} \mid x_t, d_t) \;:=\; \sum_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}} \mathbb{Q}_\omega(x_{t+1:t+\rho},\, d_{t+1:t+\rho} \mid x_t, d_t)\, f_{t+\rho}(x_{t+\rho+1} \mid x_{t+\rho}, d_{t+\rho}). \tag{6}$$

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

[2] Notation crosswalk to Arcidiacono and Miller (2019) (QE 2019, pp. 857–858). Our $V_t(x_t)$ is their $V_t(x_t)$. Our $v_t(x_t, d_t)$ is their $v_{jt}(x_t)$ (eq. 2.3). Our $u_t(x_t, d_t)$ is their $u_{jt}(x_t)$. Our $\mathbf{p}_t(x_t)$ is their $p_t(x_t) \equiv (p_{1t}(x_t), \ldots, p_{Jt}(x_t))$. Our $f_t(x_{t+1} \mid x_t, d_t)$ is their $f_{jt}(x_{t+1} \mid x_t)$. Our $\psi_{d_t}(\mathbf{p}_t(x_t))$ is their $\psi_j[p_t(x_t)]$ (eq. 2.4). Our per-step weights $\omega(d_t \mid x_t)$ correspond to their $\omega_\tau(x_\tau, j)$ in the ρ-period expansion (eqs. 2.5–2.7). The single-period identity in (3) is the $\rho = 0$ case of their Theorem 1. The derivation is in Appendix F.2.

Writing $\mathbb{Q}_\omega(x_{t+\tau}, d_{t+\tau} \mid x_t, d_t)$ for the time-$(t+\tau)$ marginal of (5), recursive substitution of the AM representation (3) at each future date expresses $v_t(x_t, d_t)$ as the expected discounted sum of augmented flow payoffs μ along a path generated by $\omega_{t+1:t+\rho}$:

$$\begin{aligned} v_t(x_t, d_t) \;&= u_t(x_t, d_t) \;+\; \sum_{\tau=1}^{\rho} \beta^\tau \sum_{x_{t+\tau}, d_{t+\tau}} \mu_{t+\tau}(x_{t+\tau}, d_{t+\tau})\, \mathbb{Q}_\omega(x_{t+\tau}, d_{t+\tau} \mid x_t, d_t) \\ &\quad +\; \beta^{\rho+1} \sum_{x_{t+\rho+1}} V_{t+\rho+1}(x_{t+\rho+1})\, \kappa_{t+\rho+1}(x_{t+\rho+1} \mid x_t, d_t). \end{aligned} \tag{7}$$

If the terminal distributions match ($\kappa_{t+\rho+1}(x_{t+\rho+1} \mid x_t, d_t) = \kappa_{t+\rho+1}(x_{t+\rho+1} \mid x_t, d'_t)$ for all $x_{t+\rho+1}$), the continuation value $V_{t+\rho+1}(x_{t+\rho+1})$ cancels from the difference $v_t(x_t, d_t) - v_t(x_t, d'_t)$, yielding an expression in flow payoffs alone.

History-conditional weights and signed weights. The AM definition restricts $\omega_{t+\tau}$ to depend only on the contemporaneous state $x_{t+\tau}$, giving the joint product form (5). It does not condition on additional lags for a tractability reason: with per-step weights, a longer history makes terminal matching a non-convex polynomial system in $\{\omega_{t+\tau}\}$, which we resolve by the flow reparameterization of Section 2.4. We address this limitation and extend their representation in two ways.

(i) History-conditional weights. We let $\omega_{t+\tau}(d_{t+\tau} \mid x_{t:t+\tau}, d_{t:t+\tau-1})$ depend on every state observed and every action taken up to and including time $t+\tau$, and only that history. The same product structure as in (5) carries through:

$$\mathbb{Q}_\omega(x_{t+1:t+\rho},\, d_{t+1:t+\rho} \mid x_t, d_t) \;=\; \prod_{\tau=1}^{\rho} f_{t+\tau-1}(x_{t+\tau} \mid x_{t+\tau-1}, d_{t+\tau-1}) \cdot \omega_{t+\tau}(d_{t+\tau} \mid x_{t:t+\tau}, d_{t:t+\tau-1}). \tag{8}$$

We recover the Arcidiacono–Miller (AM) Markov form by dropping the prefix $(x_{t:t+\tau-1}, d_{t:t+\tau-1})$ from the conditioning. The history-conditional generalization enlarges the set of triples (x_t, d_t, d'_t) at which finite dependence can hold (Section 4.1). The primitive transition is first-order Markov in x_t (Assumption 3). We allow the continuation weights to vary across the realized histories that reach a common contemporaneous state. This offsets path-dependence in the accumulated flow that no single contemporaneous-state weight can match.

(ii) Signed weights. We also drop the non-negativity restriction on $\omega_{t+\tau}$ (and on the joint flow Φ introduced in Section 2.4).

Remark 1 (Convex hull vs. affine hull)**.** *Write $\boldsymbol{\pi}(d_{t+1:t+\rho})$ for the terminal-state distribution reached from (x_t, d_t) along the fixed action sequence $d_{t+1:t+\rho}$. Non-negative weights restrict the achievable terminal distribution to the convex hull of $\{\boldsymbol{\pi}(d_{t+1:t+\rho})\}_{d_{t+1:t+\rho}}$. Signed weights expand it to the affine hull. This unlocks finite dependence where no convex combination of these trajectory endpoints agrees across $d_t \neq d'_t$. The cost is variance inflation in the GFD effective regressor (§3.6).*

2.4 Flow Parameterization and the Convex QP

A direct search over the history-conditional weights $\{\omega_{t+\tau}\}_{\tau=1}^{\rho}$ is nonlinear in the decision variables: the terminal distribution $\kappa_{t+\rho+1}$ is a high-degree polynomial in $\{\omega_{t+\tau}\}$, so terminal-

matching reduces to a non-convex system of multivariate polynomial equations. We resolve this by reparameterizing the entire product (8) as a single joint flow Φ. The resulting feasibility conditions are linear and the matching objective is quadratic. The problem becomes a convex QP, so any solution is a global optimum.

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

$$\phi_{x_{t+1:t+\rho}, d_{t+1:t+\rho}}(x_t, d_t) \;=\; \prod_{\tau=1}^{\rho} f_{t+\tau-1}(x_{t+\tau} \mid x_{t+\tau-1}, d_{t+\tau-1})\, \omega_{t+\tau}(d_{t+\tau} \mid x_{t:t+\tau}, d_{t:t+\tau-1}), \tag{9}$$

where $\omega_{t+\tau}$ is a (possibly signed) history-conditional per-step weight summing to one over $d_{t+\tau}$. The conditioning $(x_{t:t+\tau}, d_{t:t+\tau-1})$ is exactly the information available to a forward-looking agent at the moment $d_{t+\tau}$ is chosen: states observed through time $t+\tau$ and actions taken through time $t+\tau-1$. The weight cannot look ahead to future state realizations $x_{t+\tau+1:t+\rho}$ or future actions $d_{t+\tau+1:t+\rho}$ beyond the current action $d_{t+\tau}$ at which it is evaluated. We call this property non-anticipativity (Remark 2).

Linear constraints on Φ. The flow matrix Φ admits the decomposition (9) for some valid history-conditional weight sequence if and only if it satisfies two linear constraints (Theorem 2 below): an initial flow constraint, and a flow-conservation constraint. With signed flows a mild no-ghost condition is also required (Remark 3).

(i) Initial flow. For each $x_{t+1} \in \mathcal{X}$, the total flow through histories that begin at state x_{t+1} equals the one-step arrival probability:

$$\sum_{x_{t+2:t+\rho}} \sum_{d_{t+1:t+\rho}} \phi_{(x_{t+1},\, x_{t+2:t+\rho}),\, d_{t+1:t+\rho}}(x_t, d_t) \;=\; f_t(x_{t+1} \mid x_t, d_t). \tag{10}$$

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

$$\begin{aligned} &\sum_{x_{t+\tau+2:t+\rho}} \sum_{d_{t+\tau+1:t+\rho}} \phi_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}}(x_t, d_t) \\ &\qquad = f_{t+\tau}(x_{t+\tau+1} \mid x_{t+\tau}, d_{t+\tau}) \cdot \sum_{x_{t+\tau+1:t+\rho}} \sum_{d_{t+\tau+1:t+\rho}} \phi_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}}(x_t, d_t), \end{aligned} \tag{11}$$

where the left-hand sum holds the prefix $(x_{t+1:t+\tau+1}, d_{t+1:t+\tau})$ fixed and the right-hand sum holds $(x_{t+1:t+\tau}, d_{t+1:t+\tau})$ fixed. Summing (10) over x_{t+1} shows that the entries of $\Phi(x_t, d_t)$ sum to one, so $\Phi(x_t, d_t)$ is a unit-mass signed measure on path space.

Theorem 2 (Constraint–decomposition equivalence)**.** *Fix $(x_t, d_t) \in \mathcal{X}\times\mathcal{D}$ and a horizon $\rho \geq 1$. The flow matrix $\Phi(x_t, d_t) \in \mathbb{R}^{S^\rho \times D^\rho}$ satisfies the initial-flow constraint* (10)*, the flow-conservation constraint* (11)*, and the no-ghost condition if and only if it admits the multiplicative decomposition* (9) *for some history-conditional weight sequence $\{\omega_{t+\tau}\}_{\tau=1}^{\rho}$ with $\sum_{d_{t+\tau}} \omega_{t+\tau}(d_{t+\tau} \mid \cdot) = 1$ at every history of non-zero marginal flow. The no-ghost condition requires that every history with zero state-prefix marginal flow $M_\tau^{\mathrm{s}} = 0$ also has zero action-prefix subflow $M_\tau^{\mathrm{sa}} = 0$. On every history of non-zero marginal flow, the implied weights $\omega_{t+\tau}$ are unique.*

The proof is constructive. On the forward direction, each $\omega_{t+\tau}$ is recovered as a ratio of two marginal path flows. The converse direction telescopes the product form against the constraints. Full details, including the marginal-flow notation used in the construction, are deferred to Appendix F.1.

Remark 2 (Non-anticipativity as a linear constraint)**.** *The flow-conservation constraint* (11) *is the linear-algebraic encoding of the non-anticipativity property already discussed: it forces each per-step weight $\omega_{t+\tau}$ to depend on past states and actions $(x_{t:t+\tau}, d_{t:t+\tau-1})$ alone, never on the future.*

Remark 3 (Signed flows and the no-ghost condition)**.** *The no-ghost condition is automatic for non-negative flows. With signed weights the linear constraints* (10)–(11) *alone can be met by offsetting positive and negative subflows below a zero state-prefix marginal. Such a flow satisfies the constraints yet admits no per-step multiplicative representation (the ratio defining $\omega_{t+\tau}$ is $0/0$). The downstream results, however, are stated on the linear signed-flow object directly. The existence test (Theorem 3) and the value-difference identity (Theorem 4) use the flow constraints and terminal matching, not the multiplicative weights. They therefore hold for the full linear class regardless of ghost mass. The no-ghost condition is needed only to read off non-anticipative per-step weights; it isolates the residual case of endogenous zero marginals created by cancellation. These are distinct from the structural zeros removed by reachability pruning (Appendix B).*

Terminal measure and the QP. The terminal-state signed measure induced by $\Phi(x_t, d_t)$ is

$$\kappa_{t+\rho+1}(x_{t+\rho+1} \mid x_t, d_t) \;=\; \sum_{x_{t+1:t+\rho}} \sum_{d_{t+1:t+\rho}} \phi_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}}(x_t, d_t) \cdot f_{t+\rho}(x_{t+\rho+1} \mid x_{t+\rho}, d_{t+\rho}). \tag{12}$$

It is a genuine distribution only when $\Phi \geq 0$. With signed weights its entries may be negative. Matching it across d_t and d_t' is what cancels the continuation value. The path probability is already absorbed into Φ; no separate factor appears. With analyzed action d_t and reference action d_t' from state x_t, the full optimization problem is

$$\begin{aligned} \min_{\Phi(x_t,d_t),\, \Phi(x_t,d_t')} \; & \sum_{x_{t+\rho+1} \in \mathcal{X}} \big(\kappa_{t+\rho+1}(x_{t+\rho+1} \mid x_t, d_t) - \kappa_{t+\rho+1}(x_{t+\rho+1} \mid x_t, d_t') \big)^2 \\ \text{s.t.} \quad & \text{constraints (10) and (11) for both } \Phi(x_t, d_t) \text{ and } \Phi(x_t, d_t'). \end{aligned} \tag{13}$$

The two flow matrices share the same flow-conservation structure but have different initial-flow constraints: $\Phi(x_t, d_t)$ is constrained by $f_t(\cdot \mid x_t, d_t)$ on the right-hand side of (10), while $\Phi(x_t, d_t')$ is constrained by $f_t(\cdot \mid x_t, d_t')$. All constraints are linear and the objective is quadratic in $(\Phi(x_t, d_t), \Phi(x_t, d_t'))$, making this a convex QP.

Handling zero transition entries. When $f_t(x_{t+1} \mid x_t, d_t) = 0$, every length-$\rho$ path through the unreachable transition has zero visit probability, so the corresponding entries of $\Phi(x_t, d_t)$ are identically zero by (9). We prune these from the QP (13), keeping only entries indexed by reachable $(x_{t+1:t+\rho}, d_{t+1:t+\rho})$. Pruning both shrinks the optimization (often substantially when transitions are sparse; Appendix B) and prevents the flow-conservation constraints (11) from being vacuously satisfied on impossible histories.

A fully worked $S = 2$, $D = 2$, $\rho = 1$ example is given in Appendix A. It instantiates the flow matrix, the two linear constraints, the multiplicative decomposition, and the matching QP on the smallest non-trivial case.

2.5 Existence Test for Finite Dependence

This subsection states the paper's main result: pointwise ρ-step finite dependence at (x_t, d_t, d'_t) holds if and only if a single sparse linear system on the structural transitions is feasible (Theorem 3). The test is operationally checkable before any data step, applies to arbitrary transition kernels, and recovers the previously hand-constructed FD-amenable classes (Section 4.1). The supporting Theorem 2 (constraint–decomposition equivalence) underwrites this result by showing that the linear constraints on Φ are exactly those needed to admit a per-step multiplicative form.

Pointwise finite dependence and linear feasibility.

Definition 1 (Pointwise ρ-period finite dependence)**.** Fix a triple $(x_t, d_t, d'_t) \in \mathcal{X} \times \mathcal{D}^2$ with $d_t \neq d'_t$ (with d_t the analyzed action and d'_t the reference action), and a horizon $\rho \geq 1$. We say pointwise ρ-period finite dependence holds at (x_t, d_t, d'_t) if there exist flow matrices $\Phi(x_t, d_t)$ and $\Phi(x_t, d'_t)$ that satisfy initial flow (10) and flow conservation (11) and induce, via (12), identical terminal distributions:

$$\kappa_{t+\rho+1}(x_{t+\rho+1} \mid x_t, d_t) \;=\; \kappa_{t+\rho+1}(x_{t+\rho+1} \mid x_t, d'_t) \qquad \forall\, x_{t+\rho+1} \in \mathcal{X}. \tag{14}$$

The following linear system has solvability equivalent to Definition 1. Let $\operatorname{vec}\Phi(x_t, d_t) \in \mathbb{R}^{S^\rho D^\rho}$ denote the column-stack of entries of $\Phi(x_t, d_t)$, indexed by $(x_{t+1:t+\rho}, d_{t+1:t+\rho})$ in the calendar-time slice notation introduced earlier (so the future path runs over states $x_{t+1}, \ldots, x_{t+\rho}$ and actions $d_{t+1}, \ldots, d_{t+\rho}$).

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

Initial-flow block $\mathbf{A}_{\text{init}} \in \{0,1\}^{S \times S^\rho D^\rho}$, with rows indexed by $x_{t+1} \in \mathcal{X}$:

$$(\mathbf{A}_{\text{init}})_{x_{t+1},\,(x_{t+1:t+\rho},\, d_{t+1:t+\rho})} \;=\; \mathbf{1}\{(\text{first state of } x_{t+1:t+\rho}) = x_{t+1}\}. \tag{15}$$

Flow-conservation (non-anticipativity constraint, NAC) block $\mathbf{A}_{\text{nac}}$. Bring (11) to one side:

for each $\tau \in \{1, \dots, \rho-1\}$, each prefix $(x_{t+1:t+\tau}, d_{t+1:t+\tau}) \in \mathcal{X}^\tau \times \mathcal{D}^\tau$, and each next state $x_{t+\tau+1} \in \mathcal{X}$,

$$\underbrace{\sum_{x_{t+\tau+2:t+\rho}} \sum_{d_{t+\tau+1:t+\rho}} \phi_{x_{t+1:t+\rho}, d_{t+1:t+\rho}}(x_t, d_t)}_{\text{flow through } (x_{t+1:t+\tau+1}, d_{t+1:t+\tau})} - f_{t+\tau}(x_{t+\tau+1} \mid x_{t+\tau}, d_{t+\tau}) \underbrace{\sum_{x_{t+\tau+1:t+\rho}} \sum_{d_{t+\tau+1:t+\rho}} \phi_{x_{t+1:t+\rho}, d_{t+1:t+\rho}}(x_t, d_t)}_{\text{flow through } (x_{t+1:t+\tau}, d_{t+1:t+\tau})} = 0. \tag{16}$$

Each instance of (16) is one linear equation in $\operatorname{vec}\Phi$. The coefficients are +1 on every entry $(x_{t+1:t+\rho}, d_{t+1:t+\rho})$ extending the prefix $(x_{t+1:t+\tau+1}, d_{t+1:t+\tau})$ and $-f_{t+\tau}(x_{t+\tau+1} \mid x_{t+\tau}, d_{t+\tau})$ on every entry extending $(x_{t+1:t+\tau}, d_{t+1:t+\tau})$. We take $\mathbf{A}_{\text{nac}}$ to be the matrix whose rows are these coefficient vectors, indexed by the quadruple $(\tau, x_{t+1:t+\tau+1}, d_{t+1:t+\tau})$. 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). Hence $\mathbf{A}_{\text{nac}} \in \mathbb{R}^{C \times S^\rho D^\rho}$ with C as in (17):

$$C \;:=\; \text{number of NAC rows} \;=\; \sum_{\tau=1}^{\rho-1} S^{\tau+1} D^\tau. \tag{17}$$

By construction $\mathbf{A}_{\text{nac}} \operatorname{vec}\Phi = \mathbf{0}$ is exactly (11) written across all $(\tau, x_{t+1:t+\tau+1}, d_{t+1:t+\tau})$.

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

$$(\mathbf{T}_{\text{last}})_{x', (x_{t+1:t+\rho}, d_{t+1:t+\rho})} \;=\; f_{t+\rho}(x' \mid x_{t+\rho}, d_{t+\rho}). \tag{18}$$

By (12), $\mathbf{T}_{\text{last}} \operatorname{vec}\Phi(x_t, d_t) = \kappa_{t+\rho+1}(\cdot \mid x_t, d_t)$.

Stack the two scenarios into one decision variable

$$\mathbf{x} \;:=\; \begin{pmatrix} \operatorname{vec}\Phi(x_t, d_t) \\ \operatorname{vec}\Phi(x_t, d'_t) \end{pmatrix} \;\in\; \mathbb{R}^{2S^\rho D^\rho}.$$

With $\mathbf{f}^{(d_t)} := \big(f_t(x_{t+1} \mid x_t, d_t)\big)_{x_{t+1} \in \mathcal{X}} \in \mathbb{R}^S$, define

$$\mathbf{A} \;:=\; \begin{pmatrix} \mathbf{A}_{\text{init}} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\text{init}} \\ \mathbf{A}_{\text{nac}} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\text{nac}} \\ \mathbf{T}_{\text{last}} & -\mathbf{T}_{\text{last}} \end{pmatrix}, \qquad \mathbf{b} \;:=\; \begin{pmatrix} \mathbf{f}^{(d_t)} \\ \mathbf{f}^{(d'_t)} \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix}. \tag{19}$$

The matrix $\mathbf{A} \in \mathbb{R}^{(3S+2C) \times 2S^\rho D^\rho}$ depends on the transition f, the horizon ρ, and the initial state x_t (via $\mathbf{A}_{\text{init}}$ rows for x_{t+1}). The vector $\mathbf{b}$ additionally depends on the compared actions (d_t, d'_t).

Theorem 3 (Pointwise FD as linear feasibility)**.** *Fix $(x_t, d_t, d'_t) \in \mathcal{X} \times \mathcal{D}^2$ with $d_t \neq d'_t$, and let $\mathbf{A}, \mathbf{b}$ be the matrix and vector constructed in* (19)*. Then pointwise ρ-period finite dependence at (x_t, d_t, d'_t)*

(Definition 1) holds if and only if

$$\mathbf{b} \in \text{Range}(\mathbf{A}), \qquad \textit{equivalently} \qquad \text{rank}\big([\,\mathbf{A} \mid \mathbf{b}\,]\big) = \text{rank}(\mathbf{A}). \tag{20}$$

Whenever (20) *holds, any solution* $\mathbf{x} \in \mathbb{R}^{2S^\rho D^\rho}$ *of* $\mathbf{A}\,\mathbf{x} = \mathbf{b}$ *yields, by inverse vectorization, flow matrices* $\Phi(x_t, d_t)$ *and* $\Phi(x_t, d_t')$ *that achieve terminal-distribution matching with horizon* ρ*.*

The proof is in Appendix F.3.

Remark 4 (Computational interpretation)**.** *Theorem 3 reduces verification of pointwise finite dependence at* (x_t, d_t, d_t') *to a single sparse linear-algebra computation. The matrix* $\mathbf{A}$ *has* $(3S + 2C) \times 2S^\rho D^\rho$ *entries, with* C *as in* (17)*. Both* $\mathbf{A}$ *and* $\mathbf{b}$ *are built from* f *and* (x_t, d_t, d_t') *alone, with no appeal to the value function or CCPs. In practice we test* $\mathbf{b} \in \text{Range}(\mathbf{A})$ *via sparse least squares using the LSQR algorithm of Paige and Saunders (1982). This is an iterative solver for* $\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$ *that requires only matrix–vector products with the sparse* $\mathbf{A}$ *and* $\mathbf{A}^\top$*, avoiding the dense normal equations* $\mathbf{A}^\top\mathbf{A}$*. We declare 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 $(x_t, d_t, d_t') \in \mathcal{X} \times \mathcal{D}^2$ with $d_t \neq d_t'$. By Theorem 3, this is verified by running the LSQR test of Remark 4 once per triple.

3 The Generalized Finite Dependence (GFD) Estimator

Theorem 3 settles when FD is achievable. We now develop the natural estimator for the FD-feasible class. Generalized Finite Dependence (GFD) is a class of two-step CCP estimators of the structural payoff parameter θ. Each member takes any FD-feasible flow-matrix collection $\widehat{\mathbf{\Phi}}$ satisfying the test of Section 2.5 as an explicit input. Treating $\widehat{\mathbf{\Phi}}$ 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 $\widehat{\mathbf{\Phi}}$ 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. The data are generated by the DDC model of Section 2 at an unknown structural parameter $\theta_0 \in \Theta \subseteq \mathbb{R}^{d_\theta}$. The flow payoff is $u_t(x_t, d_t; \theta)$ for $d_t \in \mathcal{D}$, with reference normalization $u_t(x_t, 0; \theta) \equiv 0$.[3] Throughout this section we maintain Assumptions 1–4 together with:

Assumption 5 (Type-I extreme-value shocks)**.** *The shocks* $\epsilon_t(d_t)$ *are i.i.d. Type-I extreme-value across* (t, d_t)*, so that the Hotz–Miller correction takes the closed form given in the footnote to Lemma 1 and the equilibrium CCP is the multinomial logit* $p_t(d_t \mid x_t) = \exp(v_t(x_t, d_t))/\sum_{d_t''} \exp(v_t(x_t, d_t''))$*.*

[3] Subscript convention: a 0 in an action argument (e.g. $u_t(x_t, 0)$, Λ_0, $\Delta\phi$ relative to 0) denotes the reference action; a 0 on a parameter or CCP (e.g. θ_0, $\mathbf{p}_0$) denotes the true/population value.

First-stage estimates. The researcher constructs nonparametric estimates $\widehat{f}_t(x_{t+1} \mid x_t, d_t)$ and $\widehat{p}_t(d_t \mid x_t)$ from the panel by cell frequency, sieve, or kernel. These satisfy the $\sqrt{N}$-consistency/asymptotic-linearity condition of Assumption 7. The plug-in Hotz–Miller correction is $\psi_{d_t}(\widehat{\mathbf{p}}_t(x_t)) = \gamma_E - \ln \widehat{p}_t(d_t \mid x_t)$.

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

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

$$\widehat{\mathbf{\Phi}} \;=\; \left\{ \widehat{\Phi}(x_t, d_t) \,:\, x_t \in \mathcal{X},\; d_t \in \mathcal{D} \right\} \tag{21}$$

indexed by every (x_t, d_t) pair, where each $\widehat{\Phi}(x_t, d_t) \in \mathbb{R}^{S^\rho \times D^\rho}$ is a flow matrix in the sense of Section 2.4. The researcher supplies the entire collection as input.

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

(a) each $\widehat{\Phi}(x_t, d_t)$ satisfies the initial-flow constraint (10) and the flow-conservation (NAC) constraint (11), both with the transition $\widehat{f}$ in place of f; and

(b) the induced terminal distribution $\widehat{\kappa}_{t+\rho+1}(\cdot \mid x_t, d_t)$ from (12) is the same for every $d_t \in \mathcal{D}$:

$$\widehat{\kappa}_{t+\rho+1}(x_{t+\rho+1} \mid x_t, d_t) \;=\; \widehat{\kappa}_{t+\rho+1}(x_{t+\rho+1} \mid x_t, d_t'), \qquad \forall\, d_t, d_t' \in \mathcal{D},\; \forall\, x_{t+\rho+1} \in \mathcal{X}. \tag{22}$$

Condition (a) is the per-action structural feasibility from Theorem 2. Condition (b) extends the FD matching condition to all D actions simultaneously. It is what makes the continuation value cancel from every value-difference $v_t(x_t, d_t; \theta) - v_t(x_t, 0; \theta)$ with $d_t \neq 0$.

By default we use the canonical input: for each x_t, solve the joint feasibility QP enforcing (a) and (b) and take the Moore–Penrose minimum-ℓ_2-norm solution. This is a single sparse linear solve per x_t. The variance-optimal choice is treated in §3.7, and structure-exploiting closed forms (e.g. the Type S shift-register / trivial-flush construction) in §4.1.

3.3 The GFD Value-Difference Identity

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

Theorem 4 (GFD value-difference identity)**.** *Suppose Assumptions 1–5 hold and that $\mathbf{\Phi} = \{\Phi(x_t, d_t)\}_{x_t, d_t}$ is FD-feasible at horizon ρ under the true transitions f (Definition 2 with f in place of $\widehat{f}$). Then for every $(x_t, d_t) \in \mathcal{X} \times \mathcal{D}$ with $d_t \neq 0$,*

$$\begin{aligned} v_t(x_t, d_t; \theta) \;-\; v_t(x_t, 0; \theta) \;&= u_t(x_t, d_t; \theta) \\ &\quad + \sum_{\tau=1}^{\rho} \beta^\tau \sum_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}} \Delta\phi_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}}(x_t, d_t)\, \mu_{t+\tau}(x_{t+\tau}, d_{t+\tau}; \theta), \end{aligned} \tag{23}$$

with $\Delta\phi_{x_{t+1:t+\rho},d_{t+1:t+\rho}}(x_t,d_t) := \phi_{x_{t+1:t+\rho},d_{t+1:t+\rho}}(x_t,d_t) - \phi_{x_{t+1:t+\rho},d_{t+1:t+\rho}}(x_t,0)$*, and* $\mu_{t+\tau}$ *is the augmented flow payoff* (4)*, evaluated at the population CCPs* $\mathbf{p}$*. The right-hand side depends only on the flow input* $\mathbf{\Phi}$*, the parameter* θ*, and the population CCPs; the continuation value* $V_{t+\rho+1}$ *does not appear. The equality to the structural value-difference therefore holds at the true* θ_0*, where* $\mathbf{p}$ *is the equilibrium CCP profile. For* $\theta \neq \theta_0$ *the right-hand side defines the GFD pseudo value-difference* $\widetilde{v}^{\rho\text{-}FD}(\cdot;\theta,\mathbf{\Phi},\mathbf{p})$*, a known function of* θ *with* $\mathbf{p}$ *fixed. This pseudo value-difference is the object the pseudo-likelihood maximizes. It is not the Bellman difference of the model at* $u(\cdot;\theta)$*.*

The proof is in Appendix F.4.

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 (23):

$$\begin{aligned}\widetilde{v}^{\rho\text{-FD}}(x_t,d_t;\,\theta,\widehat{\mathbf{\Phi}},\widehat{\mathbf{p}}) \;:=\; & u_t(x_t,d_t;\theta) \\ & + \sum_{\tau=1}^{\rho}\beta^{\tau}\sum_{x_{t+1:t+\rho},\,d_{t+1:t+\rho}} \Delta\widehat{\phi}_{x_{t+1:t+\rho},\,d_{t+1:t+\rho}}(x_t,d_t)\,\widehat{\mu}_{t+\tau}(x_{t+\tau},d_{t+\tau};\theta),\end{aligned} \tag{24}$$

where $\Delta\widehat{\phi}$ and $\widehat{\mu}_t$ are the sample analogs of $\Delta\phi$ and the augmented flow payoff μ_t in (23)/(4), evaluated at $\widehat{\mathbf{p}}$. We impose the normalization $\widetilde{v}^{\rho\text{-FD}}(x_t,0;\cdot)=0$ for every x_t. By Assumption 5, the FD-implied conditional choice probability is the multinomial logit

$$\Lambda_{d_t}(x_t;\,\theta,\widehat{\mathbf{\Phi}},\widehat{\mathbf{p}}) \;:=\; \frac{\exp\big(\widetilde{v}^{\rho\text{-FD}}(x_t,d_t;\,\theta,\widehat{\mathbf{\Phi}},\widehat{\mathbf{p}})\big)}{\sum_{d_t''\in\mathcal{D}}\exp\big(\widetilde{v}^{\rho\text{-FD}}(x_t,d_t'';\,\theta,\widehat{\mathbf{\Phi}},\widehat{\mathbf{p}})\big)}. \tag{25}$$

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

$$\widehat{\theta}^{\text{GFD}}\big(\widehat{\mathbf{\Phi}}\big) \;:=\; \arg\max_{\theta\in\Theta}\;\frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T}\log\Lambda_{d_{it}}\big(x_{it};\,\theta,\widehat{\mathbf{\Phi}},\widehat{\mathbf{p}}\big). \tag{26}$$

What $\widehat{\theta}^{\text{GFD}}$ estimates. The estimator targets the structural payoff parameter θ_0 governing $u_t(x_t,d_t;\theta_0)$. Under exact ρ-period FD with $\mathbf{\Phi}$ feasible under the true f, Theorem 4 gives $\widetilde{v}^{\rho\text{-FD}}(x_t,d_t;\,\theta_0,\mathbf{\Phi},\mathbf{p}) = v_t(x_t,d_t;\theta_0) - v_t(x_t,0;\theta_0)$. Hence $\Lambda_{d_t}(x_t;\,\theta_0,\mathbf{\Phi},\mathbf{p}) = p_{0,t}(d_t \mid x_t)$ at the true parameter, and the population GFD log-likelihood (26) is the correctly specified conditional multinomial log-likelihood. The GFD estimator inherits the standard maximum-likelihood-estimator (MLE) consistency. For a $\sqrt{N}$-consistent first-stage $(\widehat{f},\widehat{\mathbf{p}})$, $\widehat{\theta}^{\text{GFD}}(\widehat{\mathbf{\Phi}}) \xrightarrow{p} \theta_0$. Two conditions are required: the generated input converges to a population FD-feasible limit $\widehat{\mathbf{\Phi}} \xrightarrow{p} \mathbf{\Phi}^* \in \mathcal{F}(\cdot;f)$, and the population criterion identifies θ_0 (non-singular GFD information $\mathbf{J}(\mathbf{\Phi}^*)$). Theorem 5 makes these conditions precise.

What $\widehat{\Phi}$ does in the estimator. The flow input $\widehat{\mathbf{\Phi}}$ enters (26) through the GFD pseudo value-difference (24). It packs the ρ-period continuation cancelation into a finite weighted sum over paths. The estimator therefore never solves the Bellman equation at any θ (see the cost discussion below). 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. The variance-minimizing choice of $\widehat{\mathbf{\Phi}}$ is characterized in Section 3.7.

3.5 The GFD Estimation Procedure

Putting the pieces together, the full GFD procedure is:

Step 1 (First stage). Construct the first-stage inputs $\widehat{f}_t, \widehat{p}_t$ and the plug-in Hotz–Miller correction $\psi_{d_t}(\widehat{\mathbf{p}}_t(x_t))$ of Section 3.1.

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

$$\min_{\{\Phi(x_t,d_t)\}_{d_t\in\mathcal{D}}} \sum_{d_t=1}^{D-1} \sum_{x_{t+\rho+1}\in\mathcal{X}} \left(\widehat{\kappa}_{t+\rho+1}(x_{t+\rho+1} \mid x_t, d_t) - \widehat{\kappa}_{t+\rho+1}(x_{t+\rho+1} \mid x_t, 0)\right)^2 \tag{27}$$

subject to (10)–(11) (with $\widehat{f}$) for each $\Phi(x_t, d_t)$. In the non-stationary case, solve it separately for each calendar time t. The constraints then use the time-specific kernels $\widehat{f}_t, \widehat{f}_{t+1}, \ldots, \widehat{f}_{t+\rho}$. Then $\widehat{\mathbf{\Phi}}$ is indexed by (x_t, t) unless stationarity is imposed. The minimum is zero iff a single common reference terminal distribution $\boldsymbol{\kappa}^* = \boldsymbol{\kappa}(\cdot \mid x_t, 0)$ is FD-matchable by every action d_t simultaneously. For $D = 2$ it reduces to pointwise FD at $(x_t, 1, 0)$ (Theorem 3). For $D > 2$ it is the stronger requirement that the D per-action achievable affine sets share a common point, which the QP tests directly. The default canonical input is as defined in Section 3.2. A strictly positive minimum signals that (ρ-step, joint) FD fails at (x_t, t). The procedure then increases ρ. If ρ is capped, it flags the cell. The resulting $\widehat{\mathbf{\Phi}}$ is then only approximately FD-feasible, and the downstream value-difference inherits the residual.

Step 3 (Pseudo-likelihood maximization). For every (x_t, d_t) in the data, the GFD pseudo value-difference $\widetilde{v}^{\rho\text{-FD}}(x_t, d_t; \theta, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}})$ is a known function of θ via (24). Solve (26) for $\widehat{\theta}^{\text{GFD}}(\widehat{\mathbf{\Phi}})$ by Newton or quasi-Newton iteration. When the flow payoff is linear in θ, Step 3 simplifies further (Section 3.6).

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. With LSQR and exploited sparsity it scales roughly as $|\mathcal{X}| \cdot S^\rho D^\rho$. This cost is paid once and does not repeat across parameter evaluations. Step 3 is a low-dimensional (d_θ-variable) optimization, concave whenever $\widetilde{v}^{\rho\text{-FD}}$ is linear in θ (the logit-type case). Crucially, none of Steps 2–3 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.

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 $\widehat{\theta}^{\text{GFD}}$. A

variance-optimized input (Section 3.7) gives another, and a non-negative AM-style input gives a third when feasible.

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

When the flow payoff is linear in θ, the GFD pseudo value-difference (24) becomes linear in θ as well, and Step 3 of the procedure reduces to a standard logit MLE on a generated design.

Assumption 6 (Linear-in-parameter payoff)**.** *The systematic flow payoff is linear in θ:*

$$u_t(x_t, d_t; \theta) \;=\; z_t(x_t, d_t)^\top \theta, \qquad d_t \in \mathcal{D}, \tag{28}$$

where $z_t : \mathcal{X} \times \mathcal{D} \to \mathbb{R}^{d_\theta}$ is a known regressor, normalized so that the reference action $d_t = 0$ has zero payoff: $z_t(x_t, 0) = \mathbf{0}$ for every x_t and t.

Under Assumption 6, substituting $z_{t+\tau}(x_{t+\tau}, d_{t+\tau})^\top \theta + \psi_{d_{t+\tau}}$ into (24) and grouping the θ-dependent and θ-independent terms gives

$$\widetilde{v}^{\rho\text{-FD}}(x_t, d_t;\, \theta, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) \;=\; \widetilde{H}_t(x_t, d_t;\, \widehat{\mathbf{\Phi}})^\top \theta \;+\; \widetilde{h}_t(x_t, d_t;\, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}), \tag{29}$$

where the effective regressor $\widetilde{H}_t \in \mathbb{R}^{d_\theta}$ and the effective offset $\widetilde{h}_t \in \mathbb{R}$ are

$$\widetilde{H}_t(x_t, d_t;\, \widehat{\mathbf{\Phi}}) \;:=\; z_t(x_t, d_t) \;+\; \sum_{\tau=1}^{\rho} \beta^\tau \sum_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}} \Delta\widehat{\phi}_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}}(x_t, d_t) \;\cdot\; z_{t+\tau}(x_{t+\tau}, d_{t+\tau}), \tag{30}$$

$$\widetilde{h}_t(x_t, d_t;\, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) \;:=\; \sum_{\tau=1}^{\rho} \beta^\tau \sum_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}} \Delta\widehat{\phi}_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}}(x_t, d_t) \;\cdot\; \psi_{d_{t+\tau}}(\widehat{\mathbf{p}}_{t+\tau}(x_{t+\tau})), \tag{31}$$

where the path-flow difference is

$$\Delta\widehat{\phi}_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}}(x_t, d_t) \;:=\; \widehat{\phi}_{x_{t+1:t+\rho}, d_{t+1:t+\rho}}(x_t, d_t) \;-\; \widehat{\phi}_{x_{t+1:t+\rho}, d_{t+1:t+\rho}}(x_t, 0).$$

Here $\Delta\widehat{\phi}$ is the contrast between the analyzed action d_t and the reference action 0 at each length-ρ continuation path. With this decomposition the GFD log-likelihood (26) becomes a standard logit log-likelihood with generated regressor $\widetilde{H}_t$ and generated offset $\widetilde{h}_t$:

$$\widehat{\theta}^{\text{GFD}}\big(\widehat{\mathbf{\Phi}}\big) \;=\; \arg\max_{\theta \in \Theta} \; \frac{1}{NT} \sum_{i,t} \log \frac{\exp\big(\widetilde{H}_t(x_{it}, d_{it})^\top \theta + \widetilde{h}_t(x_{it}, d_{it})\big)}{\sum_{d_t''} \exp\big(\widetilde{H}_t(x_{it}, d_t'')^\top \theta + \widetilde{h}_t(x_{it}, d_t'')\big)}. \tag{32}$$

The objective is concave in θ, and the dependence of the asymptotic variance on the flow input runs entirely through the effective regressor $\widetilde{H}_t$ (Section 3.7).

3.7 Variance–Input Mapping and the Variance-Optimal GFD Estimator

The GFD estimator is consistent for any FD-feasible flow input, but its asymptotic variance depends on that input. We state the asymptotic distribution and identify the variance-optimal

input within the FD-feasible class. Both results are purely theoretical. Monte Carlo evidence is deferred to Section 4; we preview two of its findings below, namely the near-optimality of the canonical input and the efficiency multiplier of Remark 5.

Null-space parameterization of FD-feasible inputs. The joint feasibility QP (27) is generically under-determined at $\rho \geq 2$: the feasible set $\mathcal{F}(x_t; f)$ defined by (10)–(11) together with the common-terminal condition (22) is an affine variety. Writing $\mathbf{\Phi}^{\mathrm{MP}}(x_t; f) \in \mathcal{F}(x_t; f)$ for the Moore–Penrose minimum-ℓ_2-norm element (the canonical input) and $\mathbf{N}(x_t; f)$ for a basis of the constraint null space, every feasible flow admits the parameterization

$$\mathbf{\Phi}(x_t; \mathbf{q}) \;=\; \mathbf{\Phi}^{\mathrm{MP}}(x_t; f) + \mathbf{N}(x_t; f)\, \mathbf{q}, \tag{33}$$

with the null-space coordinate $\mathbf{q}$ ranging over a Euclidean space whose dimension equals that of $\mathcal{F}(x_t; f)$.

We impose two standard two-step regularity conditions.

Assumption 7 (First-stage regularity)**.** *The first-stage estimators $\widehat{f}, \widehat{p}$ are $\sqrt{N}$-consistent and asymptotically linear: $\sqrt{N}\,(\widehat{f} - f,\ \widehat{p} - p) = N^{-1/2} \sum_{i=1}^{N} \boldsymbol{\psi}^{\mathrm{fs}}(\zeta_i) + o_p(1)$ for an influence function $\boldsymbol{\psi}^{\mathrm{fs}}$ with $\mathbb{E}[\boldsymbol{\psi}^{\mathrm{fs}}] = \mathbf{0}$ and finite second moment. The population CCPs are bounded away from 0 and 1 on the support, $p_t(d \mid x) \in (\underline{c}, 1 - \underline{c})$ for some $\underline{c} > 0$, so the Hotz–Miller correction $\psi_{d_t} = \gamma_E - \ln p_t(d_t \mid x_t)$ is well-defined.*

Assumption 8 (Regular flow-input domain)**.** *There is a compact set $\mathcal{Q} \subseteq \mathbb{R}^{d_q}$ on which the GFD information $\mathbf{J}(\mathbf{\Phi}(\mathbf{q}))$ is uniformly bounded away from singularity: $\inf_{\mathbf{q} \in \mathcal{Q}} \sigma_{\min}\big(\mathbf{J}(\mathbf{\Phi}(\mathbf{q}))\big) > 0$, and the meat $\mathbf{\Omega}(\mathbf{\Phi}(\mathbf{q}))$ is finite and continuous on $\mathcal{Q}$.*

Assumption 9 (Local rank/feasibility stability)**.** *In a neighbourhood of the true f, the FD-feasible set $\mathcal{F}(\cdot; f)$ and the null space $\mathcal{N}(f)$ of the feasibility system vary smoothly in f at constant rank. The population selection map $f \mapsto \mathbf{\Phi}(\cdot; f)$ is differentiable. Its derivative $\partial_f \mathbf{\Phi}$ is tangent to $\mathcal{F}(\cdot; f)$. This holds for the renewal, shift-register, and Kronecker classes. There feasibility is governed by a fixed-rank linear structure. It can fail at data-generating processes (DGPs) where finite dependence relies on knife-edge zeros, rank restrictions, or deterministic transitions. The null space of such DGPs is unstable under nonparametric estimation of f.*

Theorem 5 (Consistency and asymptotic normality of the GFD estimator)**.** *Suppose Assumptions 1–5 and 7 hold; $\widehat{\mathbf{\Phi}}$ is FD-feasible under $\widehat{f}$ for every N with $\widehat{\mathbf{\Phi}} \xrightarrow{p} \mathbf{\Phi}^*$ FD-feasible under f; and the GFD information $\mathbf{J}(\mathbf{\Phi}^*) := -\frac{1}{T} \sum_{t=1}^{T} \mathbb{E}_t[\nabla^2_\theta \log \Lambda_{d_{it}}(x_{it}; \theta_0, \mathbf{\Phi}^*, \mathbf{p})]$ is non-singular. The asymptotics are taken over $N \to \infty$ with T fixed and the agents i i.i.d.; $\mathbb{E}_t$ denotes the expectation under the period-t law of (x_{it}, d_{it}) (a time-average, since the model is non-stationary). Then $\widehat{\theta}^{\mathrm{GFD}}(\widehat{\mathbf{\Phi}}) \xrightarrow{p} \theta_0$ for every such $\widehat{\mathbf{\Phi}}$, and*

$$\sqrt{N}\big(\widehat{\theta}^{\mathrm{GFD}}(\widehat{\mathbf{\Phi}}) - \theta_0\big) \;\xrightarrow{d}\; \mathcal{N}\big(\mathbf{0},\, \mathbf{\Sigma}(\mathbf{\Phi}^*)\big), \tag{34}$$

where $\mathbf{\Sigma}(\mathbf{\Phi}) := \mathbf{J}(\mathbf{\Phi})^{-1}\, \mathbf{\Omega}(\mathbf{\Phi})\, \mathbf{J}(\mathbf{\Phi})^{-1}$ and $\mathbf{\Omega}(\mathbf{\Phi})$ is the variance of the per-agent GFD score $\frac{1}{T} \sum_{t=1}^{T} \nabla_\theta \log \Lambda_{d_{it}}(x_{it}; \theta_0, \mathbf{\Phi}, \mathbf{p})$ (which carries the within-agent serial dependence across t), augmented by the first-stage CCP plug-in correction. The generated input $\widehat{\mathbf{\Phi}}$ is FD-feasible under $\widehat{f}$, and $\widehat{f}$ enters the criterion only through $\widehat{\mathbf{\Phi}}$. The GFD value-difference identity (Theorem 4) holds exactly at the

true f for every FD-feasible $\mathbf{\Phi}$. The population score therefore has mean zero at θ_0 along the f-feasible manifold. The gradient $\partial_{\mathbf{\Phi}} \mathbb{E}[\nabla_\theta \log \Lambda_{d_t}(x_t; \theta_0, \mathbf{\Phi}, \mathbf{p})]$ then vanishes in directions tangent to it. This is a Neyman-orthogonality property. Under Assumption 9 the perturbation $\partial_f \mathbf{\Phi}$ induced by estimating f stays tangent to that manifold. The $\widehat{f}$-channel is then first-order negligible and $\mathbf{\Phi}$ may be treated as a fixed design in $\mathbf{\Sigma}$, carrying the CCP plug-in correction only. Absent that stability, $\partial_f \mathbf{\Phi}$ has an off-manifold component along which the score gradient need not vanish. A transition-induced influence term then enters $\mathbf{\Omega}(\mathbf{\Phi})$ (Appendix F.6).

The proof is in Appendix F.6.

Composing (33) with $\mathbf{\Sigma}(\cdot)$ gives the smooth (rational) variance–input mapping $\mathbf{q} \mapsto \mathbf{\Sigma}(\mathbf{\Phi}(\mathbf{q}))$. Asymptotic efficiency among GFD estimators is the problem of minimizing its trace.

Theorem 6 (Existence and FD-class trace optimality of $\mathbf{\Phi}^{\mathrm{opt}}$)**.** *Under the conditions of Theorem 5 and the compact regular sub-region $\mathcal{Q} \subseteq \mathbb{R}^{d_q}$ of Assumption 8, on which $\mathbf{J}(\mathbf{\Phi}(\mathbf{q}))$ is uniformly bounded away from singularity, the variance-optimal flow input*

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

exists and attains $\operatorname{tr}(\mathbf{\Sigma}(\mathbf{\Phi}^{\mathrm{opt}})) = \min_{\mathbf{\Phi}\in\mathcal{F}_{\mathcal{Q}}} \operatorname{tr}(\mathbf{\Sigma}(\mathbf{\Phi}))$. Here $\mathcal{F}_{\mathcal{Q}} := \{\mathbf{\Phi}^{\mathrm{MP}} + \mathbf{N}\mathbf{q} : \mathbf{q} \in \mathcal{Q}\} \subseteq \mathcal{F}(\cdot; f)$ is the image of the coordinate region $\mathcal{Q}$ under $\mathbf{q} \mapsto \mathbf{\Phi}(\mathbf{q})$. The minimization is over feasible flow inputs, not over a literal intersection of a flow set with a coordinate set. The result is the existence and attainment of a trace-minimizing flow input over the compact feasible sub-region $\mathcal{F}_{\mathcal{Q}}$; it does not assert optimality over the full FD-feasible affine variety or a semiparametric efficiency bound. Beyond the hypotheses of Theorem 5, FD-feasibility, and the regularity of $\mathcal{Q}$ (Assumption 8), the trace-minimization itself imposes no additional controllability, rank, or instrument-Jacobian conditions.

The proof is in Appendix F.7.

Sample-feasible procedure. The sample variant of (35) replaces f with $\widehat{f}$ and uses a plug-in sandwich $\widehat{\mathbf{\Sigma}}$. It is evaluated at a preliminary $\widehat{\theta}^{(0)} := \widehat{\theta}^{\mathrm{GFD}}(\widehat{\mathbf{\Phi}}^{\mathrm{MP}})$. Quasi-Newton minimization in $\mathbf{q}$ then follows, with a single re-run of Step 3 of Section 3.5 at the optimum. The trace map is rational but not generally convex in $\mathbf{q}$; multi-start is recommended with the Moore–Penrose anchor $\mathbf{q} = \mathbf{0}$ as a natural seed.

The canonical input is already near-optimal. The gain from this optimization is small in the numerical designs we study. We check this by Monte Carlo on the two worked DGPs of Section 4: the Type-K investment model ($\rho = 1$) and the Type-D entry/exit game ($\rho = 2$, with genuinely signed weights). In both, the canonical input $\mathbf{\Phi}^{\mathrm{MP}}$ is close to the within-class root mean squared error (RMSE) minimum over the searched null-space directions. The reason is structural: the minimum-ℓ_2-norm input keeps the signed flow weights small, which minimizes the first-stage plug-in contribution to $\mathbf{\Omega}$. The program (35) formalizes the variance-optimal input, while the canonical input is a stable baseline in these designs.

From within-class optimality to the semiparametric bound. Theorem 6 is optimality within the GFD class. Let $\mathcal{E}_\rho$ denote the class of regular two-step estimators of θ_0 that use only the

joint law of the ρ-window $(x_{it}, d_{it}, x_{i,t+1}, \ldots, x_{i,t+\rho})$. The first stage $(\widehat{f}, \widehat{\mathbf{p}})$ is nonparametric. Every GFD estimator lies in $\mathcal{E}_\rho$, but so do other ρ-horizon CCP estimators. Write $\mathbf{\Sigma}_\rho^{\text{eff}}$ for the Chamberlain (1987) semiparametric efficiency bound of $\mathcal{E}_\rho$. Whether $\mathbf{\Phi}^{\text{opt}}$ attains it depends on four sufficient conditions at the DGP:

(G1) Full controllability: for every (x_0, d, d'), $d \neq d'$, the pair-specific controllability matrix has full rank $S - 1$ (the achievable terminal-distribution perturbations span the simplex tangent). A linear-algebraic rank check on f.

(G2) Instrument-Jacobian rank: $\mathbf{q} \mapsto \widetilde{\mathbf{H}}(\cdot; \mathbf{\Phi}(\mathbf{q}))$ has Jacobian of full row rank d_θ at $\mathbf{q}^{\text{opt}}$.

(G3) FD-coincidence: some $\mathbf{\Phi}^\dagger$ in the fixed feasible sub-region $\mathcal{F}_Q$ of Theorem 6 (equivalently, with null-space coordinate $\mathbf{q}^\dagger$ interior to Q) realizes the efficient ρ-horizon influence function, $\mathbf{\Sigma}(\mathbf{\Phi}^\dagger) = \mathbf{\Sigma}_\rho^{\text{eff}}$. This is the load-bearing alignment condition; it is not implied by (G1)–(G2).

(G4) Oracle equivalence: the first-stage plug-in correction in (34) does not inflate the variance above its known-first-stage value (automatic for cell-frequency CCPs at the population point under exact FD; a separate check for sieve/kernel first stages).

Corollary 7 (Conditional ρ-horizon semiparametric efficiency)**.** *Under the conditions of Theorem 5 and gates (G1)–(G4) at the DGP, $\mathbf{\Sigma}(\mathbf{\Phi}^{\text{opt}}) = \mathbf{\Sigma}_\rho^{\text{eff}}$: the variance-optimal GFD estimator attains the ρ-horizon semiparametric efficiency bound.*

The proof is in Appendix F.8. The result is conditional. Within-class trace optimality (Theorem 6) does not require gates (G1)–(G4), but equality with the broader Chamberlain bound does.

Why the bound is conditional and not generic. The gates mark a genuine restriction. GFD replaces the infinite-horizon value difference with a finite-horizon flow weighting. It therefore does not use all cross-equation Bellman restrictions that a full-information estimator can exploit. Gate (G3) is the alignment condition under which the efficient ρ-horizon influence function is realized by a feasible flow input. When that alignment fails, the remaining gap is the sampling variance cost of working within the GFD class. Remark 5 quantifies it.

Remark 5 (Price of the signed-weight extension when the gates fail)**.** *When the gates (G1)–(G4) of Corollary 7 fail, the variance-optimal GFD estimator need not attain the full-information bound; any gap is the price of FD-feasibility (the affine-hull cost of Remark 1). We quantify this price by simulation. On the three-player Type-D entry/exit-game Monte Carlo (where $\rho^* = 2$ and the flow weights are genuinely signed), across* 80 *replications the trace of the GFD estimator's sampling covariance exceeds that of the full-solution NFXP maximum-likelihood estimator by a variance-inflation multiplier of* 1.24 *(per-parameter ratios* 1.09–1.29*). The signed-weight estimator thus forfeits roughly a quarter of the efficiency of the full Bellman solve (bounded and modest) in exchange for eliminating the dynamic program. The loss shrinks toward unity in renewal/Type-R DGPs, where the weights remain non-negative (Table 2).*

3.8 Counterfactual Computation via Finite Dependence

The counterfactual implication is narrow. We focus on payoff-only experiments that change the structural payoff from θ_0 to θ_1 while holding the transition kernel f_t fixed. In this case the same FD input $\mathbf{\Phi}$ remains valid. The only new object is the counterfactual CCP profile $\mathbf{p}_1 := \{\mathbf{p}_{1,t}\}_t$, where $\mathbf{p}_{1,t}(x_t) := \mathbf{p}_t(x_t; \theta_1, f)$. The contribution is therefore to carry the existing finite-dependence logic into non-stationary payoff-only counterfactuals. It is not a new identification result for transition-changing policies.

Theorem 8 (GFD counterfactual identity: payoff-only counterfactual CCPs from finite dependence)**.** *Suppose Assumptions 1–5 hold and that $\mathbf{\Phi}$ is FD-feasible at horizon ρ under f (Definition 2). Then $\mathbf{p}_{1,t}$ satisfies the fixed-point equation*

$$p_{1,t}(d_t \mid x_t) \;=\; \Lambda_{d_t}\big(x_t;\, \theta_1, \mathbf{\Phi}, \mathbf{p}_1\big), \qquad \forall\, (x_t, d_t) \in \mathcal{X} \times \mathcal{D}, \tag{36}$$

where Λ_{d_t} is the GFD multinomial-logit map (25) *evaluated at the counterfactual payoff. Moreover, (i) $\mathbf{p}_{1,t}$ exists and is unique, inherited from the Bellman contraction at θ_1 via the Hotz–Miller bijection, independently of any iteration; and (ii) the FD iteration $\mathbf{p}^{(m+1)} = \Lambda(\cdot;\, \theta_1, \mathbf{\Phi}, \mathbf{p}^{(m)})$ converges locally to $\mathbf{p}_{1,t}$ whenever the spectral radius of the composite map's Jacobian at $\mathbf{p}_{1,t}$ is below unity.*

The proof and full convergence analysis are in Appendix F.5. The theorem should be read as a reuse result. Once FD-feasibility is established for f_t, the same finite-horizon representation can be evaluated at the counterfactual payoff path. This is the sense in which the paper expands the counterfactual reach of finite dependence to non-stationary models.

Multi-agent extension. The same identity extends to payoff-only counterfactuals in dynamic games, provided the counterfactual equilibrium is well-posed. We state this separately. Unlike the single-agent case, existence and uniqueness of the counterfactual equilibrium are not guaranteed by a Bellman contraction. They must be assumed, as is standard in the empirical-games literature.

Assumption 10 (Counterfactual MPE)**.** *At the counterfactual payoff θ_1 the game admits a Markov-perfect equilibrium (MPE) with CCP profile $\mathbf{p}_1 = \{\mathbf{p}_{1,i}\}_i$, and the analyst selects one such profile by a maintained equilibrium-selection rule.*

Proposition 9 (Payoff-only counterfactual in games)**.** *Let each player i face the opponent-integrated transition $f_i^{\mathbf{p}_{-i}}$ obtained by integrating out rivals' CCPs, and suppose Assumptions 1–5 hold for each player's projected single-agent problem with $\mathbf{\Phi}_i$ FD-feasible for $f_i^{\mathbf{p}_{-i,1}}$. Under Assumption 10, the selected counterfactual profile $\mathbf{p}_1$ solves the per-player GFD fixed point evaluated at the equilibrium perceived transitions:*

$$p_{1,i,t}(d \mid x) \;=\; \Lambda_d^{(i)}\big(x;\, \theta_{1,i}, \mathbf{\Phi}_i, \mathbf{p}_1\big), \qquad \forall\, i,\ (x, d). \tag{37}$$

No Bellman solve is required at any candidate θ_1.

The proof (Appendix F.5) applies Theorem 4 player by player at the equilibrium perceived transition $f_i^{\mathbf{p}_{-i,1}}$; the per-player value-difference identity holds exactly under per-player FD, and assembling over players gives (37). Assumption 10 supplies the existence/selection that the single-agent contraction argument cannot.

This scope is complementary to existing identification results. Arcidiacono and Miller (2020) show that counterfactual CCPs in short panels are generally unidentified unless finite dependence holds. Kalouptsidi et al. (2021) give a null-space compatibility condition for transition-changing counterfactuals. Theorem 8 stays within the payoff-only class. Transition-changing experiments require a new feasibility check and the broader compatibility machinery of that literature.

3.9 Extension to Unobserved Heterogeneity

Finite dependence composes with a finite-mixture treatment of latent types. Because the existence test depends on the transition kernel alone, one flow input $\widehat{\mathbf{\Phi}}^{(\ell)}$ per type is precomputed once at $\rho = \max_\ell \rho^*_\ell$. The effective regressors are formed before the EM loop. The GFD logit log-likelihood is then averaged over posterior type weights in the E-step and maximized in the M-step. Each iteration costs a weighted logit with no Bellman re-solve. The construction extends to non-stationary mixtures and short panels. Theorem 12 gives consistency and asymptotic normality of the global maximizer under a Kasahara and Shimotsu (2009) mixture-identification condition. Corollary 11 records the per-type GFD identity. Appendix E gives the algorithm and a compact consistency statement.

4 Applications

This section verifies Theorem 3 on eight benchmark and stylized DDC examples (§4.1). It then benchmarks the GFD estimator's finite-sample behavior against four comparators (§4.2): the Hotz–Miller (HM) CCP estimator, nested pseudo-likelihood (NPL), the mathematical program with equilibrium constraints (MPEC), and the nested fixed-point algorithm (NFXP).

4.1 Existence Test on Benchmark and Stylized DDC Examples

The existence test of Theorem 3 certifies the finite-dependence horizon from the transition structure alone. When it certifies finite dependence, the value-difference reconstruction has residuals at numerical precision in the examples below. We organize the verification on eight examples by a four-type taxonomy of transition structures: renewal (Type R) with a rank-one reference action; shift register (Type S) with a p-lag register; Kronecker (Type K) with action-invariant exogenous factor; and the residual triple-connector / strategic ripple class (Type D). For the first three types the residual $\mathbf{b} \in \text{Range}(\mathbf{A})$ of Theorem 3 holds at a structurally transparent horizon, giving ρ^* in closed form; we verify each below directly from the transition structure (the Kronecker case via Lemma 10). Type D is dispatched by direct invocation of the existence test.

Type R (renewal). An action $d_t = 0$ induces a rank-one transition matrix $\mathbf{F}_0$. The residual $\mathbf{b}$ at $\rho = 1$ trivially lies in Range($\mathbf{A}$) at every triple $(x_t, d_t, 0)$, so $\rho^* = 1$ via the classical non-negative renewal solution of Arcidiacono and Miller (2011).

Type S (shift register, p-lag). The state factorizes into a p-component lag register of past actions plus an action-independent exogenous component. At $\rho = p$ the flow matrices $\Phi(x_t, d_t)$

and $\Phi(x_t, d'_t)$ that replay the same future action sequence produce identical terminal registers, so $\rho^* = p$ at every triple.

Type K (Kronecker with action-invariant exogenous factor). The transition factorizes as $\mathbf{F}_{d_t} = \mathbf{F}^{\text{endo}}_{d_t} \otimes \mathbf{F}^{\text{exo}}$ with $\mathbf{F}^{\text{exo}}$ action-invariant. The exogenous factor cancels identically between any two compared actions, so joint-state FD is equivalent to FD on the x^{endo}-projection and $\rho^*_{\text{joint}} = \rho^*_{\text{endo}}$ regardless of the persistence of $\mathbf{F}^{\text{exo}}$ (Lemma 10 in Appendix B). The structured GFD solver of Appendix B runs the flow QP on the M-dimensional endogenous projection only, yielding wall-clock savings that grow with the exogenous grid dimension Z: the structured solve avoids the Z^ρ blow-up of the joint state space (Table 10 in the appendix, shown at $\rho = 1$).

Type D (triple connector / strategic ripple). All-endogenous single-agent transitions or shared-action games have no closed-form ρ^* and generically require signed flow inputs (Remark 1). The operational route forms $(\mathbf{A}, \mathbf{b})$ at the candidate horizon, tests $\mathbf{b} \in \text{Range}(\mathbf{A})$ by sparse LSQR, and reads off ρ^* as the smallest passing horizon.

Remark 6 (Pointwise vs. universal FD in Type-D models)**.** *Single-agent Type-D models can have pointwise FD at a lower ρ than universal FD whenever the residual* $\mathbf{b}$ *falls in a low-dimensional subspace of* Range($\mathbf{A}$)*. Multi-player Type-D games carry two horizon notions. The "verified ρ^*" column of Table 1 reports the operative one, $\rho^* = 2$, not the minimal LSQR-feasible horizon. A per-player projected $\rho = 1$ certificate can pass the linear test. That certificate reflects the projected single-agent problem after integrating out opponents' equilibrium CCPs at the baseline profile. The operative horizon is set by the strategic ripple. Player i's action moves $t+1$ states. Opponents' best responses at $t+1$ move i's perceived transition at $t+2$. The continuation cancels only with terminal matching at $t+3$, so the ripple operates at $\rho = 2$. This $\rho^* = 2$ is an empirical finding verified for the two cited examples (entry/exit game, dynamic oligopoly), not a theorem for arbitrary games. Richer interaction topologies, larger network propagation, or deeper lag structure can raise the operative horizon. It must then be re-verified directly via Theorem 3.*

Results. Table 1 reports the predicted and LSQR-verified ρ^* together with the worst-case LSQR residual at the certifying horizon. Some rows are simplified versions of published dynamic discrete choice settings; others are stylized worked examples included to cover a transition class. All are implemented as `code/example_*.py` in the replication package.

Findings. Predicted and verified ρ^* coincide for every model (Table 1). The Type-D horizon convention (Remark 6) is reused in the Monte Carlo below.

4.2 Estimation Monte Carlo

We next verify the GFD estimator's finite-sample behavior against standard benchmarks. The estimator recovers the true structural parameters, with bias and root mean squared error (RMSE) competitive with the benchmark estimators. We use two DGPs: the single-agent Type-K investment model ($|\mathcal{X}| = 20$, $\rho^* = 1$) and the three-player Type-D stationary entry/exit game ($|\mathcal{X}_i| = 40$, $\rho^* = 2$). Both experiments share $N = 2{,}000$, $T = 20$, and the same first-stage CCP and transition estimates across estimators. Scripts `code/monte_carlo_investment.py` and `code/monte_carlo_game.py` accompany the replication package.

Model	Anchor / role	Type	S	D	ρ^* pred.	ρ^* verif.	LSQR resid.
Engine Replacement	Rust (1987)	R	10	2	1	1	8.7×10^{-14}
Job Search	stylized worked example	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}
Altug–Miller (p=3)	Altug and Miller (1998)	S	8	2	3	3	1.9×10^{-13}
Investment	stylized Type-K example	K	20	3	1	1	7.9×10^{-14}
Inventory Control	stylized Type-K example	K	20	3	1	1	7.9×10^{-14}
Entry/Exit Game	Ericson and Pakes (1995)	D	24	2	2	2‡	9.7×10^{-17}
Dynamic Oligopoly	Ericson and Pakes (1995)	D	16	2	2	2‡	1.9×10^{-13}

Table 1: Verification of the eight benchmark and stylized examples. S is the number of states and D the number of actions. The second column reports either the published benchmark that anchors the example or the role of a stylized example constructed in this paper. Type codes: R = renewal; S = shift register; K = Kronecker with action-invariant exogenous factor; D = triple connector / strategic ripple. Predicted ρ^* comes from the type-specific closed forms of §4.1 (R, S, K), or direct verification at $\rho = 1$ followed by $\rho = 2$ for Type D. Verified ρ^* is the smallest horizon at which $\mathbf{A}^{(x_t,d_t,d'_t)}\,\mathbf{x} = \mathbf{b}^{(x_t,d_t,d'_t)}$ is feasible at every triple. It is computed by sparse LSQR. ‡ The entry/exit game and the dynamic oligopoly are multi-player Markov-perfect models tabulated at their operative horizon $\rho^* = 2$. The per-player projected $\rho = 1$ certificate is a projected single-agent certificate, not the operative strategic horizon (Remark 6).

Single-agent Type-K investment. The state (k_t, o_t) factorises as capital $k_t \in \{0, \ldots, 4\}$ and productivity o_t on a 4-state Tauchen grid (Tauchen, 1986), with $|\mathcal{X}| = 20$. Three actions shift capital deterministically by $-1, 0, +1$. Productivity follows an action-invariant AR(1). This gives Type-K structure with $\rho^* = 1$ via Lemma 10. Flow payoff $u_t(x_t, d_t; \theta) = \theta^{\text{rev}} o_t \sqrt{k_t} - \theta^{\text{cost}} d_t - \theta^{\text{adj}} d_t^2$ is linear in $\theta_0 = (2.5, 0.3, 0.1)$. The parameters are calibrated for steady-state action shares $(47\%, 32\%, 21\%)$. The discount factor is $\beta = 0.95$. We compare GFD at the canonical Moore–Penrose flow input against Hotz–Miller CCP-2step, MPEC (Su and Judd, 2012), and the NFXP benchmark.

	$\theta^{\text{rev}} = 2.5$		$\theta^{\text{cost}} = 0.3$		$\theta^{\text{adj}} = 0.1$		
Method	Bias	RMSE	Bias	RMSE	Bias	RMSE	Time(s)
GFD	−0.003	0.044	0.001	0.023	−0.001	0.015	0.018
CCP-2step	0.002	0.043	0.000	0.023	−0.000	0.012	0.061
NFXP	0.002	0.043	0.000	0.023	−0.000	0.012	5.208
MPEC	0.002	0.043	0.001	0.023	−0.000	0.012	0.905

Table 2: Investment model Monte Carlo: 500 replications, $N = 2{,}000$, $T = 20$, Apple M3 Ultra. GFD uses Laplace-smoothed cell-frequency CCPs ($\alpha = 0.1$) for the first stage.

GFD matches CCP-2step, MPEC, and NFXP to ~ 0.001 in θ-units on θ^{rev} and θ^{cost}. A small residual efficiency difference remains on θ^{adj} (RMSE 0.015 vs 0.012), reflecting the GFD moment's particular weighting of the ρ-window rather than a suboptimal flow input: the canonical min-norm input already sits at the within-class trace minimum of Theorem 6, so re-optimizing $\mathbf{\Phi}$ over the FD-feasible class does not close it. Wall-clock: GFD is 3.4× faster than CCP-2step, 50× faster than MPEC, and 290× faster than NFXP. We sweep the joint state space from $|\mathcal{X}| = 20$ to $|\mathcal{X}| = 5{,}000$ via the structured Kronecker solver of Appendix B. The GFD/NFXP wall-time ratio then rises from ∼17× to ∼136× (Table 11 in Appendix B).

Three-player Type-D stationary entry/exit game. Three firms $i = 1, 2, 3$ each choose $d_{i,t} \in \{0, 1\}$ (exit / enter-or-stay). The market-size index is $w_t \in \{2, 4, 6, 8, 10\}$, a 5-state persistent exogenous Markov chain. The incumbency profile is $(y_{1,t}, y_{2,t}, y_{3,t}) \in \{0, 1\}^3$. Per-player state $x_{i,t} = (w_t, y_{1,t}, y_{2,t}, y_{3,t})$ has $|\mathcal{X}_i| = 40$ values. Per-period payoffs are linear in $\theta_i = (\theta_i^{\text{rev}}, \theta_i^{\text{comp}}, \theta_i^{\text{fc}}, \theta_i^{\text{entry}})$:

$$u_i(x_{i,t}, d_{i,t} = 1; \theta_i) \;=\; \theta_i^{\text{rev}} \log(w_t) \;-\; \theta_i^{\text{comp}}\, \mathbb{E}\big[\log(1 + N_{-i,t}^{\text{enter}})\big] \;-\; \theta_i^{\text{fc}} \;-\; \theta_i^{\text{entry}}\,(1 - y_{i,t}), \quad (38)$$

with $u_i(\cdot, 0; \theta_i) = 0$, where $N_{-i,t}^{\text{enter}} = \sum_{j \neq i} \mathbf{1}\{d_{j,t} = 1\}$ is the number of active rivals. The expectation in the competition term is taken under opponents' equilibrium CCPs. $\beta = 0.9$; $\theta_{0,1} = (1, 1, 1, 1)$ for Player 1; the other two players use slightly perturbed entry costs (0.9 and 0.8). The stationary MPE is solved by Gauss–Seidel best-response iteration; under the resulting perceived transitions, FD holds at $\rho^* = 2$ (Remark 6). We simulate 200 replications with $N = 2{,}000$ firms over $T = 20$ periods, estimating Player 1's parameter vector while holding the other players' true θ fixed.

	Bias / RMSE per parameter (true $\theta_{0,1} = (1, 1, 1, 1)$)				
Method	θ^{rev}	θ^{comp}	θ^{fc}	θ^{entry}	Time (s)
Two-step CCP-based (no inner MPE solve per parameter trial):					
GFD ($\mathbf{\Phi}^{\text{MP}}$)	−0.005/0.048	−0.017/0.162	+0.002/0.046	−0.001/0.022	0.18
HM (Hotz and Miller, 1993)	−0.004/0.046	−0.013/0.153	−0.000/0.043	−0.001/0.022	0.05
Iterated CCP / Bellman methods (solve the MPE per parameter trial):					
NPL (Aguirregabiria and Mira, 2007)	−0.004/0.046	−0.013/0.153	+0.001/0.043	−0.001/0.022	0.13
NFXP†	+0.008/0.045	+0.016/0.162	+0.005/0.052	+0.001/0.020	5.61

Table 3: Three-player Markov-perfect entry/exit game. 200 replications (20 for NFXP†), $N = 2{,}000$, $T = 20$, $\rho = 2$, $|\mathcal{X}_i| = 40$, $\beta = 0.9$, $\theta_{0,1} = (1, 1, 1, 1)$. Bias and RMSE in θ-units. †200-rep NFXP would take ∼ 19 minutes.

All four estimators are essentially unbiased (|bias| ≤ 0.017). Component-wise, GFD's RMSE sits within ∼7% of HM/NPL and within ∼8% of NFXP on every θ-component. This confirms the FD-feasible $\widehat{\mathbf{\Phi}}$ consistency claim (Theorem 5). The two-step CCP methods (GFD 0.18s, HM 0.05s) sit in the same fast regime because neither re-solves the MPE per parameter trial. NPL (0.13s) solves the MPE iteratively in CCP space. NFXP (5.61s) solves it via Bellman value iteration. The 31× gap is between two-step CCP methods and NFXP, not between GFD

and the other CCP-based estimators. We therefore do not claim a GFD wall-clock advantage over HM/NPL on the multi-agent estimation step; the GFD identity's value at estimation is consistency, and its payoff is the Bellman-free counterfactual identity of Theorem 8, taken to data in §5.

The verification and simulations set up the empirical application, where GFD estimates a non-stationary game without imposing a final value-function path.

5 Empirical Application: Non-stationary US Airline Entry Game

Does the value of airline market entry vary over time, and does it matter if we assume otherwise? We answer with a dynamic entry game of ten U.S. carriers over 2011–2023. The panel spans two structural breaks: the 2008–2013 merger wave and the 2020 Coronavirus Disease 2019 (COVID-19) demand shock.

We build the game on a structure close to the empirical airline-entry literature. A nested-logit demand model and Nash–Bertrand pricing deliver each carrier's variable profit in dollars. The demand estimates match Aguirregabiria and Ho (2012). We estimate a price coefficient of -1.03 per \$100 (their -1.37) and a hub coefficient of $+0.028$ (their 0.027–0.036). GFD then estimates the dynamic entry game on top. It uses one regression per carrier, with no game-theoretic fixed point and no value-function solve.

Two findings follow. First, entry values are non-stationary. A Wald test rejects stationarity for eight of the ten carriers (Table 5). Second, the non-stationarity is modest once we control for dollar profitability. The variable-profit term absorbs the COVID demand collapse. The residual entry-value movements are small, and imposing stationarity distorts predicted entry by under two percentage points on average. The large raw swings of a profit-blind specification, such as a Southwest collapse at COVID, are a profitability effect rather than a structural shift in entry value.

This setting is useful for two reasons. On computation, the GFD value-difference identity of Theorem 4 turns a thirteen-year, ten-agent, non-stationary game into one logit per carrier. On interpretation, the dollar profit control separates demand shocks from structural entry-value shifts. Once we make that separation, the apparent non-stationarity shrinks. What survives is strategic network repositioning, clearest in the 2021 recovery.

5.1 Setup

Industry background. Mergers between 2008 and 2013 consolidated the U.S. network carriers. Delta absorbed Northwest in 2008. United merged with Continental in 2010. Southwest acquired AirTran in 2011. American merged with US Airways in 2013. The COVID-19 demand shock in 2020 cut traffic on many routes. Recovery through 2023 was uneven. We group ticketing carriers into ten corporate entities and treat them as the strategic players: the four legacy and low-cost majors (DL, AA, UA, WN) and six others (Alaska, JetBlue, Spirit, Frontier, Allegiant, Hawaiian).

Structural model. At the start of year t, carrier i either serves market m ($a_{i,t} = 1$) or not. Markets are non-directional city-pairs. The per-period payoff, in dollars, is

$$u_i(a_{i,t}, x_{m,t}, t) \;=\; a_{i,t}[\, R_{i,m,t} \;-\; \mathrm{FC}_{i,m,t} \;-\; (1 - a_{i,t-1})\,\mathrm{EC}_{i,m,t}] \;+\; \varepsilon_{i,t}(a_{i,t}), \tag{39}$$

with ε independently and identically distributed Type-I extreme value and outside option $u_i(0, \cdot) = 0$. Here $R_{i,m,t}$ is the carrier's variable profit, recovered in dollars from the demand and pricing model of §5.2. The fixed cost $\mathrm{FC}_{i,m,t} = \gamma_{i,t} + g_2 \,\mathrm{hub}_{i,m,t} + g_3 \,\mathrm{dist}_m$ and the entry cost $\mathrm{EC}_{i,m,t} = \eta_1 + \eta_2 \,\mathrm{hub}_{i,m,t} + \eta_3 \,\mathrm{dist}_m$ follow Aguirregabiria and Ho (2012). The hub variable is the carrier's number of nonstop connections at the two endpoint metros. The discount factor is $\beta = 0.95$. The time-varying primitive is the carrier-year intercept $\gamma_{i,t}$, the entry value net of variable profit. A lower $\gamma_{i,t}$ raises the fixed cost and lowers the value of serving.

Data. We build the panel from the Bureau of Transportation Statistics (BTS) DB1B Market itinerary file, 2011–2023. We define markets on city-market codes. We measure market size as the geometric mean of the two endpoint metro populations from Census Core-Based Statistical Area (CBSA) estimates. After grouping carriers and dropping the missing-lag year, the estimation panel holds 1,368,275 carrier-market-year cells over 2012–2022. Across the panel, carriers serve 33.8% of cells. Non-incumbents enter in 2.6% of cells, and incumbents exit in 4.9% of cells. The constructed dynamic panel includes 2023. We use that terminal year only to form the 2022 renewal offset, because the offset needs next-year CCPs. We do not use it as an estimating year. The replication package records this convention in a saved audit file. The file maps each reported number to the corresponding script output.

5.2 The dollar structure and its literature anchor

The dollar variable profit $R_{i,m,t}$ comes from a standard demand and supply model, estimated once and held fixed in the dynamic stage. We estimate nested-logit demand (one nest for air travel) on the product-level panel of carrier-market-quarter fares and shares. Fare is endogenous. We instrument it with a Hausman shifter, the carrier's mean fare in its other markets, and Berry–Levinsohn–Pakes (BLP) shifters, the count and characteristics of competing products. The first stage is strong, with F above 10^4. Nash–Bertrand pricing then recovers each carrier's marginal cost from its first-order condition. The variable profit is $R = (\text{fare} - \text{marginal cost}) \times \text{passengers}$.

Table 4 reports the demand estimates against Aguirregabiria and Ho (2012). The price and hub coefficients have the same sign and comparable magnitudes. The implied markups (\$58–\$97, median \$68) leave positive marginal costs throughout. The mean variable profit per served carrier-market is \$630K per year, the same order as the benchmark's fixed-cost estimates. This dollar structure is the literature anchor for the dynamic stage.

5.3 Identification via the GFD value-difference identity

Finite dependence in the entry game. Each carrier has a renewal action. Not serving ($a_{i,t} = 0$) resets its own incumbency to zero regardless of the current state. Finite dependence therefore holds at $\rho = 1$ along the comparison path $a_i = 0$. We work with the per-player perceived transition. We integrate out rivals' estimated CCPs $\widehat{p}_{-i}$. Identification is conditional on these

Table 4: Demand estimates against the structural benchmark. Nested logit, two-stage least squares; standard errors in parentheses. AH2012 = Aguirregabiria and Ho (2012).

Parameter	This paper	AH2012
Price (per $100)	−1.028 (0.005)	−1.366
Nesting σ	0.411 (0.002)	(nested logit)
Hub (demand)	0.028 (0.000)	0.027–0.036
Distance (/1000mi)	0.555 (0.004)	0.228

sieve CCPs, which we hold fixed at their data-anchored values. We do not re-solve the Markov-perfect equilibrium. Rivals enter through two data-anchored channels. The first is the dollar variable profit R, which a rival's presence lowers through Nash–Bertrand competition. The second is the continuation offset below.

For each carrier i and year t, the value-difference identity of Theorem 4 gives a per-period moment condition. The log odds of serving equal the dollar payoff difference plus a known continuation offset:

$$\log \frac{\widehat{p}_i(1 \mid x_{m,t}, t)}{\widehat{p}_i(0 \mid x_{m,t}, t)} \;=\; \theta[R_{i,m,t} - \mathrm{FC}_{i,m,t} - (1 - a_{i,t-1})\,\mathrm{EC}_{i,m,t}] \;+\; b_{i,m,t}, \tag{40}$$

where $\theta = 1/\sigma_\varepsilon$ converts dollars to the choice scale. The renewal offset $b_{i,m,t} = \beta\,[\log \widehat{p}_i(0 \mid 0, t{+}1) - \log \widehat{p}_i(0 \mid 1, t{+}1)]$ is the continuation gain from incumbency. It depends only on $\widehat{p}$ and β, not on the structural parameters. Estimation of the fixed-cost and entry-cost parameters, including the year intercepts $\gamma_{i,t}$, therefore reduces to a logit regression with R as a known dollar regressor and $b_{i,m,t}$ as a known offset.

Scope: what GFD identifies, and what it does not. GFD identifies the structural payoff given the scale σ_ε. It does not pin the scale itself. The dollar variable profit R has little variation orthogonal to the entry barriers: large markets carry high R but also attract more competitors, so R does not separate entry from non-entry after we control for incumbency. The coefficient $\theta = 1/\sigma_\varepsilon$ is therefore weakly identified, and we do not report dollar fixed-cost or entry-cost levels. The scale of the shock is an intertemporal object, pinned by the trade-off between the flow payoff and the discounted value stream, which the GFD value-difference identity cancels. In the point estimate, the implied shock scale is about $18.8 million, which is too large to read as a credible dollar normalization. We accordingly report the entry-value dynamics $\gamma_{i,t}$ in the identified (normalized) units, and keep the dollar magnitudes on the variable-profit side, where they are point-identified.

5.4 Estimates

Is stationarity appropriate? Under the model, the entry-value primitives are time-invariant if and only if $\gamma_{i,t}$ is constant across years. We estimate the unrestricted model with one intercept for each carrier-year. For each carrier, we then test whether these intercepts equal the 2012 intercept. Let $\widehat{\theta}$ denote the unrestricted estimate and let $\widehat{V}$ denote its inverse-Hessian covariance matrix. We form a matrix R_i whose rows contain the restrictions $\gamma_{i,t} - \gamma_{i,2012} = 0$ for $t = 2013, \ldots, 2022$.

We report

$$W_i = (R_i\widehat{\theta})'(R_i\widehat{V}R_i')^{-1}(R_i\widehat{\theta}).$$

Under the null, W_i has an asymptotic $\chi^2(10)$ distribution. The ten degrees of freedom come from the ten year-to-2012 differences in the 2012–2022 estimation sample. The covariance matrix is conditional on the recovered variable-profit control and the first-stage CCP used in the renewal offset. Table 5 reports the result. The test rejects stationarity for eight of the ten carriers. The two exceptions are Allegiant and Hawaiian. The rejections are sharpest for the legacy carriers United, Delta, and American. Their networks moved most across the merger wave and the pandemic.

Table 5: Wald test of entry-value stationarity: $H_0\colon \gamma_{i,t} = \gamma_i \;\forall\, t$ (10 restrictions per carrier).

Carrier	$\chi^2(10)$	p	Carrier	$\chi^2(10)$	p
UA	3071.8	< 0.001	WN	213.1	< 0.001
F9	1375.9	< 0.001	AS	412.6	< 0.001
DL	1083.4	< 0.001	NK	30.9	< 0.001
AA	447.6	< 0.001	B6	30.4	< 0.001
			G4	8.9	0.54
			HA	0.2	1.00

Carrier codes: AA = American, AS = Alaska, B6 = JetBlue, DL = Delta, F9 = Frontier, G4 = Allegiant, HA = Hawaiian, NK = Spirit, UA = United, WN = Southwest.

The entry-value paths. Figure 1 plots the estimated $\widehat{\gamma}_{i,t}$ for the four majors. The paths move, but modestly. The full carrier-year range is about one unit for Delta and Southwest, and about two for United. This is an order of magnitude smaller than a profit-blind intercept would show. The COVID year illustrates the point. Net of the variable-profit term, Southwest's entry value barely moves in 2020 (+0.15), while United and American decline by 0.47 and 0.41. A profit-blind specification would instead load the entire COVID demand collapse onto the intercept, producing a large Southwest drop in the profit-blind estimates. The variable-profit term absorbs that collapse. The residual entry-value path is what remains, and it is small. The paths are not driven solely by the variable-profit control: dropping R leaves the carrier-year intercepts highly correlated with the baseline estimates (correlation 0.81).

What the residual non-stationarity captures. The surviving movements are not the demand cycle. The variable-profit term absorbs that. They track strategic network repositioning. These are footprint and fixed-cost decisions, made for reasons beyond current profit. The 2021 recovery shows the largest and most divergent movements. United and Southwest raise their entry values (+1.31 and +0.75), in line with their aggressive post-pandemic expansion. Delta lowers its value (-0.77), in line with a more conservative rebuild. American's value stays elevated through 2014–2019, tracking its post-merger network integration. A stationary model pools these divergent bets into one carrier intercept. It would read strategic repositioning as cross-sectional market heterogeneity instead.

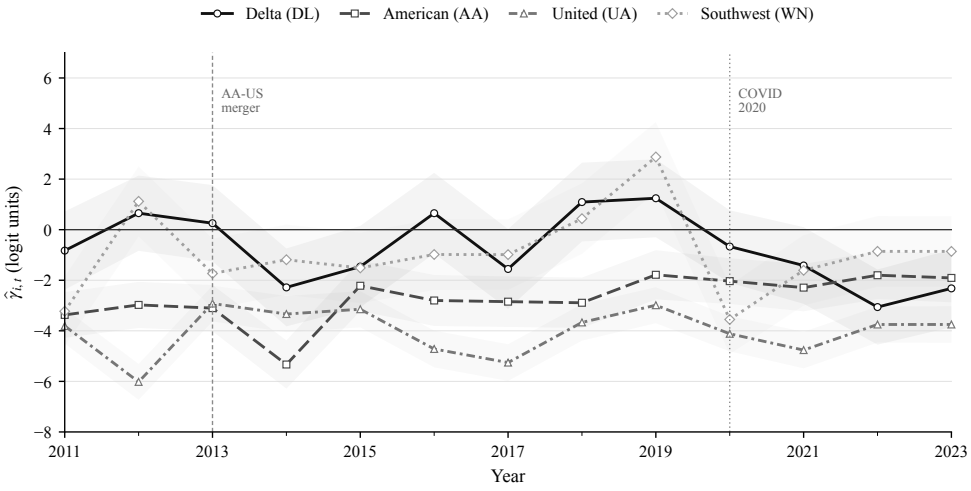

Figure 1: Estimated entry-value paths $\widehat{\gamma}_{i,t}$ (normalized units) for the four majors, 2012–2022. Dashed vertical: 2013 AA–US Airways merger. Dotted vertical: 2020 COVID demand shock.

5.5 The cost of imposing stationarity

How much does the non-stationarity matter for predictions? For each cell we form two predicted entry probabilities. The first uses the full $\widehat{\gamma}_{i,t}$. The second pools $\gamma_{i,t}$ to the carrier time-average $\bar{\gamma}_i$, holding the other primitives and the offset fixed. The gap isolates the entry-value channel. After we control for variable profit, this cost is modest. The mean absolute gap is 0.8 percentage points, and the 90th percentile is 1.9 points. The largest carrier-year distortions reach about 6 points (United in 2021). This is smaller than the gap a profit-blind specification produces, because most of the year-to-year movement in raw entry rates is profitability, which the dollar structure already captures.

This finding refines the case for non-stationarity rather than overturning it. The statistical rejection is decisive. The economic content is concentrated. It lives in the few carrier-years where the structural entry value genuinely shifts net of profit.

5.6 Replicating AH's hubbing counterfactual

Aguirregabiria and Ho (2012) use their estimated model to ask why airlines form hub-and-spoke networks. Their Table 11 reports the two-hub concentration ratio CR-2. This is the share of a carrier's nonstop connections involving its two largest hubs. We first replicate their empirical object, before comparing it with our GFD implementation. We rebuild their 2004 network sample from DB1B Market files using their 55-city and 22-carrier definitions. The rebuilt sample recovers the first hubs for the seven largest carriers and matches the published CR-2 ratios closely, with a mean absolute gap of 3.1 percentage points for those carriers. Exact agreement would require the ticket-level filters and GAUSS equilibrium objects behind their NPL implementation. We do not use those objects.

The large AH channel comes from the entry-cost hub coefficient. Table 6 reports the relevant cost parameters from their Table 9. To put magnitudes on a common scale, we also convert the

AH dollar-cost coefficients into payoff-index units. The conversion divides by their estimated shock scale σ_ε and changes the sign, since a lower cost raises the payoff from serving. The standard errors use a delta-method approximation based on the reported standard error of σ_ε.

Table 6: Cost parameters behind the AH hubbing channel

Parameter	AH2012, converted (s.e.)	This paper (s.e.)
Fixed-cost hub effect	0.122 (0.030)	0.012 (0.001)
Fixed-cost distance effect	−0.482 (0.088)	−0.015 (0.009)
Entry-cost hub effect	1.102 (0.182)	−0.032 (0.001)
Entry-cost distance effect	−0.001 (0.001)	0.319 (0.012)

Notes: Entries are effects on the normalized payoff index for serving, so the two columns are in common logit-index units. AH2012 dollar-cost coefficients are converted as $-\widehat{c}/\widehat{\sigma}_\varepsilon$, with $\widehat{\sigma}_\varepsilon = 8.402$. Delta-method standard errors ignore covariances with $\widehat{\sigma}_\varepsilon$, which AH2012 does not report.

Table 7 records the AH counterfactual benchmark for the seven carriers reported in their Table 11. Eliminating the entry-cost hub effect lowers CR-2 by 24 to 42 percentage points for the legacy carriers, under both AH implementation methods. This replicates the qualitative AH result: the entry-cost hub channel is the main force behind their hubbing counterfactual.

Table 7: AH2012 hubbing benchmark: CR-2 concentration ratios

Carrier	Observed	Method I No EC-hub	Method II No EC-hub	Method II change
WN	18.2	8.9	8.3	-9.9
AA	42.0	17.6	16.6	-25.4
UA	45.7	17.8	15.7	-30.0
DL	48.0	18.7	17.9	-30.1
CO	68.3	27.3	26.0	-42.3
NW	49.2	18.7	17.2	-32.0
US	45.3	18.1	16.4	-28.9

Notes: Method I is AH's Taylor approximation without strategic interactions. Method II iterates the counterfactual policy map starting from the Taylor approximation. The final column is Method II No EC-hub minus Observed, in percentage points.

We now compare methods. The GFD exercise in Table 8 is deliberately narrower than the original AH simulation. The observed CR-2 column is computed from BTS nonstop connections. The counterfactual columns use the GFD payoff index on carrier-market pairs that appear in the observed nonstop network at least once. We hold the baseline renewal offset and the observed hub path fixed. Thus the table is a payoff-channel analogue of AH Table 11, not a full endogenous network-transition simulation. Experiment 2 sets the hub-size coefficient in fixed costs to zero. Experiment 3 sets the hub-size coefficient in entry costs to zero.

Table 8: AH-style hubbing counterfactuals: CR-2 concentration ratios

Carrier	Observed	Baseline	No FC-hub	No EC-hub	EC-hub change
AA	22.7	9.2	9.1	9.2	0.0
AS	47.0	37.6	37.2	37.5	-0.1
B6	59.0	39.4	39.0	39.7	0.3
DL	36.5	17.4	17.1	17.4	0.0
F9	39.7	19.5	18.8	19.6	0.1
G4	37.9	40.3	39.5	40.3	0.0
HA	81.7	76.1	76.1	75.9	-0.2
NK	30.0	23.8	23.7	23.8	0.0
UA	21.3	6.4	6.3	6.4	0.0
WN	7.6	6.6	6.7	6.6	0.0

Notes: CR-2 is the share of nonstop connections involving a carrier's two largest hubs, averaged over 2012–2022. "No FC-hub" sets the fixed-cost hub coefficient to zero. "No EC-hub" sets the entry-cost hub coefficient to zero. The final column is No EC-hub minus Baseline, in percentage points.

The comparison is informative because the two counterfactual objects are different. AH changes the structural cost of route entry and then updates the equilibrium network. Our GFD exercise changes the payoff index on the observed support, holding the renewal offset and hub path fixed. In AH, eliminating the entry-cost hub effect substantially lowers CR-2 for the legacy carriers. In our dollar GFD specification, the EC-hub experiment moves CR-2 by less than one percentage point for every carrier. The fixed-cost hub experiment is similarly small. The baseline also underpredicts the observed CR-2 of American, Delta, and United. The gap therefore points to a network mechanism outside our current payoff object. The likely candidate is AH's fourth experiment: local managers internalize profits from one-stop routes that use a city-pair as a segment. Our current variable-profit control is a carrier-market profit object, so it does not reproduce that segment-complementarity mechanism.

5.7 What we learn about airline entry

Four points emerge. First, GFD scales. One logit per carrier estimates a ten-agent, thirteen-year, non-stationary game across two structural breaks, and recovers dollar structure comparable to the structural benchmark of Aguirregabiria and Ho (2012). Second, the raw non-stationarity of entry is largely a profitability effect. The dollar variable-profit term absorbs the demand shocks, and the residual structural entry value moves much less than a profit-blind intercept suggests. Third, what remains is strategic, not cyclical. The residual non-stationarity tracks divergent network bets. The clearest case is the 2021 recovery, when United and Southwest expanded while Delta held back. A stationary model would hide this repositioning, reading it as cross-sectional heterogeneity. Fourth, the AH replication clarifies what our method-comparison exercise is missing. AH's large hubbing response comes from an endogenous network simulation with segment-profit complementarity. Our current carrier-market payoff analogue does not include that mechanism, so it does not recover the large entry-cost hub channel.

Relation to the non-stationary estimation literature. Prior non-stationary dynamic game studies (Igami, 2017, 2018; Igami and Uetake, 2020; Doraszelski et al., 2018) solve the value function by backward induction at each candidate parameter; Sweeting (2013) combines backward induction with CCP elements. GFD replaces the value-function solve with one regression per carrier. The backward-induction computation reduces to computing the renewal offset once from the sieve CCPs. The cost is a boundary requirement at the last period, where the offset needs terminal-period CCPs or an explicit extrapolation rule. We drop the terminal year from estimation for this reason.

6 Conclusion

We develop a computable framework for finite dependence in dynamic discrete choice models. The existence test (Theorem 3) reduces ρ-period finite dependence to a pointwise linear feasibility test $\mathbf{Ax} = \mathbf{b}$ on the structural transitions. It is a pre-estimation diagnostic at every (x_t, d_t, d'_t) triple. The payoff-only counterfactual identity (Theorem 8) yields a Bellman-free fixed-point system for counterfactual CCPs, with restart-based validation against extraneous fixed points. FD-class trace optimality (Theorem 6) characterizes the variance-optimal flow input as the trace-minimizing solution over a compact regular sub-region. The fourth contribution is the supporting infrastructure. The GFD estimator (Definition 3) takes any FD-feasible flow input as an explicit argument, and the computational machinery of Appendix B makes the test and estimator operational at large state spaces.

Three limitations bound the present scope. First, the framework is developed for discrete state spaces. Extension to fully continuous state spaces via sieve approximation of the transition kernel remains open. Forming the feasibility test from a sieve estimate $\widehat{f}_{K_N}$ turns pointwise FD into an approximate condition. Its joint bias–variance rate (sieve error against signed-weight norm inflation) is the open theoretical question. Recent continuous-state DDC methods (Blevins, 2025; Adusumilli and Eckardt, 2025) are natural comparison points. By contrast, non-stationary, time-varying transitions stay within scope throughout. The existence test, the GFD identity, and the airline application all allow f_t to vary with t. Second, the GFD counterfactual identity (Theorem 8) 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 B keep this manageable for the eight examples verified in Section 4.1. Pushing ρ well beyond the values verified here, in dense-transition models, would require either a sparser reformulation or stronger structural restrictions.

Implications for applied work. The results suggest a simple sequence for empirical DDC work. First, one can run the pointwise feasibility test of Theorem 3 on the estimated transition kernel. In our examples, seven of eight pass at $\rho \leq 2$. When the test passes, GFD provides a two-step CCP estimator with the same observable inputs as other CCP methods and without repeated Bellman solves. Second, for payoff-only counterfactuals, the GFD counterfactual identity of Theorem 8 computes counterfactual CCPs without re-solving the dynamic program. The restart and spectral-radius diagnostics of Section 3.8 check the resulting fixed point. Third, when finite dependence admits multiple feasible flow inputs, Theorem 6 selects the trace-minimizing

member of the GFD class over the regular sub-region. The replication package implements these steps together with the eight-example verification, the Monte Carlo benchmarks, and the airline application.

Outlook. The pre-estimation feasibility test changes the role of finite dependence from a model-specific construction into a checkable property of the transition kernel. This matters in settings where large state spaces, non-stationarity, or strategic interaction make repeated Bellman solves costly. The examples in this paper cover entry/exit and dynamic oligopoly, multi-period labor histories, multi-component asset and demographic state vectors, and multi-player Markov-perfect games. The framework does not remove the need to verify finite dependence in each new application. It makes that verification an explicit linear-algebraic step.

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 C.-Y. Ho (2012). A dynamic oligopoly game of the us airline industry: Estimation and policy experiments. *Journal of Econometrics 168*(1), 156–173.

Aguirregabiria, V. and A. Magesan (2018). Solution and estimation of dynamic discrete choice structural models using euler equations. *Working Paper, University of Toronto*.

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., E. Aucejo, A. Maurel, and T. Ransom (2016). College attrition and the dynamics of information revelation. *Working paper*.

Arcidiacono, P. and P. B. Ellickson (2011). Practical methods for estimation of dynamic discrete choice models. *Annual Review of Economics 3*, 363–394.

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.

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.

Doraszelski, U., G. Lewis, and A. Pakes (2018). Just starting out: Learning and equilibrium in a new market. *American Economic Review 108*(3), 565–615.

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

Gayle, G.-L. and L. Golan (2012). Life-cycle fertility and human capital accumulation. *Review of Economic Studies 79*(1), 148–181.

Gayle, G.-L., A. Hincapie, and R. A. Miller (2018). Fertility and female labor supply. *Working paper*.

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

Hao, Y., H. Kasahara, and K. Shimotsu (2026). Semiparametric identification of the discount factor and payoff function in dynamic discrete choice models. 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.

Igami, M. (2017). Estimating the innovator's dilemma: Structural analysis of creative destruction in the hard disk drive industry, 1981–1998. *Journal of Political Economy 125*(3), 798–847.

Igami, M. (2018). Industry dynamics of offshoring: The case of hard disk drives. *American Economic Journal: Microeconomics 10*(1), 67–101.

Igami, M. and K. Uetake (2020). Mergers, innovation, and entry-exit dynamics: Consolidation of the hard disk drive industry, 1996–2016. *Review of Economic Studies 87*(6), 2672–2702.

James, J. (2014). Occupational choice and finite dependence. *Working paper*.

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.

McLachlan, G. J. and T. Krishnan (2008). *The EM Algorithm and Extensions* (2nd ed.). Wiley.

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.

Pesendorfer, M. and P. Schmidt-Dengler (2008). Asymptotic least squares estimators for dynamic games. *Review of Economic Studies 75*(3), 901–928.

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.

Su, C.-L. and K. L. Judd (2012). Constrained optimization approaches to estimation of structural models. *Econometrica 80*(5), 2213–2230.

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

Sweeting, A. (2013). Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry. *Econometrica 81*(5), 1763–1803.

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

Wu, C. F. J. (1983). On the convergence properties of the EM algorithm. *The Annals of Statistics 11*(1), 95–103.

A 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 the agent observes $x_t = l$ and chooses an initial action $d_t \in \mathcal{D}$. At time $t + 1$ the state $x_{t+1} \in \mathcal{X}$ is drawn from $f_t(\cdot \mid x_t, d_t)$ and the agent chooses $d_{t+1} \in \mathcal{D}$ according to the per-step weight $\omega_{t+1}(\cdot \mid x_t, x_{t+1}, d_t)$. The terminal state $x_{t+2} \in \mathcal{X}$ is drawn from $f_{t+1}(\cdot \mid x_{t+1}, d_{t+1})$; $\rho = 1$ means we evaluate the matching condition on the distribution of x_{t+2}. The QP compares two scenarios that differ only in the time-t action: in scenario A the agent picks d_t, in scenario B the agent picks d'_t. After time t the two scenarios share the same linear constraints, but produce distinct flow matrices because $f_t(\cdot \mid x_t, d_t)$ and $f_t(\cdot \mid x_t, d'_t)$ differ.

Flow matrix and initial-flow constraint. With $\rho = 1$ the only path index is (x_{t+1}, d_{t+1}), and the flow matrix is

$$\Phi(x_t, a) \;=\; \begin{pmatrix} \phi_{l,0}(x_t, a) & \phi_{l,1}(x_t, a) \\ \phi_{h,0}(x_t, a) & \phi_{h,1}(x_t, a) \end{pmatrix}, \qquad a \in \{d_t, d'_t\},\ x_t = l.$$

Since $\rho = 1$ there is no flow conservation, and the initial-flow constraint (10) reduces to

$$\begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} \phi_{l,0}(x_t, a) \\ \phi_{l,1}(x_t, a) \\ \phi_{h,0}(x_t, a) \\ \phi_{h,1}(x_t, a) \end{pmatrix} \;=\; \begin{pmatrix} f_t(l \mid x_t, a) \\ f_t(h \mid x_t, a) \end{pmatrix}.$$

The right-hand side is $\big(f_t(l \mid x_t, d_t),\ f_t(h \mid x_t, d_t)\big)^\top$ when $a = d_t$, and $\big(f_t(l \mid x_t, d'_t),\ f_t(h \mid x_t, d'_t)\big)^\top$ when $a = d'_t$. The constraint matrix is identical in the two scenarios.

Decomposition into the per-step weight at $\tau = 1$. The decomposition (9) collapses to

$$\phi_{x_{t+1},\, d_{t+1}}(x_t, d_t) \;=\; f_t(x_{t+1} \mid x_t, d_t) \cdot \omega_{t+1}(d_{t+1} \mid x_t, x_{t+1}, d_t), \qquad x_{t+1} \in \{l, h\},\ d_{t+1} \in \{0, 1\}, \tag{41}$$

where $x_t = l$ and $a \in \{d_t, d'_t\}$ (replacing d_t in the formulas above) is the scenario-specific time-t action. The row sums of $\Phi(x_t, d_t)$ recover the one-step arrival probabilities (the initial-flow constraint), and the within-row ratios $\phi_{x_{t+1},d_{t+1}}(x_t, d_t) / \big(\phi_{x_{t+1},0}(x_t, d_t) + \phi_{x_{t+1},1}(x_t, d_t)\big) = \omega_{t+1}(d_{t+1} \mid x_t, x_{t+1}, d_t)$ recover the time-1 per-step weight.

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

$$\begin{pmatrix} \kappa_{t+2}(l \mid x_t, d_t) \\ \kappa_{t+2}(h \mid x_t, d_t) \end{pmatrix} = \begin{pmatrix} f_{t+1}(l \mid l, 0) & f_{t+1}(l \mid l, 1) & f_{t+1}(l \mid h, 0) & f_{t+1}(l \mid h, 1) \\ f_{t+1}(h \mid l, 0) & f_{t+1}(h \mid l, 1) & f_{t+1}(h \mid h, 0) & f_{t+1}(h \mid h, 1) \end{pmatrix} \begin{pmatrix} \phi_{l,0}(x_t, d_t) \\ \phi_{l,1}(x_t, d_t) \\ \phi_{h,0}(x_t, d_t) \\ \phi_{h,1}(x_t, d_t) \end{pmatrix},$$

where the terminal state x_{t+2} is generated by the time-$(t+1)$ transition $f_{t+1}(\cdot \mid x_{t+1}, d_{t+1})$. The conditioning state $x_{t+1} \in \{l, h\}$ selects columns: columns 1–2 for $x_{t+1} = l$ and columns 3–4 for $x_{t+1} = h$. These columns match the flow-vector entries $\phi_{x_{t+1}, d_{t+1}}(x_t, d_t)$ on the right. The QP (13) matches $\kappa_{t+2}(\cdot \mid x_t, d_t)$ against $\kappa_{t+2}(\cdot \mid x_t, d'_t)$. Matching terminal distributions across the two time-t actions d_t, d'_t reduces to a finite-dimensional linear-quadratic problem.

B Scalable Computation for Large State Spaces

The GFD flow-QP is computationally practical at large state spaces once we exploit the reachability and Kronecker structure of applied models.

The joint feasibility QP (27) admits two parameterizations. The dense form has decision variable $\operatorname{vec} \mathbf{\Phi}(x) \in \mathbb{R}^{D \cdot S^\rho \cdot D^\rho}$ per state x, which for $S = 20$, $D = 3$, and $\rho = 3$ is 648,000 entries. The equivalent flow form parameterizes the same object on the history tree, which has $\sum_{\tau=0}^{\rho-1} S \cdot (D \cdot S)^\tau$ nodes. For the same configuration this is 73,220 nodes, i.e. $219{,}660 = 3 \times 73{,}220$ flow variables per action path (D weights per node); the joint QP over the two compared actions doubles this to $n_{\text{tot}} = 439{,}320$. Reachability pruning and iterative LSQR solving make this count tractable.

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

$$\mathcal{R}(d_t, x) \;:=\; \{x' \in \mathcal{X} : f_t(x' \mid x, d_t) > 0\}, \tag{42}$$

and let $b := \max_{d_t, x} |\mathcal{R}(d_t, x)|$ 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 $\{(\mathbf{h}, d_t, x') : x' \in \mathcal{R}(d_t, x_{t+\tau})\}$; 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. The investment model has $S = 20$, $D = 3$, and $\rho = 3$, with deterministic capital transitions and one Tauchen-stochastic productivity factor. Its reduction is reported in Table 9.

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.

Solving the flow-QP: direct KKT vs. iterative LSQR

For small problems the QP (27) is solved by forming the KKT system

$$\begin{pmatrix} 2\,\mathbf{C}^\top\mathbf{C} + \varepsilon\,\mathbf{I} & \mathbf{A}_{\text{eq}}^\top \\ \mathbf{A}_{\text{eq}} & -\delta\,\mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{w} \\ \boldsymbol{\lambda} \end{pmatrix} \;=\; \begin{pmatrix} \mathbf{0} \\ \mathbf{b}_{\text{eq}} \end{pmatrix}. \tag{43}$$

Here $\mathbf{C}$ maps flow variables to the terminal-distribution difference $\kappa_{t+\rho+1}(\cdot \mid x_t, d_t) - \kappa_{t+\rho+1}(\cdot \mid x_t, d'_t)$. The term ε is a Tikhonov regularizer. The term $\delta \approx 0$ is a tiny stabilizer used only when $\mathbf{A}_{\text{eq}}$ is rank-deficient (the batched solve below uses the exact-zero form (45)). Here n_{tot} denotes the total flow-variable count across both action paths, i.e. $n_{\text{tot}} = 2D \times$ (node count) (D weights per node, two paths). The solver uses the direct KKT system (43) for $n_{\text{tot}} \le 5{,}000$ and the iterative formulation otherwise. At $\rho = 2$ the 1,220-node tree of Table 9 gives $n_{\text{tot}} = 7{,}320$, already above the 5,000 switch point. This requires forming $\mathbf{C}^\top\mathbf{C}$ explicitly. The matrix has 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

$$\min_{\mathbf{w}} \left\| \begin{pmatrix} \sqrt{2}\,\mathbf{C} \\ \sqrt{\varepsilon}\,\mathbf{I} \\ \eta\,\mathbf{A}_{\text{eq}} \end{pmatrix} \mathbf{w} \;-\; \begin{pmatrix} \mathbf{0} \\ \mathbf{0} \\ \eta\,\mathbf{b}_{\text{eq}} \end{pmatrix} \right\|_2^2, \tag{44}$$

with $\eta \gg 1$ enforcing the equality constraints. Solving (44) via LSQR (Paige and Saunders, 1982) uses only matrix–vector products with the sparse stacked operator. Its memory cost is $O(\text{nnz})$ rather than $O(n_{\text{tot}}^2)$. The batched-solve discussion below explains why we avoid normal equations.

Empirical performance

Table 9 reports the combined effect of pruning and iterative solving on the investment model. At $\rho = 3$ the unpruned full tree exceeds available memory on a standard workstation. With pruning and LSQR the same problem solves in 0.50 s (Table 9). The terminal-distribution matching error stays below 10^{-15}.

ρ	Nodes (full)	Nodes (pruned)	Reduction	Solver	Time (s)
1	20	20	0%	Direct KKT	0.01
2	1,220	1,220	0%	LSQR	0.18
3	73,220	3,140	95.7%	LSQR	0.50

Table 9: Combined effect of reachability pruning and iterative LSQR solving on the investment model ($S = 20$, $D = 3$, $\beta = 0.95$). The terminal-distribution matching error is measured as $\|\kappa_{t+\rho+1}(\cdot \mid x_t, d_t) - \kappa_{t+\rho+1}(\cdot \mid x_t, d'_t)\|_1$.

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

The pruning-LSQR machinery above 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. The binding consideration is how to reuse a factorisation across the $|\mathcal{X}| \cdot (D-1)$ right-hand-side vectors $b^{(x_t,k)}$. These share a common constraint matrix A_{eq}, since only the initial-flow block depends on x_t. Two paths are implemented:

Dense singular-value-decomposition (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 the Linear Algebra PACKage (LAPACK) routine `dgesvd`, and apply it to the $|\mathcal{X}| \cdot (D-1)$ right-hand sides as a single matrix multiplication. For up to a few thousand columns the level-3 Basic Linear Algebra Subprograms (BLAS) path is typically the fastest option on current hardware, since the per-column marginal cost is dominated by the one-time SVD. The SVD cost is $O((|\mathcal{X}| \cdot D)^3)$, 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

$$\begin{pmatrix} 2\,\mathbf{C}^\top\mathbf{C} + \varepsilon\mathbf{I} & \mathbf{A}_{\text{eq}}^\top \\ \mathbf{A}_{\text{eq}} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \boldsymbol{\lambda} \end{pmatrix} = \begin{pmatrix} \mathbf{0} \\ \mathbf{b} \end{pmatrix}. \tag{45}$$

Both $\mathbf{C}$ and $\mathbf{A}_{\text{eq}}$ are sparse for any sparse-transition DGP, e.g. Tauchen-discretised AR(1), deterministic shifts, or Kronecker-separable factor models. Then $\mathbf{C}^\top\mathbf{C} + \varepsilon\mathbf{I}$ remains sparse. The KKT block is therefore symmetric-indefinite but sparse. We factor it once via `scipy.sparse.linalg.splu` (using the default Unsymmetric MultiFrontal PACKage (UMFPACK) backend) and then triangular-solve for each right-hand side $(\mathbf{0}; \mathbf{b}^{(x_t,k)})$. The factorisation cost scales sub-cubically in $|\mathcal{X}|$. The triangular-solve cost is essentially linear in $|\mathcal{X}|$ per right-hand side (RHS). End-to-end, the path scales considerably better than the dense pseudoinverse for general (non-Kronecker-separable) DGPs. Some DGPs admit the Kronecker factorisation (46) below (the case in Table 11). For these the structured solver dominates this KKT-LU path by a further factor of $\sim Z^c$. The total reduction over the dense full-state pseudoinverse is then $\sim Z^{1+c}$, derived below. Here Z is the exogenous-factor dimension and c the per-solve scaling exponent.

Why not normal equations. The standard alternative solves $\mathbf{x} = (\mathbf{A}_{\text{aug}}^\top\mathbf{A}_{\text{aug}})^{-1}\mathbf{A}_{\text{aug}}^\top\mathbf{b}$ via sparse `splu`. This squares the condition number of $\mathbf{A}_{\text{aug}}$. The penalty weight $w = 10^6$ is already needed to enforce the constraints to acceptable precision. Then $\text{cond}(\mathbf{A}_{\text{aug}}) \sim 10^6$ becomes $\text{cond}(\mathbf{A}_{\text{aug}}^\top\mathbf{A}_{\text{aug}}) \sim 10^{12}$. A small Tikhonov ridge cannot regularise this without introducing significant bias. In our trials this path produced estimates that diverged numerically at $|\mathcal{X}| = 200$. The KKT formulation (45) avoids the squaring entirely and is the recommended sparse path.

Structured Kronecker solver: when the state factors as endo × exo

The principal source of the scaling gains in Table 11 is the recognition that many applied DDC transitions factor as a Kronecker product

$$f(x' \mid x, a) \;=\; T_{\text{endo}}(k' \mid k, a) \;\cdot\; T_{\text{exo}}(o' \mid o), \tag{46}$$

with $x = (k, o)$, 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 10 formalises the reduction.

Lemma 10 (Action-invariant Kronecker cancellation)**.** *If* $\mathbf{F}_{d_t} = \mathbf{F}^{\text{endo}}_{d_t} \otimes \mathbf{F}^{\text{exo}}$ *with* $\mathbf{F}^{\text{exo}}$ *action-invariant, joint-state FD is equivalent to FD on the* x^{endo}*-projection, and the joint flow input factorizes as* $\Phi(x_t, d_t) = \Phi^{\text{endo}}(x^{\text{endo}}_t, d_t) \otimes \Phi^{\text{exo}}(x^{\text{exo}}_t)$*. Consequently* $\rho^*_{\text{joint}} = \rho^*_{\text{endo}}$ *regardless of the persistence of* $\mathbf{F}^{\text{exo}}$*.*

Proof. Under $\mathbf{F}_{d_t} = \mathbf{F}^{\text{endo}}_{d_t} \otimes \mathbf{F}^{\text{exo}}$ the ρ-step joint transition factorizes by the mixed-product property of the Kronecker product. It equals the product of ρ endogenous steps Kronecker-times the product of ρ exogenous steps. Because $\mathbf{F}^{\text{exo}}$ is action-invariant, the exogenous factor is the same matrix along both compared action paths d_t, d'_t. Hence the terminal joint measure factorizes as $\boldsymbol{\kappa}_{\text{joint}}(d) = \boldsymbol{\kappa}_{\text{endo}}(d) \otimes \boldsymbol{\kappa}^{\text{exo}}$ with the exogenous factor $\boldsymbol{\kappa}^{\text{exo}}$ common to d and d'. The terminal-matching condition $\boldsymbol{\kappa}_{\text{joint}}(d) = \boldsymbol{\kappa}_{\text{joint}}(d')$ is therefore equivalent, by cancelling the common Kronecker factor $\boldsymbol{\kappa}^{\text{exo}}$ (identical on both action paths by action-invariance, with unit total mass $\mathbf{1}^\top \boldsymbol{\kappa}^{\text{exo}} = 1$, so it neither vanishes nor annihilates the difference), to $\boldsymbol{\kappa}_{\text{endo}}(d) = \boldsymbol{\kappa}_{\text{endo}}(d')$. Thus joint FD at horizon ρ holds iff endogenous FD at horizon ρ holds. The equivalence at each fixed ρ makes the two sets of feasible horizons identical, so their minima coincide ($\rho^*_{\text{joint}} = \rho^*_{\text{endo}}$). A joint flow input realizing it is $\Phi^{\text{endo}} \otimes \Phi^{\text{exo}}$ with Φ^{exo} the action-invariant exogenous path measure. □

The optimal flow weight at the full state then factors as

$$\boldsymbol{\Phi}^{\text{full}}\big((k, o),\, a;\, (k', o')\big) \;=\; \boldsymbol{\Phi}^{\text{endo}}(k,\, a;\, k') \;\cdot\; T_{\text{exo}}(o' \mid o), \tag{47}$$

so the structured solver computes $\boldsymbol{\Phi}^{\text{endo}}$ on the M-state endogenous problem ($M := |\{k\}|$) and never materialises the $|\mathcal{X}|^2$ full-state weight tensor. Equation (47) is written for $\rho = 1$, where the exogenous path measure $\boldsymbol{\Phi}^{\text{exo}}$ of Lemma 10 reduces to the single-step kernel T_{exo}; for $\rho > 1$, T_{exo} is replaced by the ρ-step exogenous path measure.

Effective regressors via factored Kronecker form. With $\boldsymbol{\Phi}^{\text{endo}}$ in hand, the GFD-MLE effective regressor $\widetilde{H}[x, a, r]$ at full state $x = (k, o)$ 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 $z[(k, o), a, \cdot] = (o \cdot$

$k, -a, -a^2$):

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

$$\widetilde{H}[x,a,1] = -(a - a_{\text{ref}}) \;-\; \beta \sum_{k',a'} \Delta\mathbf{\Phi}^{\text{endo}}(k,a;k',a') \cdot a' \qquad \text{(cost)}$$

$$\widetilde{H}[x,a,2] = -(a^2 - a_{\text{ref}}^2) \;-\; \beta \sum_{k',a'} \Delta\mathbf{\Phi}^{\text{endo}}(k,a;k',a') \cdot a'^2 \qquad \text{(adjustment)}$$

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

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

where $\widehat{\psi}(a,x) = -\log \widehat{p}(a \mid x) + \gamma_E$ 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 $|\mathcal{X}| = 5{,}000 = 100 \cdot 50$, the asymptotic flop-count ratio is large: $\sim Z^4 = 6{,}250{,}000\times$ relative to dense pinv ($c = 3$), $\sim Z^3 = 125{,}000\times$ relative to KKT-LU ($c = 2$), and $\sim Z^{2.5} \approx 18{,}000\times$ at $c = 1.5$. These are order-of-magnitude bounds at large Z, not realized wall-clock figures. In the measured range ($Z \le 40$) constant-factor overhead (matrix construction, basis precomputation) dominates. Table 10 reports the structured-vs-standard flow-QP speedup, reaching 71.5× at $Z = 40$; Table 11 reports the full GFD estimator at $|\mathcal{X}| = 5{,}000$ running in 1.39s, 136× faster than NFXP.

Table 10 reports the standard-vs-structured flow-QP wall-clock on the investment model at $\rho = 1$ and endogenous dimension $M = 5$ across a sweep $Z \in \{4, 10, 20, 40\}$ of the exogenous productivity grid.

Table 11 sweeps the joint state space from $|\mathcal{X}| = 20$ to $|\mathcal{X}| = 5{,}000$ at fixed sample size ($N = 1000, T = 15$). Its baseline differs from Table 10: that table times the structured solver against the standard flow QP, whereas this one times the full GFD estimator against the NFXP benchmark.

GFD's wall-clock advantage over NFXP grows with the state space (Table 11). Estimator accuracy is comparable to NFXP across this sweep, but first-stage-limited at the largest $|\mathcal{X}|$. The smoothed CCP estimator of Appendix C addresses that last case.

Z	$S = MZ$	Standard (ms)	Structured (ms)	Speedup
4	20	30.7	5.4	5.7×
10	50	254	7.1	35.9×
20	100	908	18.6	48.7×
40	200	5,611	78.5	71.5×

Table 10: Kronecker scaling of the GFD flow QP on the Type-K investment model at $\rho = 1$. Endogenous dimension $M = 5$ fixed; exogenous productivity grid Z varies. Standard mode loops over all $S = MZ$ initial states; structured mode solves M endogenous-only QPs and tensors with $\mathbf{F}^{\text{exo}}$ via Lemma 10. Wall-clock is mean of 3 repeats on Apple M3 Ultra.

$\lvert\mathcal{X}\rvert$	GFD time (s)	NFXP time (s)	NFXP / GFD
20	0.017	0.30	17×
60	0.025	0.32	13×
200	0.054	0.89	17×
600	0.11	2.10	19×
2,000	0.31	34.6	112×
5,000	1.39	189.2	136×

Table 11: Scaling of GFD vs. NFXP wall-clock on the investment DGP across joint state-space size. Same DGP as Table 2; $|\mathcal{X}|$ varied via capital and productivity grid sizes ($5\times4, 10\times6, 20\times10, 30\times20, 50\times40, 100\times50$). Mean of 5 replications per row (3 at $|\mathcal{X}| = 5{,}000$). GFD uses the smoothed first-stage CCP estimator (Appendix C) and the structured Kronecker flow-QP solver of this appendix.

C 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 $p_{0,t}(d_t \mid x_t)$, which enters the value-difference identity through the Hotz–Miller correction $\widehat{\psi}_{d_t}(x_t) = -\log \widehat{p}(d_t \mid x_t) + \gamma_E$. 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 $\widetilde{h}_t$ of (32). With the empirical-frequency estimator

$$\widehat{p}^{\text{emp}}(d_t \mid x_t) \;=\; \frac{n_{x,d_t}}{n_x}, \qquad n_{x,d_t} := \sum_{i,t} \mathbf{1}\{x_{it} = x,\; d_{it} = d_t\}, \tag{48}$$

the per-state variance is $\text{Var}[\widehat{p}^{\text{emp}}(d_t \mid x_t)] = p_{0,t}(d_t \mid x_t)(1 - p_{0,t}(d_t \mid x_t))/n_x$, 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 $|\mathcal{X}| = 2{,}000$. As $|\mathcal{X}|$ grows the empirical-CCP first stage becomes the binding source of GFD finite-sample error, which the smoother below addresses.

Multinomial-logit smoother on a low-dimensional basis

We replace (48) with a parametric multinomial logit fit on a low-dimensional basis $\mathbf{z}(x) \in \mathbb{R}^{p_{\text{dim}}}$ of the state. For the investment DGP of the investment Monte Carlo we use

$$\mathbf{z}(x) \;=\; \big(1,\; k(x),\; o(x),\; k(x)\cdot o(x)\big)^{\top}, \qquad p_{\text{dim}} = 4, \tag{49}$$

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

$$\ell(\boldsymbol{c}) \;=\; \sum_{i,t} \log \frac{\exp\big(\mathbf{z}(x_{it})^{\top}\boldsymbol{c}_{d_{it}}\big)}{\sum_{a'=0}^{D-1}\exp\big(\mathbf{z}(x_{it})^{\top}\boldsymbol{c}_{a'}\big)}, \qquad \boldsymbol{c}_0 \equiv \mathbf{0}. \tag{50}$$

The maximizer is computed by Newton iteration on the joint Hessian. Because $\ell(\boldsymbol{c})$ 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:

$$\widehat{p}^{\,\text{logit}}(d_t \mid x_t) \;=\; \frac{\exp\big(\mathbf{z}(x)^{\top}\widehat{\boldsymbol{c}}_{d_t}\big)}{\sum_{a'=0}^{D-1}\exp\big(\mathbf{z}(x)^{\top}\widehat{\boldsymbol{c}}_{a'}\big)}, \qquad x \in \mathcal{X}. \tag{51}$$

The entire smoother runs in $O\big(NT\cdot[p_{\text{dim}}(D-1)]^2\big)$ time per Newton iteration, dominated by the assembly of the $p_{\text{dim}}(D-1)\times p_{\text{dim}}(D-1)$ Hessian; an implementation accompanies the replication package.

Variance properties

The pointwise variance of $\widehat{p}^{\,\text{logit}}(d_t \mid x_t)$ scales as $p_{\text{dim}}/(NT)$ rather than as $|\mathcal{X}|/(NT)$, because all NT observations contribute to the estimation of every $\widehat{\boldsymbol{c}}_{d_t}$. For the investment calibration of the investment Monte Carlo, this gives a per-state variance reduction factor of approximately $|\mathcal{X}|/p_{\text{dim}} = 500$ at $|\mathcal{X}| = 2{,}000$, i.e. a $\sqrt{500} \approx 22\times$ reduction in RMSE. That reduction keeps the smoothed first-stage RMSE more than an order of magnitude below the $|\mathcal{X}|/(NT)$ scaling of the raw empirical-frequency estimator as the state space grows (Table 11).

When parametric smoothing is appropriate

The logit smoother is consistent for the true CCP function whenever $p_{0,t}(d_t \mid x_t) = \Lambda_{d_t}(\mathbf{z}(x)^{\top}\boldsymbol{c}_{d_t})$ for some $\boldsymbol{c}$, i.e. when the basis $\mathbf{z}(\cdot)$ is rich enough to span the log-odds surface $\log[p_{0,t}(d_t \mid x_t)/p_{0,t}(0 \mid x_t)]$. In our DGPs this is satisfied exactly because the equilibrium CCP is itself a multinomial logit in the basis $\mathbf{z}$ of (49) (the structural payoff is linear in the same regressors). When this exact-parametric assumption fails, three robust alternatives apply, with varying degrees of variance reduction:

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

3. Frequency estimator with Laplace smoothing. $\widehat{p}^{\text{lap}}(d_t \mid x_t) = (n_{x,d_t} + \alpha)/(n_x + \alpha D)$ with $\alpha \in (0, 1]$. This is cheap and partly mitigates the zero-cell problem. It 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. Under an oracle (known-CCP) second stage, or for first stages with asymptotically equivalent influence functions, the GFD large-sample efficiency is invariant to this choice. With the finite-sample CCP plug-in correction of Theorem 5, different first stages can move the feasible asymptotic variance. Finite-sample RMSE is in any case highly sensitive when $|\mathcal{X}|$ grows relative to NT.

D Applied Illustration: Altug–Miller (1998) Female Labour Supply

This appendix illustrates the GFD framework on a published structural labour supply model. We use Altug–Miller (AltM) to denote Altug and Miller (1998), an early application of the finite-dependence cancellation to estimation of a dynamic discrete choice female labour-supply model. It develops a specific path-based flow input $\mathbf{\Phi}$ tailored to their three-lag-register state space. We apply our canonical Moore–Penrose $\mathbf{\Phi}$ (Section 2.4) to the same model on the same PSID 1968–1985 dataset and report the resulting structural estimates. We do not claim numerical replication of AltM 1998's published Table VI 2.a values. We report on which parameters our estimates agree with AltM in sign and which do not. We then enumerate the specific implementation simplifications relative to AltM's full pipeline that drive the residual magnitude gaps.

Model. A woman at year t chooses to participate in market work ($d_{nt} = 1$) or not ($d_{nt} = 0$). State $x_{nt} = (\mathbb{1}\{l_{n,t-1} > 0\}, \mathbb{1}\{l_{n,t-2} > 0\}, \mathbb{1}\{l_{n,t-3} > 0\}) \in \{0, 1\}^3$ encodes the past three years of participation, giving $|\mathcal{X}| = 8$ and $\rho^* = 3$ (the canonical Type-S shift-register horizon). Here l_{nt} denotes hours worked (AltM's continuous labour-supply measure, in thousands). The FD state is the binary participation history $\mathbb{1}\{l > 0\}$. It is distinct from the hours terms entering the log-odds. The participation log-odds equation (Altug and Miller, 1998, eq. 6.7) is

$$\ln\frac{p_{nt}}{1 - p_{nt}} \;=\; \mathbf{z}'_{nt}\,\mathbf{B}_0 \;+\; (\mathbf{z}_{nt}\,l_{nt})'\,\mathbf{B}_1 \;+\; \delta_0 l_{nt}^2 \;+\; \sum_{s=1}^{3} \delta_s\, l_{nt}\, l_{n,t-s} \;+\; \eta_n\,\lambda_t\,w_{nt}\,l_{nt} \;+\; C(x_{nt}),$$

with $\mathbf{z}_{nt} = (1, \text{KIDS}_{nt}, \text{KIDS}_{nt}{\cdot}\text{MAR}_{nt}, \text{AGE}_{nt}, \text{AGE}_{nt}^2)$. The parameter is $\theta = (\mathbf{B}_0, \mathbf{B}_1, \delta_0, \delta_1, \delta_2, \delta_3) \in \mathbb{R}^{14}$. The continuation correction $C(x_{nt})$ is the only object that distinguishes the two estimators we compare: AltM 1998's truncated-CCP form sums $\rho = 3$ Hotz–Miller log-ratios along a single path, while the GFD form uses the canonical Moore–Penrose flow weights integrated over all ρ-step paths.

Data. We download raw Panel Study of Income Dynamics (PSID) waves 1968–1985 and parse SAS dictionaries to per-year CSVs. We construct a female panel with the AltM 1998 sample-selection criteria. This restricts to married women aged 25–65 with non-missing participation, labour hours, and demographic regressors. After attaching three-period participation lags and dropping observations with incomplete lag histories, the working sample contains 11,915 woman-year observations spanning 1,849 women across 1973–1985. The accompanying scripts in `code/altug_miller_replication/` document the full extraction pipeline.

Estimation results. We run AltM-CCP and GFD as two CCP-based two-step estimators on the same individual-level participation likelihood. They share Steps A (hours Euler equation), A.5 (latent hours for non-workers), and the first-stage CCPs. Table 12 reports the 14-parameter vector together with AltM 1998's published Table VI 2.a values.

Parameter	AltM-CCP	GFD	AltM 1998 Tab. VI 2.a	sign-match?	GFD/AltM-CCP
B_{00} (const)	+2.11	−0.14	−14.0	n/a (not comparable)	n/a
B_{01} KIDS	−0.19	−0.21	+204	×	1.08
B_{02} KIDS·MAR	+0.14	+0.12	−187	×	0.85
B_{03} AGE	−0.17	−0.16	+31	×	0.97
B_{04} AGE^2	+0.002	+0.001	−0.03	×	0.96
B_{10} l	+125.0	+129.3	−5.6	×	1.03
B_{11} KIDS·l	+26.0	+26.2	+0.16	✓	1.01
B_{12} KIDS·MAR·l	+38.1	+38.0	+0.18	✓	1.00
B_{13} AGE·l	+1.99	+2.10	+0.31	✓	1.06
B_{14} AGE$^2 \cdot l$	+0.029	+0.031	+0.002	✓	1.07
δ_0 l^2	−35.9	−37.6	+0.002	×	1.05
δ_1 $l \cdot l_{t-1}$	−7.78	−8.24	−0.0005	✓	1.06
δ_2 $l \cdot l_{t-2}$	+5.94	+6.51	−0.0007	×	1.10
δ_3 $l \cdot l_{t-3}$	−15.5	−16.3	−0.0007	✓	1.05

Table 12: 14-parameter participation logit estimates on PSID 1973–1985: AltM-CCP (the truncated-path continuation correction of Altug and Miller, 1998) vs. GFD (the canonical Moore–Penrose flow input of Section 2.4). Both estimators share Step A, Step A.5, and the first-stage CCPs; they differ only in $C(x_{nt})$. AltM 1998 published values are from Table VI col 2.a. GFD and AltM-CCP agree to within ∼10–15% on the 13 slope coefficients (rightmost column). The intercept B_{00} is the exception. It absorbs the differing offset-level normalization of the two FD inputs and so is not directly comparable. This is consistent with the consistency-for-any-FD-input result of Theorem 5 on real PSID data.

Findings. The slope agreement reported in Table 12 shows that two different FD-feasible $\mathbf{\Phi}$ choices (AltM 1998's path-based form and our canonical Moore–Penrose form) deliver the same slope θ in finite sample up to first-stage noise. Both estimators match the published AltM 1998 Table VI 2.a values in sign on the parameters checkmarked in Table 12. The residual magnitude gap is traceable to a small number of AltM-specific implementation choices our pipeline does not reproduce.

Sources of the residual magnitude gap. The following AltM-specific implementation details account for the gap. Each is present in AltM's full pipeline but not in our minimum-faithful reproduction: (i) kernel-density derivative terms f'/f in Step A's hours Euler equation (AltM eq. 6.6) that account for the truncation of hours at zero for non-workers; our Step A runs ordinary least squares on workers only. (ii) MaCurdy-style nonparametric projection of woman-specific fixed effects η_n onto time-invariant observables (race, education, kids, marital, region, age at panel start). Our implementation uses the simpler traditional fixed-effects approach (per-woman time averages of level residuals). (iii) Implicit hours-unit conventions in AltM's published Table VI that we cannot fully infer from the published text. Our pipeline uses hours in thousands ($l_k = l/1000$). If AltM's published coefficients are on raw annual hours scale, the δ coefficients would differ by 10^6 from the conversion alone. That conversion partially offsets the published-vs-our gap. (iv) An ad-hoc η-scale calibration of $k = 5/\text{median}(\eta\lambda wl)$ that we apply to bring the wage offset onto AltM's implicit scale; AltM's actual normalization may differ. The empirical content of this appendix is therefore the consistency-for-any-FD-input check between two GFD-class estimators (AltM-CCP and GFD) on PSID data.

E The GFD-EM Extension: Latent Unobserved Heterogeneity

This appendix records how the GFD identity extends to finite mixtures. The extension is useful when agents differ by latent type. It is not part of the paper's main empirical contribution.

E.1 The GFD-EM algorithm

Agents of unobserved type $\ell \in \{1, \ldots, L\}$ may differ in preferences $\theta^{(\ell)}$, transitions $f^{(\ell)}$, or both. The observed CCP is the mixture $p_t(d_t \mid x_t) = \sum_\ell \pi_\ell \, p_{\ell,t}(d_t \mid x_t)$. The FD existence test depends only on f. A single flow input $\widehat{\mathbf{\Phi}}$ therefore works across types when transitions are common. When transitions differ, the test of Theorem 3 is applied per type. The model then runs at $\rho = \max_\ell \rho^*_\ell$ with type-specific $\widehat{\mathbf{\Phi}}^{(\ell)}$.

Before the EM loop, solve the flow QP for $\widehat{\mathbf{\Phi}}^{(\ell)}$ and form the effective regressors $\widehat{H}^{(\ell)}$. The EM iteration alternates three steps. Step (E) refreshes the posterior type weights $q_{i\ell}^{(m)} \propto \widehat{\pi}_\ell^{(m-1)} \widehat{p}_0^{(\ell)}(x_{i,t_0}) \prod_t \widehat{p}_\ell^{(m-1)}(d_{it} \mid x_{it}) \prod_t \widehat{f}_t^{(\ell)}(x_{i,t+1} \mid x_{it}, d_{it})$. It also forms the type-specific Hotz–Miller offsets $\widehat{h}^{(\ell),(m)}$. This posterior is the choice likelihood multiplied by the type-specific initial-state and transition probabilities. These factors do not cancel when transitions differ by type; they drop out only under common transitions. Step (M) maximizes the type-weighted GFD pseudo-log-likelihood $\widehat{\theta}^{(m)} = \arg\max_\theta \sum_{i,t,\ell} q_{i\ell}^{(m)} \log \Lambda_{d_{it}}(x_{it}; \theta^{(\ell)}, \widehat{\mathbf{\Phi}}^{(\ell)}, \widehat{\mathbf{p}}_\ell^{(m-1)})$. Step (U) updates the mixture weights $\widehat{\pi}_\ell^{(m)} = N^{-1} \sum_i q_{i\ell}^{(m)}$ and the type-specific CCPs via Λ at the new $(\widehat{\theta}^{(m)}, \widehat{\mathbf{\Phi}}^{(\ell)})$.

The effective regressor $\widehat{H}^{(\ell)}$ depends on $(\widehat{f}^{(\ell)}, \widehat{\mathbf{\Phi}}^{(\ell)})$ but not on the iterating CCPs. It is therefore constructed once before the EM loop and reused across iterations. Only the Hotz–Miller offset $\widehat{h}^{(\ell),(m)}$ is refreshed each iteration. The Arcidiacono–Miller EM estimator (AM-EM; Arcidiacono and Miller, 2011) instead recomputes the matrix inversion $(I - \beta F_{\widehat{\mathbf{p}}})^{-1}$ each iteration because the CCP-weighted transition depends on the iterating CCPs.

E.2 Non-stationary mixtures

For non-stationary finite-horizon models, the GFD identity of Theorem 4 applies type by type.

Corollary 11 (GFD identity under unobserved heterogeneity)**.** *Let the data generating process be a finite-mixture DDC with type prior $\boldsymbol{\pi}$ and type-conditional primitives $(u_t^{(\ell)}, f_t^{(\ell)}, \theta^{(\ell)})$. Suppose type-conditional pointwise FD holds at horizon ρ_ℓ^* for every $\ell \in \{1, \dots, L\}$ in the sense of Theorem 3. Then for each (x_t, d_t, d_t') with $d_t \neq d_t'$ and each ℓ, there exists a type-conditional flow $\boldsymbol{\Phi}^{(\ell)}(x_t)$ such that*

$$v_t^{(\ell)}(x_t, d_t) - v_t^{(\ell)}(x_t, d_t') = u_t^{(\ell)}(d_t, x_t; \theta^{(\ell)}) - u_t^{(\ell)}(d_t', x_t; \theta^{(\ell)}) + \sum_{\tau=1}^{\rho_\ell^*} \beta^\tau \sum_{(k_\tau, j_\tau)} \boldsymbol{\Phi}^{(\ell)}(x_t; d_t, d_t')[(k_\tau), (j_\tau)] \left[u_{t+\tau}^{(\ell)}(k_\tau, j_\tau; \theta^{(\ell)}) + \psi_{t+\tau}^{(\ell)}(k_\tau, j_\tau) \right], \tag{52}$$

where $\psi_t^{(\ell)}(k, j) = \gamma_E - \log p_t^{(\ell)}(k \mid j)$ is the type-conditional Hotz–Miller correction. The term $\boldsymbol{\Phi}^{(\ell)}(x_t; d_t, d_t')[(k_\tau), (j_\tau)]$ denotes the time-$(t+\tau)$ marginal over the action–state pair (k_τ, j_τ) at horizon τ. This marginal is taken over the differenced flow $\Delta\boldsymbol{\phi}^{(\ell)} := \boldsymbol{\phi}^{(\ell)}(x_t, d_t) - \boldsymbol{\phi}^{(\ell)}(x_t, d_t')$. This is the τ-step projection of the full path-flow difference of Theorem 4. The mixture log-likelihood inherits this representation and is computable from $(\widehat{\boldsymbol{\Phi}}^{(\ell)}, \widehat{p}_\tau^{(\ell)})$ without invoking any type-conditional Bellman fixed point.

The corollary follows by applying Theorem 4 (with feasibility from Theorem 3) to each type-conditional submodel. The mixture aggregation does not interact with the GFD identity. Each type's flow weights and CCP corrections are computed within the type. Proofs are in Appendix F.9.

Theorem 12 (GFD-EM consistency, non-stationary mixtures)**.** *Let $\{(d_{n,t}, x_{n,t})\}$ be an i.i.d. short panel over $t \in [t_{start}, t_{end}]$ of length T_{obs} generated by the mixture DDC of Corollary 11. Assume:*

(i) (FD feasibility) For each type ℓ, the type-conditional pointwise FD condition holds at some horizon $\rho_\ell^* < T_{obs}$ for every (x, d, d') on a set of positive sample mass.

(ii) (Mixture identification) The true type-conditional CCPs $p_t^{(\ell)}(d \mid x; \theta_0^{(\ell)})$ are pairwise distinct in ℓ on a set of positive sample mass at some $t \in [t_{start}, t_{end} - \rho_\ell^*]$. In addition, the finite-mixture rank condition of Kasahara and Shimotsu (2009) holds. That is, the matrices of type-conditional choice/transition probabilities across the repeated within-panel measurements are linearly independent. This identifies the mixture weights $\boldsymbol{\pi}$ and the component-specific CCPs. Pairwise distinctness alone is a no-collapse condition, not a mixture identification condition.

(iii) (Regularity) The parameter space $\Theta \times \Delta^{L-1}$ is compact; the type-conditional payoffs $u_t^{(\ell)}(d, x; \theta)$ are continuously differentiable in θ with bounded derivatives; the empirical CCP anchor $\widehat{p}_t^{(\ell)}(d \mid x)$ at the boundary periods is $\sqrt{N}$-consistent for the truth.

Then the global maximizer $(\widehat{\boldsymbol{\pi}}, \widehat{\boldsymbol{\theta}})$ of the GFD-representation pseudo-log-likelihood converges in probability to $(\boldsymbol{\pi}_0, \boldsymbol{\theta}_0)$ as $N \to \infty$ (up to label permutation) and is asymptotically normal with sandwich variance characterized by the GFD-based score and Hessian. The EM iteration itself converges only to a stationary point of this criterion (Wu, 1983). Consistency is therefore a property of the argmax. We use a multi-start protocol to select the global optimum among the EM limit points.

Proof sketch. The GFD representation (52) expresses the type-conditional choice probabilities as continuous functions of $(\theta^{(\ell)}, \mathbf{\Phi}^{(\ell)}, p^{(\ell)})$. The GFD-pseudo likelihood therefore converges uniformly to a population analogue. Under (ii) the unique maximizer of that analogue is the truth modulo label switching. The full four-step argument (uniform anchor consistency, identification, EM convergence, asymptotic normality) appears in Appendix F.10.

F Proofs

F.1 Proof of Theorem 2 (constraint–decomposition equivalence)

The proof uses two marginal flows. Throughout, $\tau = 1, \dots, \rho$ indexes the offset from the base time t. For each τ and each prefix, define

$$M^{\mathrm{s}}_{\tau}(x_{t+1:t+\tau}, d_{t+1:t+\tau-1}) := \sum_{x_{t+\tau+1:t+\rho}} \sum_{d_{t+\tau:t+\rho}} \phi_{x_{t+1:t+\rho}, d_{t+1:t+\rho}}(x_t, d_t),$$
$$M^{\mathrm{sa}}_{\tau}(x_{t+1:t+\tau}, d_{t+1:t+\tau}) := \sum_{x_{t+\tau+1:t+\rho}} \sum_{d_{t+\tau+1:t+\rho}} \phi_{x_{t+1:t+\rho}, d_{t+1:t+\rho}}(x_t, d_t).$$

M^{s}_{τ} is the marginal flow through the state-action prefix that has consumed states $x_{t+1:t+\tau}$ and actions $d_{t+1:t+\tau-1}$, so the time-$(t+\tau)$ action $d_{t+\tau}$ is still free. M^{sa}_{τ} is the same marginal with $d_{t+\tau}$ also fixed.

($\Leftarrow$). Substitute the product form (9) into the left-hand side (LHS) of the initial-flow constraint (10). Sum over $(x_{t+2:t+\rho}, d_{t+1:t+\rho})$. The transition factors $f_{t+\tau-1}(x_{t+\tau} \mid x_{t+\tau-1}, d_{t+\tau-1})$ for $\tau \geq 2$ sum to one over the next state. The weights $\omega_{t+\tau}$ for $\tau \geq 1$ sum to one over the next action. The only surviving factor is $f_t(x_{t+1} \mid x_t, d_t)$, which matches the RHS. An analogous telescoping verifies the flow-conservation constraint (11). Finally, the product form gives $M^{\mathrm{sa}}_{\tau} = \omega_{t+\tau} M^{\mathrm{s}}_{\tau}$ at every prefix, so $M^{\mathrm{s}}_{\tau} = 0$ forces $M^{\mathrm{sa}}_{\tau} = 0$: the no-ghost condition holds automatically.

($\Rightarrow$). On every history with $M^{\mathrm{s}}_{\tau} \neq 0$, define

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

Summing the numerator over $d_{t+\tau}$ recovers the denominator, so $\sum_{d_{t+\tau}} \omega_{t+\tau}(d_{t+\tau} \mid \cdot) = 1$. Rewrite flow conservation (11) as $M^{\mathrm{s}}_{\tau+1}(x_{t+1:t+\tau+1}, d_{t+1:t+\tau}) = f_{t+\tau}(x_{t+\tau+1} \mid x_{t+\tau}, d_{t+\tau}) \cdot M^{\mathrm{sa}}_{\tau}(x_{t+1:t+\tau}, d_{t+1:t+\tau})$. Combined with the initial-flow base case $M^{\mathrm{s}}_{1}(x_{t+1}) = f_t(x_{t+1} \mid x_t, d_t)$, telescoping yields the chain identity $\phi = \prod_{\tau} f_{t+\tau-1} \cdot \omega_{t+\tau}$. Uniqueness on non-zero-marginal histories is immediate from the definition of $\omega_{t+\tau}$ as a ratio of marginals. On a zero-marginal history ($M^{\mathrm{s}}_{\tau} = 0$) the ratio is undefined. The no-ghost condition forces $M^{\mathrm{sa}}_{\tau} = 0$ there, so every supported ϕ entry below that history is zero. The chain identity then holds with $\omega_{t+\tau}$ assigned by any unit-sum convention. □

F.2 Proof of Lemma 1 (Arcidiacono–Miller representation)

Substituting $\psi_{d_t}(\mathbf{p}_t(x_t)) = V_t(x_t) - v_t(x_t, d_t)$ into (3) gives $\sum_{d_t} \omega(d_t \mid x_t)\big[v_t(x_t, d_t) + V_t(x_t) - v_t(x_t, d_t)\big] = V_t(x_t) \sum_{d_t} \omega(d_t \mid x_t) = V_t(x_t)$, since the weights sum to unity. □

F.3 Proof of Theorem 3 (pointwise FD as linear feasibility)

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 of Definition 1 (initial flow, flow conservation, terminal matching). Below, $a \in \{d_t, d'_t\}$ ranges over the two scenarios (the analyzed action and the reference action); the upper sub-block of each pair corresponds to $a = d_t$ and the lower to $a = d'_t$.

Initial-flow blocks. For each $a \in \{d_t, d'_t\}$ and each $x_{t+1} \in \mathcal{X}$,

$$\big(\mathbf{A}_{\text{init}} \operatorname{vec}\Phi(x_t, a)\big)_{x_{t+1}} \;=\; \sum_{x_{t+2:t+\rho}} \sum_{d_{t+1:t+\rho}} \phi_{(x_{t+1},\, x_{t+2:t+\rho}),\, d_{t+1:t+\rho}}(x_t, a),$$

which equals $\mathbf{f}^{(a)}_{x_{t+1}} = f_t(x_{t+1} \mid x_t, a)$ exactly when $\Phi(x_t, a)$ satisfies the initial-flow constraint (10).

Flow-conservation (NAC) blocks. The row of $\mathbf{A}_{\text{nac}}$ indexed by $(\tau, x_{t+1:t+\tau}, d_{t+1:t+\tau}, x_{t+\tau+1})$, applied to $\operatorname{vec}\Phi(x_t, a)$ for $a \in \{d_t, d'_t\}$, evaluates to the LHS minus the RHS of (11) at that prefix. The block evaluates to $\mathbf{0}$ for both $a = d_t$ and $a = d'_t$ exactly when both $\Phi(x_t, d_t)$ and $\Phi(x_t, d'_t)$ satisfy flow conservation.

Terminal-matching block. By (12), $\mathbf{T}_{\text{last}} \operatorname{vec}\Phi(x_t, d_t) = \kappa_{t+\rho+1}(\cdot \mid x_t, d_t)$. By the same identity, $\mathbf{T}_{\text{last}} \operatorname{vec}\Phi(x_t, d'_t) = \kappa_{t+\rho+1}(\cdot \mid x_t, d'_t)$. Hence the last block evaluates to $\kappa_{t+\rho+1}(\cdot \mid x_t, d_t) - \kappa_{t+\rho+1}(\cdot \mid x_t, d'_t)$, which equals $\mathbf{0}$ exactly when (14) holds.

Hence $\mathbf{A}\,\mathbf{x} = \mathbf{b}$ is equivalent to the three conditions of Definition 1. Therefore pointwise FD holds iff there exists $\mathbf{x}$ with $\mathbf{A}\,\mathbf{x} = \mathbf{b}$, i.e. iff $\mathbf{b} \in \text{Range}(\mathbf{A})$, which is (20). The constructive statement (second sentence of the theorem) is immediate from the equivalence. □

F.4 Proof of Theorem 4 (GFD value-difference identity)

Sketch. We write the per-period payoff in the linear form $z_t(x_t, d_t)^\top \theta$ for concreteness. The argument uses only additive separability (Assumption 1) and the cancellation of the continuation term under the common-terminal condition. The same argument applies to a general $u_t(x_t, d_t; \theta)$ with $z_t^\top \theta$ replaced by $u_t(\cdot; \theta)$ throughout. 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 $v_t(x_t, d_t; \theta)$ becomes

$$\begin{aligned} v_t(x_t, d_t; \theta) \;=\; & z_t(x_t, d_t)^\top \theta \;+\; \sum_{\tau=1}^{\rho} \beta^\tau \sum_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}} \phi_{x_{t+1:t+\rho}, d_{t+1:t+\rho}}(x_t, d_t) \cdot \mu_{t+\tau}(x_{t+\tau}, d_{t+\tau}; \theta) \\ & +\; \beta^{\rho+1} \sum_{x_{t+\rho+1}} \kappa_{t+\rho+1}(x_{t+\rho+1} \mid x_t, d_t)\, V_{t+\rho+1}(x_{t+\rho+1}), \end{aligned}$$

where $\mu_{t+\tau}$ is the augmented flow payoff (4). Because each $\mu_{t+\tau}$ depends on the path only through the time-$(t+\tau)$ coordinate $(x_{t+\tau}, d_{t+\tau})$, summing the full-path weight ϕ over the remaining coordinates collapses each term to the time-$(t+\tau)$ marginal of Φ. Taking the difference of this identity at d_t and the reference action 0 gives the marginal contrast $\Delta\phi := \phi(x_t, d_t) - \phi(x_t, 0)$. Under the reference normalization $u_t(x_t, 0; \theta) = 0$, the leading payoff term of the difference reduces to $u_t(x_t, d_t; \theta)$, matching the right-hand side of (23). Invoking the common-terminal condition (22) then yields (23). □

F.5 Proof of Theorem 8 (payoff-only counterfactual CCPs from FD)

Single-agent case, existence and uniqueness. The Bellman contraction at θ_1 has a unique fixed point V_1; by the Hotz–Miller bijection it induces a unique equilibrium CCP profile $\mathbf{p}_1$. This existence and uniqueness rest on the contraction and the bijection alone, independently of the FD iteration in the convergence step below. Fixed-point characterization. Because $\mathbf{p}_1$ is the equilibrium profile at θ_1, Theorem 4 applies at $(\theta_1, \mathbf{p}_1)$, so the structural value-difference admits the FD representation (23). Substituting it into the Hotz–Miller logit-CCP map at θ_1 shows that $\mathbf{p}_1$ satisfies (36). Local convergence. Suppose the spectral radius of the composite map's Jacobian at $\mathbf{p}_{1,t}$ is below unity. The FD iteration is then a local contraction (spectral-radius theorem), so the Banach fixed-point theorem gives local isolation and linear convergence at rate equal to that spectral radius.

Game case (Proposition 9). Fix the selected counterfactual profile $\mathbf{p}_1$ of Assumption 10. For each player i, hold rivals' counterfactual CCPs $\mathbf{p}_{-i,1}$ fixed and form the opponent-integrated transition $f_i^{\mathbf{p}_{-i,1}}$. Player i's problem is then a single-agent payoff-only counterfactual. Theorem 4 applies because $\mathbf{\Phi}_i$ is FD-feasible for $f_i^{\mathbf{p}_{-i,1}}$. This yields the per-player identity (37). Assembling the I player-wise identities gives the joint system, of which $\mathbf{p}_1$ is a solution. Existence and selection of $\mathbf{p}_1$ are supplied by Assumption 10 (not by a contraction). □

F.6 Proof of Theorem 5 (consistency and asymptotic normality)

The GFD estimator is a two-step M-estimator, $\widehat{\theta}^{\mathrm{GFD}}(\widehat{\mathbf{\Phi}}) = \arg\max_{\theta\in\Theta} Q_N(\theta; \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}})$. Its criterion is $Q_N(\theta; \mathbf{\Phi}, \mathbf{p}) = (NT)^{-1} \sum_{n=1}^{N} \sum_{t=1}^{T} \log \Lambda_{d_{nt}}(x_{nt}; \theta, \mathbf{\Phi}, \mathbf{p})$, where Λ is the GFD logit map of Definition 3. The first-stage nuisance is $\eta = (f, \mathbf{p})$, entering the criterion through $\mathbf{\Phi} = \mathbf{\Phi}(\cdot; f)$ and through the Hotz–Miller correction $\psi(\mathbf{p})$ inside Λ. Write $s(\theta; \mathbf{\Phi}, \mathbf{p}) = \nabla_\theta \log \Lambda_{d_t}(x_t; \theta, \mathbf{\Phi}, \mathbf{p})$ for the score and $Q_0(\theta; \mathbf{\Phi}, \mathbf{p}) = \mathbb{E}[\log \Lambda_{d_t}(x_t; \theta, \mathbf{\Phi}, \mathbf{p})]$ for its population limit.

Step 1 (consistency, any FD-feasible $\mathbf{\Phi}$). Fix any FD-feasible $\mathbf{\Phi}$ under the true f. By Theorem 4, at θ_0 the GFD pseudo value-difference equals the true value-difference. Hence $\Lambda_d(x; \theta_0, \mathbf{\Phi}, \mathbf{p}_0) = p_0(d \mid x)$, the true CCP. This holds regardless of which FD-feasible $\mathbf{\Phi}$ is used. Hence $\theta \mapsto Q_0(\theta; \mathbf{\Phi}, \mathbf{p}_0)$ is a Kullback–Leibler criterion maximized uniquely at θ_0 under the global identification condition that $\theta \mapsto \Lambda(\cdot; \theta, \mathbf{\Phi}, \mathbf{p}_0)$ is injective on Θ. Non-singularity of $\mathbf{J}(\mathbf{\Phi})$ at θ_0 is the local rank/curvature counterpart. It is necessary for a well-separated maximum but not on its own sufficient for global injectivity, which we therefore assume (it holds for the linear-index logit specifications used throughout). Four regularity conditions enter here: continuity of Λ in $(\theta, \mathbf{\Phi}, \mathbf{p})$; compactness of Θ; a uniform law of large numbers; and $(\widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) \xrightarrow{p} (\mathbf{\Phi}^*, \mathbf{p}_0)$. They give $\sup_\theta |Q_N(\theta; \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) - Q_0(\theta; \mathbf{\Phi}^*, \mathbf{p}_0)| \xrightarrow{p} 0$. Hence $\widehat{\theta}^{\mathrm{GFD}} \xrightarrow{p} \theta_0$ by the standard argmax/consistency theorem (Newey and McFadden, 1994, Thm. 2.1). This holds for every FD-feasible $\widehat{\mathbf{\Phi}}$.

Step 2 (Neyman orthogonality in $\mathbf{\Phi}$). The mean-value/argmax identity just used holds for every FD-feasible $\mathbf{\Phi}$: $\mathbb{E}[s(\theta_0; \mathbf{\Phi}, \mathbf{p}_0)] = \mathbf{0}$ for all such $\mathbf{\Phi}$. Indeed, with $d \sim p_0(\cdot \mid x)$ and $\Lambda_d(\theta_0) = p_0(d \mid x)$,

$$\mathbb{E}[s(\theta_0; \mathbf{\Phi}, \mathbf{p}_0)] = \mathbb{E}_x \sum_d p_0(d \mid x) \frac{\nabla_\theta \Lambda_d}{\Lambda_d} = \mathbb{E}_x \sum_d \nabla_\theta \Lambda_d(x; \theta_0, \mathbf{\Phi}, \mathbf{p}_0) = \mathbb{E}_x \nabla_\theta \Big(\sum_d \Lambda_d \Big) = \mathbf{0},$$

since $\sum_d \Lambda_d \equiv 1$. The identity $\Lambda_d(\theta_0) = p_0(d \mid x)$ used here holds only for $\mathbf{\Phi}$ FD-feasible under the true f; thus $\mathbb{E}[s(\theta_0; \mathbf{\Phi}, \mathbf{p}_0)] = \mathbf{0}$ on the f-feasible manifold $\mathcal{F}(\cdot; f)$, and $\partial_{\mathbf{\Phi}}\mathbb{E}[s]$ vanishes in directions tangent to $\mathcal{F}(\cdot; f)$. It is not an identity in the ambient $\mathbf{\Phi}$: for $\mathbf{\Phi}$ infeasible under f, $\Lambda_d(\theta_0) \neq p_0$ and $\mathbb{E}[s] \neq \mathbf{0}$. Now f enters only through $\mathbf{\Phi}(\cdot; f)$, so $\partial_f \mathbb{E}[s] = \partial_{\mathbf{\Phi}}\mathbb{E}[s] \cdot \partial_f \mathbf{\Phi}$. The generated input is feasible under $\widehat{f}$, so $\partial_f \mathbf{\Phi}$ generally has a component off $\mathcal{F}(\cdot; f)$, along which $\partial_{\mathbf{\Phi}}\mathbb{E}[s]$ need not vanish. The $\widehat{f}$-channel then contributes a transition-induced influence term $\mathbf{D}_f \, \boldsymbol{\psi}^f$, with

$$\mathbf{D}_f := \partial_f \, \mathbb{E}\big[s(\theta_0; \mathbf{\Phi}(\cdot; f), \mathbf{p}_0)\big] = \partial_{\mathbf{\Phi}}\mathbb{E}[s] \cdot \partial_f \mathbf{\Phi},$$

where $\partial_{\mathbf{\Phi}}\mathbb{E}[s]$ annihilates the directions tangent to $\mathcal{F}(\cdot; f)$, so only the off-tangent part of $\partial_f \mathbf{\Phi}$ survives, and $\boldsymbol{\psi}^f$ is the transition block of $\boldsymbol{\psi}^{\mathrm{fs}}$. Under Assumption 9 the feasible set has locally constant rank, so $\partial_f \mathbf{\Phi}$ stays tangent to $\mathcal{F}(\cdot; f)$ and $\mathbf{D}_f = \mathbf{0}$: the $\widehat{f}$-channel is first-order negligible for $\widehat{\theta}^{\mathrm{GFD}}$. Only the direct dependence of Λ on $\mathbf{p}$ through the Hotz–Miller corrections then contributes a first-stage term. (The renewal, shift-register, and Kronecker classes satisfy rank stability; knife-edge/deterministic-zero DGPs retain $\mathbf{D}_f \, \boldsymbol{\psi}^f$ via the constrained delta-method.)

Step 3 (asymptotic linearization). Assume $\theta_0 \in \operatorname{int}\Theta$. A mean-value expansion of the first-order condition $\nabla_\theta Q_N(\widehat{\theta}; \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) = \mathbf{0}$ around θ_0, with $\mathbf{J}(\mathbf{\Phi}^*)$ non-singular and the Hessian obeying a uniform law of large numbers in a neighbourhood of θ_0 (standard stochastic-equicontinuity conditions, Newey and McFadden, 1994, Thm. 6.1), gives

$$\sqrt{N}(\widehat{\theta}^{\mathrm{GFD}} - \theta_0) = \mathbf{J}(\mathbf{\Phi}^*)^{-1} \sqrt{N} \, \nabla_\theta Q_N(\theta_0; \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) + o_p(1).$$

Write $\bar{s}_n := \frac{1}{T}\sum_{t=1}^{T} s(\theta_0; \mathbf{\Phi}^*, \mathbf{p}_0; x_{nt}, d_{nt})$ for the per-agent score. Expanding the score in the nuisance $(\widehat{f}, \widehat{\mathbf{p}})$ via Assumption 7 and collecting the two channels of Step 2,

$$\sqrt{N} \, \nabla_\theta Q_N(\theta_0; \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) = \frac{1}{\sqrt{N}} \sum_{n=1}^{N} \Big[\bar{s}_n + \mathbf{D}_p \, \boldsymbol{\psi}^{\mathrm{p}}(\zeta_n) + \mathbf{D}_f \, \boldsymbol{\psi}^f(\zeta_n) \Big] + o_p(1),$$

where $\mathbf{D}_p = \partial_{\mathbf{p}}\mathbb{E}[s(\theta_0; \mathbf{\Phi}^*, \mathbf{p})]\big|_{\mathbf{p}_0}$ is the total derivative collecting every occurrence of $\mathbf{p}$ in Λ (the period-t Hotz–Miller offset and the future-period corrections in the augmented payoffs $\mu_{t+\tau}$), finite by the CCP positivity of Assumption 7; $\mathbf{D}_f$ is the transition-induced derivative of Step 2; and $\boldsymbol{\psi}^{\mathrm{p}}, \boldsymbol{\psi}^f$ are the CCP and transition blocks of $\boldsymbol{\psi}^{\mathrm{fs}}$. Under Assumption 9, $\mathbf{D}_f = \mathbf{0}$ and the last term drops. The bracket is i.i.d. across agents n, mean zero, with finite variance

$$\mathbf{\Omega}(\mathbf{\Phi}^*) = \operatorname{Var}\big(\bar{s}_n + \mathbf{D}_p \boldsymbol{\psi}^{\mathrm{p}}(\zeta_n) + \mathbf{D}_f \boldsymbol{\psi}^f(\zeta_n)\big),$$

which reduces to $\operatorname{Var}(\bar{s}_n + \mathbf{D}_p \boldsymbol{\psi}^{\mathrm{p}})$ under rank stability and carries the within-agent serial dependence through $\bar{s}_n$.

Step 4 (CLT). The Lindeberg–Lévy CLT applied to the bracket, combined with Step 3 and Slutsky's theorem, yields $\sqrt{N}(\widehat{\theta}^{\mathrm{GFD}} - \theta_0) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}(\mathbf{\Phi}^*))$ with $\mathbf{\Sigma}(\mathbf{\Phi}^*) = \mathbf{J}(\mathbf{\Phi}^*)^{-1}\mathbf{\Omega}(\mathbf{\Phi}^*)\mathbf{J}(\mathbf{\Phi}^*)^{-1}$. This is the stated sandwich form (Newey and McFadden, 1994, Ch. 6), with $\mathbf{\Omega}$ carrying the CCP plug-in correction only under Assumption 9; without it, $\mathbf{\Omega}$ additionally carries the transition-induced term of Step 2. □

F.7 Proof of Theorem 6 (FD-class trace optimality)

The trace map $\mathbf{q} \mapsto \operatorname{tr}(\boldsymbol{\Sigma}(\boldsymbol{\Phi}(\mathbf{q})))$ is continuous on the compact set $\mathcal{Q}$. The input map $\mathbf{q} \mapsto \boldsymbol{\Phi}(\mathbf{q})$ is affine. The sandwich $\boldsymbol{\Sigma} = \mathbf{J}^{-1}\boldsymbol{\Omega}\mathbf{J}^{-1}$ is rational in $\boldsymbol{\Phi}$, since $\mathbf{J}$ and $\boldsymbol{\Omega}$ are. Assumption 8 keeps $\mathbf{J}$ uniformly non-singular on $\mathcal{Q}$. The composition is therefore continuous. The extreme value theorem gives a minimizer $\mathbf{q}^{\mathrm{opt}} \in \mathcal{Q}$. By construction (35) the minimum is attained at $\boldsymbol{\Phi}^{\mathrm{opt}} = \boldsymbol{\Phi}(\mathbf{q}^{\mathrm{opt}})$. Hence $\operatorname{tr}(\boldsymbol{\Sigma}(\boldsymbol{\Phi}^{\mathrm{opt}}))$ equals the minimum over $\mathcal{F}_{\mathcal{Q}}$, the feasible sub-region of the theorem statement. □

F.8 Proof of Corollary 7 (conditional ρ-horizon semiparametric efficiency)

By (G1) and (G2), $\boldsymbol{\Phi}^{\mathrm{opt}}$ is well-defined and interior with $\mathbf{J}$ non-singular, so $\boldsymbol{\Sigma}(\boldsymbol{\Phi}^{\mathrm{opt}})$ is finite and Theorem 6 applies. Upper bound. By (G3) the efficient input $\boldsymbol{\Phi}^{\dagger}$ lies in the fixed feasible sub-region $\mathcal{F}_{\mathcal{Q}}$ of Theorem 6, with $\boldsymbol{\Sigma}(\boldsymbol{\Phi}^{\dagger}) = \boldsymbol{\Sigma}_{\rho}^{\mathrm{eff}}$. The Corollary's $\boldsymbol{\Phi}^{\mathrm{opt}}$ is the trace minimizer over that same $\mathcal{Q}$ (Theorem 6), so $\operatorname{tr}\boldsymbol{\Sigma}(\boldsymbol{\Phi}^{\mathrm{opt}}) \le \operatorname{tr}\boldsymbol{\Sigma}(\boldsymbol{\Phi}^{\dagger}) = \operatorname{tr}\boldsymbol{\Sigma}_{\rho}^{\mathrm{eff}}$. Lower bound. Every GFD estimator uses only the ρ-window joint law, hence lies in $\mathcal{E}_{\rho}$. By (G4) the feasible (plug-in) estimator has the same asymptotic variance as its oracle counterpart. The Chamberlain bound for $\mathcal{E}_{\rho}$ therefore applies to $\widehat{\theta}^{\mathrm{GFD}}(\boldsymbol{\Phi}^{\mathrm{opt}})$: $\boldsymbol{\Sigma}(\boldsymbol{\Phi}^{\mathrm{opt}}) \succeq \boldsymbol{\Sigma}_{\rho}^{\mathrm{eff}}$ in the Löwner order. Equality. The difference $\boldsymbol{\Sigma}(\boldsymbol{\Phi}^{\mathrm{opt}}) - \boldsymbol{\Sigma}_{\rho}^{\mathrm{eff}}$ dominates $\mathbf{0}$ in the Löwner order. Its trace is $\operatorname{tr}\boldsymbol{\Sigma}(\boldsymbol{\Phi}^{\mathrm{opt}}) - \operatorname{tr}\boldsymbol{\Sigma}_{\rho}^{\mathrm{eff}}$, which is non-positive by the upper bound. A positive-semidefinite matrix with non-positive trace is the zero matrix. Therefore $\boldsymbol{\Sigma}(\boldsymbol{\Phi}^{\mathrm{opt}}) = \boldsymbol{\Sigma}_{\rho}^{\mathrm{eff}}$. □

F.9 Proof of Corollary 11 (GFD identity under unobserved heterogeneity)

Sketch. Fix type ℓ and condition on $L_n = \ell$. The type-conditional submodel $(u_t^{(\ell)}, f_t^{(\ell)}, \theta^{(\ell)})$ is itself a non-stationary single-agent DDC. Apply the existence test (Theorem 3) and the value-difference identity (Theorem 4) to this submodel at (x_t, d_t, d_t'): pointwise FD at horizon ρ_ℓ^* delivers a flow $\boldsymbol{\Phi}^{(\ell)}(x_t)$ satisfying (52). Since this construction is independent across types, the mixture log-likelihood $\log \sum_\ell \pi_\ell \, p_t^{(\ell)}(d_{n,t} \mid x_{n,t})$ aggregates the type-conditional GFD representations. □

F.10 Proof of Theorem 12 (GFD-EM consistency)

Sketch. The argument has four steps.

Step 1 (uniform consistency of the anchor). Under (iii), with the oracle (true) posterior weights, the cell-frequency CCP estimator $\widehat{p}_t^{(\ell)}(d \mid x)$ at each $t \in [t_{\mathrm{start}} + 1, t_{\mathrm{end}}]$ satisfies $\widehat{p}_t^{(\ell)} = p_t^{(\ell)} + O_p(N^{-1/2})$ uniformly in (d, x) (McLachlan and Krishnan, 2008). The M-estimator argument of Step 2 uses this oracle-posterior anchor. The data-dependent EM posterior inherits the same uniform rate once $(\widehat{\theta}, \widehat{\pi})$ is consistent (Step 3): the posterior weights then converge and enter $\widehat{p}_t^{(\ell)}$ smoothly, so the oracle rate propagates by a standard stochastic-equicontinuity argument.

Step 2 (continuity and identification). The GFD-representation pseudo-log-likelihood

$$\widetilde{Q}_N(\theta, \pi; \widehat{p}) \;=\; N^{-1} \sum_{n,t} \log \sum_{\ell} \pi_\ell \, p_t^{(\ell)}\big(d_{n,t} \,\big|\, x_{n,t};\, \theta^{(\ell)}, \widehat{\boldsymbol{\Phi}}^{(\ell)}, \widehat{p}^{(\ell)}\big)$$

is continuous in (θ, π) on the compact set $\Theta \times \Delta^{L-1}$ by (iii), and by Step 1 converges uniformly to its population analogue $\widetilde{Q}(\theta, \pi)$. Under (ii) the population objective has a unique maximizer $(\boldsymbol{\theta}_0, \boldsymbol{\pi}_0)$ modulo label switching. The rank condition of (ii) identifies the type-conditional components $\{(\pi_\ell, p^{(\ell)})\}_\ell$ (pairwise distinctness ruling out an interior-of-simplex collapse of components); the maintained injectivity of $\theta^{(\ell)} \mapsto p_t^{(\ell)}(\cdot; \theta^{(\ell)})$ on Θ then identifies each structural $\theta^{(\ell)}$ from its component CCPs.

Step 3 (EM convergence). Standard EM theory (Wu, 1983, Theorem 3.1) delivers monotone increase of $\widetilde{Q}_N$ along iterations and convergence of the EM iterates to a stationary point, not necessarily the global maximizer. The consistency claim is for the arg max of $\widetilde{Q}_N$, which is consistent for (θ_0, π_0) by the standard argmax argument (compactness from (iii), and the uniform convergence and unique population maximizer of Step 2); we use multi-start to select it among the EM limit points.

Step 4 (asymptotic normality). Let $\widehat{\boldsymbol{\eta}} = (\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\pi}})$ and $\boldsymbol{\eta}_0$ the truth. Taylor-expand $\partial \widetilde{Q}_N / \partial \boldsymbol{\eta}$ at $\boldsymbol{\eta}_0$. By Step 1 the empirical anchor enters $\widetilde{Q}_N$ through a smooth GFD-anchor map with bounded derivatives, so $\widehat{\boldsymbol{\eta}}$ admits the asymptotically linear expansion

$$\sqrt{N}(\widehat{\boldsymbol{\eta}} - \boldsymbol{\eta}_0) = \mathbf{J}^{-1} \frac{1}{\sqrt{N}} \sum_{n=1}^{N} \left[\, \bar{s}_n^{\eta} + \mathbf{D}_p \, \boldsymbol{\psi}^{\mathrm{p}}(\zeta_n) \, \right] + o_p(1),$$

where $\bar{s}_n^{\eta}$ is the per-agent score in $\boldsymbol{\eta} = (\boldsymbol{\theta}, \boldsymbol{\pi})$, and the prior $\boldsymbol{\pi}$ is taken in free coordinates on the simplex interior Δ^{L-1} (so that $\mathbf{J}$, the information in these reduced coordinates, is non-singular). The CCP correction is itself asymptotically linear (Step 1, Assumption 7), so the bracket is a single mean-zero i.i.d. average across agents. The Lindeberg–Lévy CLT and Slutsky's theorem yield $\sqrt{N}(\widehat{\boldsymbol{\eta}} - \boldsymbol{\eta}_0) \Rightarrow N(0, \boldsymbol{\Sigma}^{\mathrm{GFD}})$ with $\boldsymbol{\Sigma}^{\mathrm{GFD}} = \mathbf{J}^{-1} \boldsymbol{\Omega} \, \mathbf{J}^{-1}$ and $\boldsymbol{\Omega} = \mathrm{Var}(\bar{s}_n^{\eta} + \mathbf{D}_p \, \boldsymbol{\psi}^{\mathrm{p}})$. As in the proof of Theorem 5, the influence function takes this clean plug-in form along the feasible manifold under Assumption 9; absent rank stability, $\boldsymbol{\Omega}$ additionally carries the transition-induced influence term of Theorem 5 (the off-manifold $\widehat{f}$-channel entering through Step 1's anchor). □