A Generalized Finite Dependence Framework for Dynamic Discrete Choice Models*

Yu Hao† Hiroyuki Kasahara‡ Katsumi Shimotsu§

June 7, 2026

Abstract

Altug and Miller (1998); Arcidiacono and Miller (2011, 2019) introduced finite dependence (FD) as a route to estimating dynamic discrete choice models without solving the Bellman equation, and demonstrated it on several structured examples. Whether ρ-step FD—in the generalized sense that admits signed, history-conditional per-step weights (a strict superset of the Markov-weight notion)—holds for an arbitrary transition kernel has remained open. **We close this gap**: pointwise ρ-step FD at every state-action triple is equivalent to a single sparse linear-feasibility test on the structural transitions, $\mathbf{b} \in \text{Range}(\mathbf{A})$, constructed from f alone (Theorem 3). The condition is operationally checkable before any estimation step. Three of the four canonical FD-amenable transition structures—renewal, shift register, and Kronecker—are recovered in closed form with explicit $\rho^\star$ formulas; the triple-connector / strategic-ripple class carries no closed form and is dispatched by direct invocation of the rank test at its smallest passing horizon (typically $\rho = 2$).

We develop the natural *generalized finite dependence* (GFD) estimator that takes any FD-feasible flow input as an explicit argument, attains the minimum trace of the asymptotic variance within the FD-feasible class on a compact regular domain (Theorem 6), and admits a Bellman-free fixed-point system for payoff-only counterfactual CCPs (Theorem 7). Monte Carlo evidence verifies the existence test on nine canonical DDC models at machine precision and demonstrates the GFD identity across single-agent and multi-agent counterfactual settings, including a non-stationary multi-agent regime in which the GFD identity yields a computable counterfactual from an empirical CCP anchor, without committing to a specific far-future trajectory as a nested fixed-point (NFXP) solve would require. The framework extends to unobserved heterogeneity through a GFD expectation–maximization

*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".

†Faculty of Business and Economics, The University of Hong Kong. Email: haoyu@hku.hk.

‡Vancouver School of Economics, The University of British Columbia. Email: hkasahar@mail.ubc.ca.

§Faculty of Economics, The University of Tokyo. Email: shimotsu@e.u-tokyo.ac.jp.

scheme that reuses the same flow-weight machinery (Section 3.9). An empirical application to a non-stationary entry game of the four major U.S. carriers (2010–2023) rejects entry-cost stationarity at the 0.1% level and shows that imposing stationarity misstates individual carrier-year entry probabilities by up to 50 percentage points (Section 5). Appendix D applies the framework to Altug and Miller (1998)'s female labour-supply model on PSID data and verifies the consistency-for-any-FD-input result against two FD-feasible flow inputs.

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

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

1 Introduction

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

Finite dependence (FD) circumvents the Bellman fixed point. If the future state distributions induced by two distinct current actions can be matched after a short horizon ρ, the continuation values cancel from the value-function difference, which reduces to a finite sum of flow utilities and conditional choice probabilities (CCPs; Hotz and Miller, 1993). The idea originates with Altug and Miller (1998), was extended to non-negative per-step weights at small ρ by Arcidiacono and Miller (2011), and was generalized to ρ-period per-step Markov weights by Arcidiacono and Miller (2019), who demonstrate the representation on a handful of structured examples: single-period renewals at $\rho = 1$, shift registers of memory p at $\rho = p$, and a few hand-constructed cases. The general question *for which transition kernels, and at which minimal horizon $\rho^\star$, does ρ-step FD admit a feasible flow representation?*—left explicit in their work—has remained open.

We close this gap. Our headline result (Theorem 3) reduces pointwise ρ-step FD at every triple (x_t, d_t, d'_t) to a single sparse linear-feasibility condition $\mathbf{b} \in \text{Range}(\mathbf{A})$, constructed from the transition kernel f alone with no appeal to the value function or CCPs. The condition is a necessary and sufficient test runnable before any estimation step. It unifies the previously hand-constructed FD-amenable cases: three of the four canonical transition classes—renewal, shift-register, and Kronecker with action-invariant exogenous factor—are recovered as special cases of Theorem 3 with explicit closed-form $\rho^\star$, while the fourth, the triple-connector / strategic-ripple class, carries no closed form and is dispatched by direct invocation of the rank test at its smallest passing horizon (typically $\rho = 2$). We verify the existence test on nine canonical DDC models at machine precision (Section 4.1). Once existence is established, the *generalized*

finite dependence (GFD) estimator of Section 3 takes any FD-feasible flow input as an explicit argument, attains the minimum trace of the asymptotic variance within the FD-feasible class, and admits a Bellman-free fixed-point system for payoff-only counterfactual CCPs.

A key advantage of GFD is its natural accommodation of non-stationary environments. In many labor economics applications—life-cycle models of occupational choice (James, 2014), education (Arcidiacono et al., 2016), or female labor supply and fertility (Altug and Miller, 1998; Gayle and Golan, 2012; Gayle et al., 2018)—the environment is inherently non-stationary due to finite horizons, aging, or aggregate shocks. Traditional full-solution methods struggle in these settings because the value function cannot be computed via a stationary Bellman contraction; researchers must often make strong assumptions about the length of the time horizon and the evolution of state variables far beyond the sample period. In contrast, finite dependence only requires matching state distributions over a short horizon ρ, completely bypassing the need to compute the continuation value V. As emphasized by Arcidiacono and Miller (2019), this allows for the estimation of non-stationary models with the same computational ease as stationary ones, making it a particularly powerful tool for labor market dynamics.

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

This paper makes four contributions: the existence/rank-test characterization, the payoff-only GFD counterfactual identity, FD-class trace optimality, and the structural and numerical infrastructure that makes those results operational.

The headline result is the existence test. Theorem 3 reduces pointwise ρ-step finite dependence at every triple (x_t, d_t, d_t') to a single sparse linear-feasibility condition $\mathbf{b} \in \text{Range}(\mathbf{A})$ on the structural transitions f alone—no value function, no CCPs, no estimation required. The test is necessary and sufficient, operationally checkable, and recovers the four previously hand-constructed FD-amenable classes: renewal at $\rho^\star = 1$, shift registers with p-lag memory at $\rho^\star = p$, and Kronecker structures with action-invariant exogenous factor at $\rho^\star = \rho^\star_{\text{endo}}$ (Lemma 11 in Appendix B) in closed form, with the triple-connector / strategic-ripple class carrying no closed form and verified directly by Theorem 3 at the typical horizon $\rho = 2$. Section 4.1 states the four cases as a classification of FD-amenable transition structures; the supporting constraint–decomposition equivalence (Theorem 2) underwrites the QP reformulation on which the test rests.

The remainder of the paper develops the natural downstream estimator and its uses. Section 3 introduces the *generalized finite dependence (GFD) estimator*, a two-step CCP estimator that takes any FD-feasible flow input $\widehat{\mathbf{\Phi}}$ as an explicit argument; it is consistent for any feasible choice (Theorem 5) and attains the minimum trace of the asymptotic variance within the FD-feasible

class on a compact regular sub-region at the trace-optimal $\widehat{\mathbf{\Phi}}^{\text{opt}}$ (Theorem 6). When finite dependence holds, payoff-only counterfactual CCPs satisfy a Bellman-free fixed-point system (Theorem 7); empirically the system recovers the Bellman MLE at $\sim 7 \times 10^{-7}$ sup-norm accuracy with a $\sim 10\times$ wall-clock speedup amortized over a sweep of counterfactual regimes (Section 3.8). Section 3.9 extends the framework to unobserved heterogeneity through a GFD expectation–maximization scheme. Appendix B collects the numerical infrastructure that makes the existence test and downstream estimator operational at large state spaces: reachability pruning, iterative LSQR for deep-tree flow QPs, KKT-LU batched solving, and the structured Kronecker solver that exploits action-invariant exogenous factors.

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. Prior to this paper, verifying finite dependence in general required problem-specific arguments: demonstrating a renewal action (Arcidiacono and Miller, 2011) or a terminal path structure (Arcidiacono and Miller, 2019) for each model specification—as in applied finite-dependence work on occupational choice (James, 2014) and agricultural land use (Scott, 2013). As a result, several influential empirical studies resort to full-solution methods precisely because their transition structures lack these special features: occupational-choice models with multi-dimensional human capital (Keane and Wolpin, 1997), industrial-dynamics models with action-dependent state persistence (Ryan, 2012), and sectoral mobility models where human capital depreciates at action-specific rates (Dix-Carneiro, 2014). Our contribution is to replace this case-by-case verification with a general algebraic framework that accommodates these settings. Arcidiacono and Miller (2020) show that counterfactual CCPs are generally not identified off short panels unless finite dependence holds, reinforcing the importance of our characterization for identification beyond estimation convenience.

The term "generalized finite dependence" is due to Gayle (2021), who introduces it for dynamic discrete and continuous choice models; we adopt the terminology and use "generalized" in the same spirit, for the flow parameterization over arbitrary dependence horizons $\rho \geq 1$ with signed weights. Relative to that work, the contributions specific to this paper are the arbitrary-ρ existence/rank characterization with closed-form $\rho^\star$ (Theorem 3), the convex-QP feasibility test for signed weights, and the FD-class trace-optimality result (Theorem 6). An earlier working paper by two of the present authors (Hao and Kasahara, 2024) treats the two-period case using sequential weight optimization with Kronecker structure; it does not supply an arbitrary-ρ existence characterization, general asymptotic theory for signed GFD weights, or the efficiency architecture developed here.

The closest methodological comparison is Kalouptsidi et al. (2021) (KSSR), who provide a linear-representation approach to counterfactual analysis in DDC models. KSSR's central result

establishes a null-space compatibility condition between a block-Kronecker matrix built from observed transitions and the Jacobian of the counterfactual payoff transformation. Our framework is complementary rather than substitutive: Theorem 7 gives a checkable computational route within the payoff-only class, with an explicit fixed-point algorithm and restart-based validation protocol once finite dependence is verified. Magnac and Thesmar (2002) establish that DDC models are generically underidentified absent normalization restrictions and study which restrictions (exclusion, parameter-sharing) restore identification at the level of the model + data distribution; our framework contributes a different operational route via finite dependence rather than "providing the constructive conditions" for the Magnac–Thesmar identification results, which rest on independent assumptions. Berry and Compiani (2023) develop IV approaches for dynamic models with serially correlated unobservables; our framework applies to the standard conditional-independence setting but extends its computational reach.

Recent work addresses DDC tractability through complementary margins. Bunting and Ura (2025) exploit index-invertibility of reduced-form CCPs whose indices combine flow-utility and transition components, and Blevins (2025) exploit continuous-time structure. Our framework contributes the rank-test characterization (Theorem 3) that applies to arbitrary transitions, the flow parameterization that converts weight-finding into a convex QP, and the efficiency result (Theorem 6). Our approach uses transition structure in discrete time. In dynamic games, Aguirregabiria and Marcoux (2021) emphasize equilibrium restrictions; our extension exploits player-specific transition decomposition to reduce computational burden under separability assumptions. Adusumilli and Eckardt (2025) use temporal-difference learning to approximate value functions in DDC models with continuous or high-dimensional state spaces, avoiding explicit computation of transition matrices; their approach works for any DDC model but introduces approximation error, while ours achieves exact elimination of continuation values when finite dependence holds. Chen (2025) develops model-adaptive sieve approximations to the Bellman equation that improve convergence rates in large state spaces. Both approaches approximate value functions more efficiently; finite dependence eliminates them entirely. The closest Bellman-free alternative is the Euler-equation approach of Aguirregabiria and Magesan (2018), which replaces the Bellman fixed point with a system of local first-order (Euler) conditions in CCP space that still bind the policy through derivatives of the continuation value; GFD instead cancels the continuation value *exactly* over a finite horizon via signed transition weights, eliminating it rather than characterizing its local variation.

The use of negative weights connects to the broader econometrics literature on signed measures. Borusyak and Hull (2024) argue that negative weights are not a concern in design-based specifications (which include shift-share designs), because the target estimand is a convex combination of treatment effects with positive ex-ante weights even when realized weights can be negative. de Chaisemartin and D'Haultfœuille (2020) demonstrate that in two-way fixed-effects estimators, negative weights on heterogeneous treatment effects can be problematic because the

estimand becomes a non-convex combination. Within synthetic control methods, Abadie et al. (2019) likewise do not require donor weights to be non-negative when computing the synthetic match, and they provide a recursive procedure for the weights; the non-linearity of that recursion makes the computation expensive. The GFD parallel is to keep the signed-weight feasible set but to replace the non-linear recursion with a single *linear* feasibility test (Theorem 3): the constraint matrix is shared across initial states, so the entire weight-search problem collapses to one factorisation applied to all right-hand sides at once (Section 4.2.1). Our setting differs in kind from the design-based and TWFE literatures: negative weights are applied to structural transition primitives—probability distributions over future states—not to heterogeneous treatment effects. The finite dependence condition ensures exact cancellation of the continuation value, so the resulting estimand is the true payoff difference regardless of weight signs. Under approximate finite dependence, signed weights amplify the approximation bias through the weight-norm factor, creating a variance concern analogous to weak instruments but not a sign-reversal concern.

The asymptotic theory of the GFD estimator builds on standard two-step M-estimation (Murphy and Topel, 1985; Newey and McFadden, 1994); the variance-input mapping connects to Newey (1990)'s efficient instrumental-variables framework. The counterfactual computation route of Theorem 7 intersects with the single-agent nested pseudo-likelihood estimator of Aguirregabiria and Mira (2002), the dynamic-game nested pseudo-likelihood (NPL) framework of Aguirregabiria and Mira (2007), the asymptotic least-squares estimator of Pesendorfer and Schmidt-Dengler (2008), the forward-simulation estimator of Hotz et al. (1994), and the two-step game estimators of Bajari et al. (2007); the GFD QP/least-squares flow-weight machinery is a structurally constrained instance of the asymptotic-least-squares class. The unobserved-heterogeneity extension of Section 3.9 interfaces with the nonparametric mixture identification of Kasahara and Shimotsu (2009). The control-theoretic provenance of the controllability recursion follows the discrete-time free-switching reachability literature (Liberzon, 2003; Sun and Ge, 2005).

The remainder of the paper proceeds as follows. Section 2 introduces the model and the generalized finite-dependence framework. Section 3 develops the CCP estimator, the payoff-only counterfactual GFD identity (Section 3.8), and the unobserved-heterogeneity extension (Section 3.9). 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 nine canonical DDC models, Monte Carlo evidence on a single-agent investment model and a three-player Markov-perfect entry/exit game, and counterfactual demonstrations that include a non-stationary multi-agent setting in which the GFD identity is computable under future-trajectory ambiguity. Section 5 takes the framework to data, estimating a non-stationary entry game of the four major U.S. carriers over 2010–2023: a Wald test rejects entry-cost stationarity at the 0.1% level for every carrier, and a stationary estimator is shown to misstate

individual carrier-year entry probabilities by up to 50 percentage points. Section 6 concludes. Appendix D applies the GFD framework to Altug and Miller (1998)'s 14-parameter female labour-supply model on PSID data. Appendix B collects the scalable computation machinery.

2 Existence of Finite Dependence

2.1 Dynamic Discrete Choice Framework

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

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

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

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 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 $\mathcal{D}$ and 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 $G(\epsilon)$.[1][2]

The proof is in Appendix E.2.

The representation (3) holds for *any* weights summing to unity, including signed weights—the structural opening that generalized finite dependence exploits. Generalized finite dependence extends this by parameterizing the joint flow Φ, which serves as the ρ-period version 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') as the existence of Markov per-step weights $\omega_{t+\tau}(d_{t+\tau} \mid x_{t+\tau})$, $\tau = 1, \ldots, \rho$, satisfying $\sum_{d_{t+\tau}} \omega_{t+\tau}(d_{t+\tau} \mid x_{t+\tau}) = 1$, such that the two paths starting from (x_t, d_t) and (x_t, d_t') reach the same terminal state distribution after ρ steps.

To see how the weight enters the value-function difference, consider the ρ-period expansion of the choice-specific value function $v_t(x_t, d_t)$. Under a sequence of Markov per-step weights $\omega_{t+1}, \ldots, \omega_{t+\rho}$, the joint probability of a length-ρ continuation path admits the explicit product form

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

[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); and 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 and follows from multiplying Arcidiacono and Miller's eq. (2.4) by $\omega(d_t \mid x_t)$ and summing over d_t.

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

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

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). Arcidiacono and Miller (2019) note the possibility of conditioning $\omega_{t+\tau}$ on additional lags but do not pursue this direction operationally: with per-step weights, conditioning on a longer history makes the terminal-distribution-matching equation a multivariate polynomial in $\{\omega_{t+\tau}\}$, which is computationally intractable to solve directly. We acknowledge that 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. In plain language, the agent "knows the past but not the future" at the moment $d_{t+\tau}$ is chosen; the weight can reference any state realization or decision strictly before time $t+\tau+1$ but cannot peek at $x_{t+\tau+1:t+\rho}$ or $d_{t+\tau:t+\rho}$. The same product structure as in (5) carries through:

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

The AM Markov form is recovered by dropping the prefix $(x_{t:t+\tau-1}, d_{t:t+\tau-1})$ from the conditioning. The generalization strictly enlarges the set of triples (x_t, d_t, d'_t) at which finite dependence can be achieved, because history-conditional weights can offset history-dependent transition dynamics that no 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)**.** *Non-negative weights restrict the achievable terminal distribution to the* convex hull *of* $\{\boldsymbol{\pi}(d_{t+1:\rho})\}_{d_{t+1:\rho}}$*; signed weights expand it to the* affine hull*, unlocking finite dependence where no convex combination of trajectories 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, the matching objective is quadratic, and the problem becomes a convex QP that can be solved efficiently and globally.

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

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

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 $d_{t+\tau:t+\rho}$. This *non-anticipativity* is precisely what the linear constraints below enforce.

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.

(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:

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

(ii) Flow conservation. For each $\tau \in \{1, \dots, \rho - 1\}$, 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 signed probability 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) *and the flow-conservation constraint* (11) *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. On every such history, 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 E.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.*

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

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

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 dramatically when transitions are sparse; Appendix B) and prevents the flow-conservation constraints (11) from being vacuously satisfied on impossible histories.

2.5 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

$$\phi_{x_{t+1},\, d_{t+1}}(x_t, d_t) \;=\; f_t(x_{t+1} \mid x_t, d_t) \cdot \omega_{t+1}(d_{t+1} \mid x_t, x_{t+1}, d_t), \qquad x_{t+1} \in \{l, h\},\ d_{t+1} \in \{0, 1\}, \tag{14}$$

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

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

where the conditioning argument of f_t in the transition matrix is $x_{t+1} \in \{l, h\}$: columns 1–2 condition on $x_{t+1} = l$ and columns 3–4 condition on $x_{t+1} = h$, matching the flow-vector entries $\phi_{x_{t+1},d_{t+1}}(x_t, d_t)$ on the right. and the QP (13) matches $\kappa_{t+2}(\cdot \mid x_t, d_t)$ against $\kappa_{t+2}(\cdot \mid x_t, d_t')$. Equation (14) is the simplest instance of the flow decomposition that drives the entire framework: each ϕ-entry is the per-state arrival probability scaled by the time-1 per-step weight, and matching terminal distributions across the two time-0 actions reduces to a finite-dimensional linear-quadratic problem.

2.6 Existence Test for Finite Dependence

This subsection states the paper's headline 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 (renewal, shift-register, Kronecker, triple-connector) as closed-form corollaries in 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 identical terminal distributions:

$$\kappa_{t+\rho+1}(x_{\rho+1} \mid x_t, d_t) \;=\; \kappa_{t+\rho+1}(x_{\rho+1} \mid x_t, d_t') \qquad \forall\, x_{\rho+1} \in \mathcal{X}. \tag{15}$$

We now construct an explicit linear system whose solvability is 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 $(x_{t+1:\rho}, d_{t+1:\rho})$.

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 $x_{t+1} \in \mathcal{X}$:

$$(\mathbf{A}_{\text{init}})_{x_{t+1},\,(x_{t+1:\rho},\,d_{t+1:\rho})} \;=\; \mathbf{1}\{(\text{first state of } x_{t+1:\rho}) = x_{t+1}\}. \tag{16}$$

Flow-conservation (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 $x_{t+\tau+1} \in \mathcal{X}$,

$$\begin{aligned} &\underbrace{\sum_{x_{t+\tau+2:t+\rho}} \sum_{d_{t+\tau+1:t+\rho}} \phi_{x_{t+1:t+\rho},\,d_{t+1:t+\rho}}(x_t, d_t)}_{\text{flow through } (x_{t+1:t+\tau+1},\,d_{t+1:t+\tau})} \\ &\quad - \; f_{t+\tau}(x_{t+\tau+1} \mid x_{t+\tau}, d_{t+\tau}) \underbrace{\sum_{x_{t+\tau+1:t+\rho}} \sum_{d_{t+\tau+1:t+\rho}} \phi_{x_{t+1:t+\rho},\,d_{t+1:t+\rho}}(x_t, d_t)}_{\text{flow through } (x_{t+1:t+\tau},\,d_{t+1:t+\tau})} \;=\; 0. \end{aligned} \tag{17}$$

Each instance of (17) is one linear equation in $\operatorname{vec}\Phi$, with coefficients $+1$ 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—so $\mathbf{A}_{\text{nac}} \in \mathbb{R}^{C \times S^\rho D^\rho}$ with

$$C \;:=\; \text{number of NAC rows} \;=\; \sum_{\tau=1}^{\rho-1} S^{\tau+1} D^\tau. \tag{18}$$

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

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

$$(\mathbf{T}_{\text{last}})_{x',\,(x_{t+1:\rho},\,d_{t+1:\rho})} \;=\; f_{t+\rho}(x' \mid x_\rho, d_\rho). \tag{19}$$

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

$$\mathbf{A} := \begin{pmatrix} \mathbf{A}_{\text{init}} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\text{init}} \\ \mathbf{A}_{\text{nac}} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\text{nac}} \\ \mathbf{T}_{\text{last}} & -\mathbf{T}_{\text{last}} \end{pmatrix}, \qquad \mathbf{b} := \begin{pmatrix} \mathbf{f}^{(d_t)} \\ \mathbf{f}^{(d'_t)} \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix}. \tag{20}$$

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* (20)*. Then pointwise ρ-period finite dependence at (x_t, d_t, d'_t) (Definition 1) holds if and only if*

$$\mathbf{b} \in \text{Range}(\mathbf{A}), \qquad \textit{equivalently} \qquad \text{rank}\big([\,\mathbf{A} \mid \mathbf{b}\,]\big) = \text{rank}(\mathbf{A}). \tag{21}$$

Whenever (21) *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 E.3.

Remark 3 (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 $(3S + 2C)\times 2NM$ entries with row count $3S+2\sum_{\tau=1}^{\rho-1} S^{\tau+1}D^\tau$; both $\mathbf{A}$ and $\mathbf{b}$ are constructed from f and (x_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)—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, universal FD is equivalent to pointwise feasibility of $\mathbf{A}^{(x_t,d_t,d'_t)}\,\mathbf{x} = \mathbf{b}^{(x_t,d_t,d'_t)}$ at every triple, which can be verified by running the LSQR test of Remark 3 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: the *Generalized Finite Dependence (GFD) estimator* is a two-step CCP estimator of the structural payoff parameter θ that takes any FD-feasible flow-matrix collection $\widehat{\mathbf{\Phi}}$ satisfying

the test of Section 2.6 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, generated by the DDC model of Section 2 at an unknown structural parameter $\theta_0 \in \Theta \subseteq \mathbb{R}^{d_\theta}$, with flow payoff $u_t(x_t, d_t; \theta)$ for $d_t \in \mathcal{D}$ and reference normalization $u_t(x_t, 0; \theta) \equiv 0$. 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 $\psi_{d_t}(\mathbf{p}_t(x_t)) = \gamma_E - \ln p_t(d_t \mid x_t)$ and the equilibrium CCP is the multinomial logit $p_t(d_t \mid x_t) = \exp(v_t(x_t, d_t)) / \sum_{d_t''} \exp(v_t(x_t, d_t''))$.*

First-stage estimates. From the panel, the researcher constructs $\sqrt{N}$-consistent nonparametric estimates $\widehat{f}_t(x_{t+1} \mid x_t, d_t)$ and $\widehat{p}_t(d_t \mid x_t)$ (cell-frequency, sieve, or kernel). 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*

$$\widehat{\mathbf{\Phi}} \;=\; \big\{\, \widehat{\Phi}(x_t, d_t) \;:\; x_t \in \mathcal{X},\; d_t \in \mathcal{D} \,\big\} \tag{22}$$

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 $d_t \in \mathcal{D}$*:

$$\widehat{\kappa}_{t+\rho+1}(x_{t+\rho+1} \mid x_t, d_t) \;=\; \widehat{\kappa}_{t+\rho+1}(x_{t+\rho+1} \mid x_t, d_t'), \qquad \forall\, d_t, d_t' \in \mathcal{D},\; \forall\, x_{t+\rho+1} \in \mathcal{X}. \tag{23}$$

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

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—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 trivial flush) in §4.1. Consistency holds for any FD-feasible $\widehat{\mathbf{\Phi}}$; variance depends on the specific input.

3.3 The GFD Value-Difference Identity

The GFD estimator builds on a closed-form representation of the value-function difference $v_t(x_t, d_t; \theta) - v_t(x_t, 0; \theta)$ 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* $d_t \neq 0$*,*

$$\begin{aligned} v_t(x_t, d_t; \theta) \; - \; v_t(x_t, 0; \theta) \; &= u_t(x_t, d_t; \theta) \\ &\quad + \sum_{\tau=1}^{\rho} \beta^{\tau} \sum_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}} \Delta\phi_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}}(x_t)\, \mu_{t+\tau}(x_{t+\tau}, d_{t+\tau}; \theta), \end{aligned} \tag{24}$$

with $\Delta\phi_{x_{t+1:t+\rho}, d_{t+1:t+\rho}}(x_t) := \phi_{x_{t+1:t+\rho}, d_{t+1:t+\rho}}(x_t, d_t) - \phi_{x_{t+1:t+\rho}, d_{t+1:t+\rho}}(x_t, 0)$*, and* $\mu_t(x_t, d_t; \theta) := u_t(x_t, d_t; \theta) + \psi_{d_t}(\mathbf{p}_t(x_t))$ *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 proof is in Appendix E.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 (24):

$$\begin{aligned} \widetilde{v}^{\rho\text{-FD}}(x_t, d_t;\, \theta, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) \; &:= u_t(x_t, d_t; \theta) \\ &\quad + \sum_{\tau=1}^{\rho} \beta^{\tau} \sum_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}} \Delta\widehat{\phi}_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}}(x_t)\, \widehat{\mu}_{t+\tau}(x_{t+\tau}, d_{t+\tau}; \theta), \end{aligned} \tag{25}$$

with $\Delta\widehat{\phi}$ defined as in (24) (sample analog) and $\widehat{\mu}_t(x_t, d_t; \theta) := u_t(x_t, d_t; \theta) + \psi_{d_t}(\widehat{\mathbf{p}}_t(x_t))$. with 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

$$\Lambda_{d_t}(x_t;\, \theta, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) \; := \; \frac{\exp\big(\widetilde{v}^{\rho\text{-FD}}(x_t, d_t;\, \theta, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}})\big)}{\sum_{d_t'' \in \mathcal{D}} \exp\big(\widetilde{v}^{\rho\text{-FD}}(x_t, d_t'';\, \theta, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}})\big)}. \tag{26}$$

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

$$\widehat{\theta}^{\mathrm{GFD}}\big(\widehat{\mathbf{\Phi}}\big) \;:=\; \arg\max_{\theta\in\Theta}\; \frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T} \log \Lambda_{d_{it}}\big(x_{it};\, \theta, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}\big). \tag{27}$$

What $\widehat{\theta}^{\mathrm{GFD}}$ 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 $\widetilde{v}^{\rho\text{-FD}}(x_t, d_t;\, \theta_0, \mathbf{\Phi}, \mathbf{p}) = v_t(x_t, d_t; \theta_0) - v_t(x_t, 0; \theta_0)$, so $\Lambda_{d_t}(x_t;\, \theta_0, \mathbf{\Phi}, \mathbf{p}) = p_{0,t}(d_t \mid x_t)$ at the true parameter and the population GFD log-likelihood (27) is the correctly specified conditional multinomial log-likelihood. The GFD estimator inherits the standard maximum-likelihood-estimator (MLE) consistency: $\widehat{\theta}^{\mathrm{GFD}}(\widehat{\mathbf{\Phi}}) \xrightarrow{p} \theta_0$ for any FD-feasible $\widehat{\mathbf{\Phi}}$ and any $\sqrt{N}$-consistent first-stage $(\widehat{f}, \widehat{p})$.

What $\widehat{\mathbf{\Phi}}$ does in the estimator. The flow input $\widehat{\mathbf{\Phi}}$ enters (27) through the GFD pseudo value-difference (25), which packs the ρ-period continuation cancelation into a finite weighted sum over paths. *The estimator does not solve the Bellman equation at any θ.* Different feasible choices of $\widehat{\mathbf{\Phi}}$ produce different $\widehat{\theta}^{\mathrm{GFD}}(\widehat{\mathbf{\Phi}})$, all consistent for θ_0 but with different asymptotic variances; choosing $\widehat{\mathbf{\Phi}}$ to minimize the asymptotic-variance trace is the subject of Section 3.7.

3.5 The Two-Step GFD Procedure

Putting the pieces together, the full GFD procedure is:

Step 1 (*First stage*). From the panel $\{(x_{it}, d_{it})\}_{i,t}$, compute nonparametric estimates $\widehat{f}_t(x_{t+1} \mid x_t, d_t)$ and $\widehat{p}_t(d_t \mid x_t)$ together with the plug-in Hotz–Miller correction $\psi_{d_t}(\widehat{\mathbf{p}}_t(x_t))$.

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

$$\min_{\{\Phi(x_t,d_t)\}_{d_t\in\mathcal{D}}} \sum_{d_t=1}^{D-1} \sum_{x_{t+\rho+1}\in\mathcal{X}} \Big(\widehat{\kappa}_{t+\rho+1}(x_{t+\rho+1} \mid x_t, d_t) - \widehat{\kappa}_{t+\rho+1}(x_{t+\rho+1} \mid x_t, 0)\Big)^2 \tag{28}$$

subject to (10)–(11) (with $\widehat{f}$) for each $\Phi(x_t, d_t)$. By Theorem 3 (applied pairwise against the reference $d_t = 0$), the minimum is zero whenever pointwise FD holds at each triple $(x_t, d_t, 0)$. The default *canonical input* is the Moore–Penrose minimum-ℓ_2-norm solution.

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 (25). Solve (27) for $\widehat{\theta}^{\mathrm{GFD}}(\widehat{\mathbf{\Phi}})$ by Newton or quasi-Newton iteration. Section 3.6 below records the simplification that obtains when the flow payoff is linear in θ.

Computational cost. Step 1 is $O(NT)$ in the sample size. Step 2 is dominated by $|\mathcal{X}|$ sparse linear solves of dimension $D \cdot S^\rho D^\rho$ each, scaling roughly as $|\mathcal{X}| \cdot S^\rho D^\rho$ when LSQR is used and sparsity is exploited; this cost is paid once and does not repeat across parameter evaluations. Step 3 is a low-dimensional (d_θ-variable) concave optimization (logit-type with respect to the index $\widetilde{v}^{\rho\text{-FD}}$). 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 (when feasible) gives a third. All are consistent for θ_0, but the asymptotic variance $\mathbf{\Sigma}(\widehat{\mathbf{\Phi}})$ depends explicitly on the input. The *variance–input* mapping $\widehat{\mathbf{\Phi}} \mapsto \mathbf{\Sigma}(\widehat{\mathbf{\Phi}})$ is the central object of Section 3.7.

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

When the flow payoff is linear in θ, the GFD pseudo value-difference (25) 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:*

$$u_t(x_t, d_t; \theta) \;=\; z_t(x_t, d_t)^\top \theta, \qquad d_t \in \mathcal{D}, \tag{29}$$

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 (25) and grouping the θ-dependent and θ-independent terms gives

$$\widetilde{v}^{\rho\text{-FD}}(x_t, d_t;\, \theta, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) \;=\; \widetilde{H}_t(x_t, d_t;\, \widehat{\mathbf{\Phi}})^\top \theta \;+\; \widetilde{h}_t(x_t, d_t;\, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}), \tag{30}$$

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

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

$$\widetilde{h}_t(x_t, d_t;\, \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) \;:=\; \sum_{\tau=1}^{\rho} \beta^\tau \sum_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}} \Delta\widehat{\phi}_{x_{t+1:t+\rho},\, d_{t+1:t+\rho}}(x_t) \;\cdot\; \psi_{d_{t+\tau}}(\widehat{\mathbf{p}}_{t+\tau}(x_{t+\tau})), \tag{32}$$

where the path-flow difference is

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

the contrast between analyzed action d_t and reference action 0 at each length-ρ continuation path. With this decomposition the GFD log-likelihood (27) becomes a standard logit log-likelihood with generated regressor $\widetilde{H}_t$ and generated offset $\widetilde{h}_t$:

$$\widehat{\theta}^{\mathrm{GFD}}\big(\widehat{\mathbf{\Phi}}\big) \;=\; \arg\max_{\theta\in\Theta}\; \frac{1}{NT}\sum_{i,t}\log\frac{\exp\big(\widetilde{H}_t(x_{it},d_{it})^\top\theta+\widetilde{h}_t(x_{it},d_{it})\big)}{\sum_{d_t''}\exp\big(\widetilde{H}_t(x_{it},d_t'')^\top\theta+\widetilde{h}_t(x_{it},d_t'')\big)}, \tag{33}$$

which is concave in θ and admits standard Newton–Raphson or BFGS optimization. The asymptotic variance $\mathbf{\Sigma}(\widehat{\mathbf{\Phi}})$ studied in the next subsection depends on the flow input $\widehat{\mathbf{\Phi}}$ through $\widetilde{H}_t$.

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—numerical illustration is deferred to the Monte Carlo of Section 4.

Null-space parameterization of FD-feasible inputs. The joint feasibility QP (28) 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 (23) 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

$$\mathbf{\Phi}(x_t;\mathbf{q}) \;=\; \mathbf{\Phi}^{\mathrm{MP}}(x_t;f)+\mathbf{N}(x_t;f)\,\mathbf{q}, \tag{34}$$

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: $\sqrt{N}\,(\widehat{f}-f,\ \widehat{p}-p) = N^{-1/2}\sum_{i=1}^{N}\boldsymbol{\psi}^{\mathrm{fs}}(z_i)+o_p(1)$ for an influence function $\boldsymbol{\psi}^{\mathrm{fs}}$ with $\mathbb{E}[\boldsymbol{\psi}^{\mathrm{fs}}]=\mathbf{0}$ and finite second moment.*

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: $\inf_{\mathbf{q}\in\mathcal{Q}}\sigma_{\min}\big(\mathbf{J}(\mathbf{\Phi}(\mathbf{q}))\big)>0$.*

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 $\mathbf{J}(\mathbf{\Phi}^*) := -\mathbb{E}[\nabla^2_\theta \log\Lambda_{d_t}(x_t;\theta_0,\mathbf{\Phi}^*,\mathbf{p})]$ is non-singular. Then $\widehat{\theta}^{\mathrm{GFD}}(\widehat{\mathbf{\Phi}}) \xrightarrow{p} \theta_0$ for* every *such $\widehat{\mathbf{\Phi}}$, and*

$$\sqrt{N}\big(\widehat{\theta}^{\mathrm{GFD}}(\widehat{\mathbf{\Phi}})-\theta_0\big) \;\xrightarrow{d}\; \mathcal{N}\big(\mathbf{0},\,\mathbf{\Sigma}(\mathbf{\Phi}^*)\big), \tag{35}$$

where $\boldsymbol{\Sigma}(\boldsymbol{\Phi}) := \mathbf{J}(\boldsymbol{\Phi})^{-1}\,\mathbf{V}(\boldsymbol{\Phi})\,\mathbf{J}(\boldsymbol{\Phi})^{-1}$ *and* $\mathbf{V}(\boldsymbol{\Phi})$ *is the variance of the GFD score augmented by the* first-stage CCP *plug-in correction only. The estimation error in* $\widehat{\boldsymbol{\Phi}}$*—and hence in* $\widehat{f}$*, which enters the criterion only through* $\widehat{\boldsymbol{\Phi}}$*—is first-order negligible: because the GFD value-difference identity (Theorem 4) holds* exactly *at the true* f *for every FD-feasible* $\boldsymbol{\Phi}$*, the population score has mean zero at* θ_0 *for every such* $\boldsymbol{\Phi}$*, so* $\partial_{\boldsymbol{\Phi}}\,\mathbb{E}[\nabla_\theta \log \Lambda_{d_t}(x_t;\theta_0,\boldsymbol{\Phi},\mathbf{p})] = \mathbf{0}$*, a Neyman-orthogonality property that justifies treating* $\boldsymbol{\Phi}$ *as a fixed design in* $\boldsymbol{\Sigma}$*.*

The proof is in Appendix E.6.

Composing (34) with $\boldsymbol{\Sigma}(\cdot)$ gives the smooth (rational) *variance–input mapping* $\mathbf{q} \mapsto \boldsymbol{\Sigma}(\boldsymbol{\Phi}(\mathbf{q}))$. Asymptotic efficiency among GFD estimators is the problem of minimizing its trace.

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

$$\boldsymbol{\Phi}^{\mathrm{opt}} \;:=\; \boldsymbol{\Phi}^{\mathrm{MP}} + \mathbf{N}\,\mathbf{q}^{\mathrm{opt}}, \qquad \mathbf{q}^{\mathrm{opt}} \;:=\; \arg\min_{\mathbf{q}\in\mathcal{Q}} \operatorname{tr}\big(\boldsymbol{\Sigma}(\boldsymbol{\Phi}(\mathbf{q}))\big) \tag{36}$$

exists and attains $\operatorname{tr}(\boldsymbol{\Sigma}(\boldsymbol{\Phi}^{\mathrm{opt}})) = \min_{\boldsymbol{\Phi}\in\mathcal{F}(\cdot;f)\cap\mathcal{Q}} \operatorname{tr}(\boldsymbol{\Sigma}(\boldsymbol{\Phi}))$*. The result is the existence and attainment of a trace-minimizing flow input over the compact feasible sub-region* $\mathcal{F}(\cdot;f)\cap\mathcal{Q}$*; it does not assert optimality over the full FD-feasible affine variety or a semiparametric efficiency bound. Beyond FD-feasibility and the regularity of* $\mathcal{Q}$ *(Assumption 8), no controllability, rank, or instrument-Jacobian conditions are imposed.*

The proof is in Appendix E.7.

Sample-feasible procedure. The sample variant of (36) replaces f with $\widehat{f}$ and uses a plug-in sandwich $\widehat{\boldsymbol{\Sigma}}$ evaluated at a preliminary $\widehat{\theta}^{(0)} := \widehat{\theta}^{\mathrm{GFD}}(\widehat{\boldsymbol{\Phi}}^{\mathrm{MP}})$, followed by quasi-Newton minimization in $\mathbf{q}$ and 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.

3.8 Counterfactual Computation via Finite Dependence

The GFD identity (24) also delivers a route to counterfactual choice probabilities that does not solve the Bellman equation. We focus on *payoff-only* counterfactuals—those that modify the structural payoff parameter from θ_0 to θ_1 while holding the transition kernel f_t fixed—and write $\mathbf{p}_{1,t}(x_t) := \mathbf{p}_t(x_t;\,\theta_1, f)$ for the new equilibrium CCP.

Theorem 7 (GFD counterfactual identity: payoff-only counterfactual CCPs from finite dependence)**.** *Suppose Assumptions 1–5 hold and that* $\boldsymbol{\Phi}$ *is FD-feasible at horizon* ρ *under* f *(Definition 2). Then* $\mathbf{p}_{1,t}$ *satisfies the fixed-point equation*

$$p_{1,t}(d_t \mid x_t) \;=\; \Lambda_{d_t}\big(x_t;\,\theta_1, \boldsymbol{\Phi}, \mathbf{p}_1\big), \qquad \forall\,(x_t,d_t) \in \mathcal{X}\times\mathcal{D}, \tag{37}$$

where Λ_{d_t} is the GFD multinomial-logit map (26) *evaluated at the counterfactual payoff. Uniqueness of $\mathbf{p}_{1,t}$ is inherited from the Bellman contraction at θ_1 (via the Hotz–Miller bijection); the FD iteration $\mathbf{p}^{(m+1)} = \Lambda(\widetilde{v}^{\rho\text{-}FD}(\cdot;\theta_1,\mathbf{\Phi},\mathbf{p}^{(m)}))$ 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 E.5; the multi-agent extension is stated as Proposition 8 below. In practice we recommend restarting (37) from several initializations and rejecting the run if the limits disagree; the spectral-radius check at the converged point flags cases where ρ needs to be increased.[3]

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 because, unlike the single-agent case, existence and uniqueness of the counterfactual equilibrium are not guaranteed by a Bellman contraction and must be assumed (as is standard in the empirical-games literature).

Assumption 9 (Counterfactual MPE)**.** *At the counterfactual payoff θ_1 the game admits a Markov-perfect equilibrium with CCP profile $\mathbf{p}_1 = \{\mathbf{p}_{1,i}\}_i$, and the analyst selects one such profile (e.g. by the restart protocol below).*

Proposition 8 (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 9, the selected counterfactual profile $\mathbf{p}_1$ solves the per-player GFD fixed point evaluated at the equilibrium perceived transitions,*

$$p_{1,i,t}(d \mid x) \;=\; \Lambda_d^{(i)}\big(x;\,\theta_{1,i},\mathbf{\Phi}_i,\mathbf{p}_1\big), \qquad \forall\, i,\ (x,d), \tag{38}$$

and no Bellman solve is required at any candidate θ_1.

The proof (Appendix E.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 (38). Assumption 9 supplies the existence/selection that the single-agent contraction argument cannot; in practice the restart protocol both selects among equilibria and flags non-uniqueness.

The GFD counterfactual identity is complementary to existing identification results: in short panels counterfactual CCPs are generically unidentified unless finite dependence holds (Arcidiacono and Miller, 2020), and Kalouptsidi et al. (2021) give a null-space compatibility

[3] A worked numerical example on the 3-player entry/exit DGP of Section 4.2.2, with five restarts agreeing to maximum pairwise difference 8.7×10^{-7} and matching the NFXP MPE to 8.2×10^{-3}, accompanies the replication package.

condition for general (transition-changing) counterfactuals. Theorem 7 stays within the payoff-only class but gives a checkable computational route once FD is established.

Empirical demonstrations. Theorem 7 is demonstrated empirically on four DGPs in Section 4.3: single-agent Type-K (Investment) and single-agent Type-D (Entry/Exit) in §4.3.1, and multi-agent stationary MPE together with multi-agent non-stationary under future ambiguity in §4.3.2.

3.9 Extension to Unobserved Heterogeneity

Many applied DDC models feature unobserved heterogeneity, with agents of unobserved type $\ell \in \{1, \ldots, L\}$ differing in preferences $\theta^{(\ell)}$ or transitions $f^{(\ell)}$, or both. The observed CCP is the mixture $p_t(d_t \mid x_t) = \sum_\ell \pi_\ell \, p_{\ell,t}(d_t \mid x_t)$. Since the FD existence test depends only on f, a single flow input $\widehat{\mathbf{\Phi}}$ works across types when transitions are common; when transitions differ, the test of Theorem 3 is applied per type and the model runs at $\rho = \max_\ell \rho^\star_\ell$ with type-specific $\widehat{\mathbf{\Phi}}^{(\ell)}$.

GFD-EM algorithm. 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 (E) refreshing the posterior type weights $q_{i\ell}^{(m)} \propto \widehat{\pi}_\ell^{(m-1)} \prod_t \widehat{p}_\ell^{(m-1)}(d_{it} \mid x_{it})$ together with the type-specific Hotz–Miller offsets $\widehat{h}^{(\ell),(m)}$; (M) maximising the type-weighted GFD pseudo-log-likelihood $\widehat{\theta}^{(m)} = \arg\max_\theta \sum_{i,t,\ell} q_{i\ell}^{(m)} \log \Lambda_{d_{it}}(x_{it}; \theta^{(\ell)}, \widehat{\mathbf{\Phi}}^{(\ell)}, \widehat{\mathbf{p}}_\ell^{(m-1)})$; and (U) updating the mixture weights $\widehat{\pi}_\ell^{(m)} = N^{-1} \sum_i q_{i\ell}^{(m)}$ and type-specific CCPs via Λ at the new $(\widehat{\theta}^{(m)}, \widehat{\mathbf{\Phi}}^{(\ell)})$.

Computational advantage. The effective regressor $\widehat{H}^{(\ell)}$ depends on $(\widehat{f}^{(\ell)}, \widehat{\mathbf{\Phi}}^{(\ell)})$ but *not* on the iterating CCPs, so it is constructed once before the EM loop and reused across iterations; only the Hotz–Miller offset $\widehat{h}^{(\ell),(m)}$ is refreshed each iteration. AM-EM, by contrast, recomputes the matrix inversion $(I - \beta F_{\widehat{\mathbf{p}}})^{-1}$ per iteration (an $O(|\mathcal{X}|^3)$ operation) because the CCP-weighted transition depends on the iterating CCPs. When $L \cdot |\mathcal{X}|$ is large, GFD-EM's per-iteration cost is asymptotically negligible.

3.10 Non-Stationary Lifecycle and Short Panels

The setup above accommodates stationary mixtures of the kind studied by Arcidiacono and Miller (2011). For non-stationary finite-horizon models—lifecycle labour supply, education choice, retirement—the GFD identity of Theorem 7 applies elementwise per type, yielding a corollary that we record because it underwrites the consistency result below.

Corollary 9 (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 $\rho^\star_\ell$ for every $\ell \in \{1, \ldots, L\}$ 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 $\mathbf{\Phi}^{(\ell)}(x_t)$ such that

$$v_t^{(\ell)}(x_t, d_t) - v_t^{(\ell)}(x_t, d_t') = u_t^{(\ell)}(d_t, x_t; \theta^{(\ell)}) - u_t^{(\ell)}(d_t', x_t; \theta^{(\ell)}) + \sum_{\tau=1}^{\rho_\ell^\star} \beta^\tau \sum_{(k_\tau, j_\tau)} \mathbf{\Phi}^{(\ell)}(x_t; d_t, d_t')[(k_\tau), (j_\tau)] \big[u_{t+\tau}^{(\ell)}(k_\tau, j_\tau; \theta^{(\ell)}) + \psi_{t+\tau}^{(\ell)}(k_\tau, j_\tau)\big], \tag{39}$$

where $\psi_t^{(\ell)}(k, j) = -\log p_t^{(\ell)}(k \mid j)$ is the type-conditional Hotz–Miller correction. The mixture log-likelihood inherits this representation and is computable from $(\widehat{\mathbf{\Phi}}^{(\ell)}, \widehat{p}_\tau^{(\ell)})$ without invoking any type-conditional Bellman fixed point.

The corollary follows by applying Theorem 7 to each type-conditional submodel; the mixture aggregation does not interact with the GFD identity since each type's flow weights and CCP corrections are computed within the type. Proofs are in Appendix E.8.

Relation to the AM 1998 representative-CCP estimator. The GFD identity of Corollary 9 is the finite-mixture extension of Altug and Miller (1998)'s Proposition 1. The two estimators share an identical *mechanism*—finite dependence—and differ in two implementation choices. The mechanism is the terminal-distribution matching that Theorem 3 characterizes: at horizon $\rho^\star$ the future state distributions induced by the compared actions coincide, so the continuation value $V_0(\cdot)$ in the GFD expansion cancels in the difference $V_1(\cdot) - V_0(\cdot)$ (AM 1998 invoke this cancellation immediately below their (5.5); it is the substantive content of finite dependence specialised to a ρ-lag shift register). The two implementation choices are:

(a) *Flow-weight construction.* AM 1998 use the Type S trivial-flush construction tailored to a ρ-lag shift register state. Corollary 9 uses the canonical Moore–Penrose flow input of §3.4, which is available whenever pointwise FD holds at any horizon $\rho^\star < \infty$ (Theorem 3), including all four type taxa of §4.1 and the non-classical residual class.

(b) *Unobserved heterogeneity.* AM 1998 handle the individual-specific effect $\nu_n \eta_n$ via the projection method of MaCurdy (1981)—a deterministic (nonparametric) function of observables x_n. GFD-EM replaces the projection with a finite-mixture EM step. When the relevant dimensions of heterogeneity are not captured by observables (latent "types" orthogonal to demographics), the MaCurdy projection is misspecified; the mixture step restores consistency.

In short, AM 1998 is the GFD specialisation to Type S finite dependence with MaCurdy projection of heterogeneity; GFD-EM is the GFD identity for any FD-feasible transition structure with a finite-mixture treatment of latent heterogeneity. Finite dependence is the common mechanism.

Theorem 10 (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 9. Assume:*

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

*(ii) (*Type 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 $t \in [t_{start}, t_{end} - \rho^\star_\ell]$.*

*(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 GFD-EM estimator $(\widehat{\boldsymbol{\pi}}, \widehat{\boldsymbol{\theta}})$ that maximises the GFD-representation pseudo-log-likelihood through E-step posterior averaging and M-step optimization converges in probability to $(\boldsymbol{\pi}_0, \boldsymbol{\theta}_0)$ as $N \to \infty$ and is asymptotically normal with sandwich variance characterized by the GFD-based score and Hessian.

Proof sketch. The GFD representation (39) expresses the type-conditional choice probabilities as continuous functions of $(\theta^{(\ell)}, \mathbf{\Phi}^{(\ell)}, p^{(\ell)})$ where $p^{(\ell)}$ enters only through the boundary periods $t + \tau$, $\tau \le \rho^\star_\ell$. Under (iii) the empirical anchor $\widehat{p}_t^{(\ell)}$ is consistent, so the GFD-pseudo likelihood converges uniformly to its population analogue. Under (ii) the population pseudo-log-likelihood has a unique maximiser at $(\boldsymbol{\pi}_0, \boldsymbol{\theta}_0)$ modulo label switching; condition (i) ensures the GFD representation is well-defined. Standard EM convergence theory (Wu, 1983; McLachlan and Krishnan, 2008, § 3) delivers convergence of EM iterates to a stationary point of the observed-data likelihood (not necessarily the global maximiser); consistency is of the arg max estimator, which under (ii) is the unique population maximiser modulo label-switching by the standard M-estimator argument. Asymptotic normality follows by a Taylor expansion of the score around $(\boldsymbol{\pi}_0, \boldsymbol{\theta}_0)$ with the GFD-induced influence function. The full argument appears in Appendix E.9.

Relation to the literature. Theorem 10 extends the EM consistency results of Arcidiacono and Miller (2011) for stationary finite-mixture DDC to the non-stationary lifecycle case, and extends the GFD two-step of Altug and Miller (1998) from the MaCurdy-projection treatment of individual effects (MaCurdy, 1981) to a finite-mixture treatment of latent types. The identification condition (ii) is the analogue of the nonparametric mixture identification of Kasahara and Shimotsu (2009) when the type-conditional CCPs are distinct on a positive-mass set. For HMM-style heterogeneity with type evolution (Hu and Shum, 2012), the same machinery applies with type-state-conditional GFD weights.

A finite-sample Monte Carlo illustration appears in §4.2.3: with $T = 10$, $K = 2$ unobserved types orthogonal to observables, the AM 1998 GFD two-step without a mixture step

delivers a single-θ MLE that is biased by $|-1.96|$ in b against type 1 and $|+0.89|$ in g against type 2, while GFD-EM with the identical GFD identity plus a mixture EM step recovers both type-conditional parameters and the type prior to within $|0.05|$ across all components (data `data/gfd_em_am_benchmark_mc.json`).

Positioning summary. Table 1 summarises GFD-EM's coverage relative to the closest alternatives in the dynamic discrete choice literature. The four columns correspond to substantive features that, in combination, distinguish the dynamic discrete choice problems we target: multi-agent strategic interaction, latent unobserved heterogeneity treated via mixture (rather than via projection onto observables), non-stationary panels (finite-horizon lifecycle or short observation window inside a longer model horizon), and counterfactual analysis that does not require re-solving the Markov-perfect equilibrium for every candidate parameter value. NPL and Hotz–Miller both support latent-UH treatment when wrapped in EM (the Hotz–Miller+EM combination is canonical in Arcidiacono and Miller, 2011). To our knowledge GFD-EM is the only estimator that covers the first three columns jointly; the multi-agent counterfactual column carries an asterisk for GFD because exact-counterfactual solutions still require iterating the GFD anchor while approximate (frozen-anchor) counterfactuals are obtained in closed form.

4 Applications

This section verifies Theorem 3 on nine canonical DDC models (§4.1), benchmarks the GFD estimator's finite-sample behavior against the Hotz–Miller (HM) CCP estimator, nested pseudo-likelihood (NPL), the mathematical program with equilibrium constraints (MPEC), and the nested fixed-point algorithm (NFXP) (§4.2), and demonstrates Theorem 7's counterfactual GFD identity in single-agent and multi-agent settings (§4.3).

4.1 Existence Test on Canonical DDC Models

We organize the verification on nine canonical DDC models 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 $\rho^\star$ in closed form; we verify each below directly from the transition structure (the Kronecker case via Lemma 11). 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 $\text{Range}(\mathbf{A})$ at every triple $(x_t, d_t, 0)$, so $\rho^\star = 1$ via the classical AM non-negative solution.

Type S (shift register, p-lag). The state factorizes into a p-component lag register of past

Method	Multi-agent games	Latent UH (EM)	Non-stationary panel	Bellman-free counterfactual
NFXP–EM	✓[†]	✓	✓[‡]	×
NPL (Aguirregabiria and Mira, 2007)	✓	✓[♭]	×	partial
Hotz–Miller (Hotz and Miller, 1993)	×	✓[♭]	partial	✓
AM 1998 (Altug and Miller, 1998)	×	✓[◇]	partial	×
AM 2011 EM (Arcidiacono and Miller, 2011)	×	✓	×	×
AM 2019 (Arcidiacono and Miller, 2019)	✓	×[§]	✓	partial
KSSR 2021 (Kalouptsidi et al., 2021)	✓	×	✓	✓
GFD–EM (this paper)	✓	✓	✓	✓[*]

Table 1: Method-coverage matrix for dynamic discrete choice estimators. ✓ = method handles the feature; × = does not (or has known limitations). [†]NFXP–EM solves the Markov-perfect equilibrium at every candidate parameter value; multi-agent + non-stationary applications are tractable only at small state spaces, and identification of structural parameters becomes sensitive to the terminal value-function specification when the observation window is strictly inside the model horizon (Table 6). [‡]NFXP–EM accommodates non-stationarity via backward induction from a chosen terminal V, but in short-panel applications the choice of that terminal V is not identified from data and introduces systematic bias. [◇]AM 1998 treat the time-invariant individual-specific effect via the projection method of MaCurdy (1981), which assumes the heterogeneity dimension is a deterministic function of observables; this is misspecified for latent type heterogeneity orthogonal to the observables (Table 5). [♭]NPL and Hotz–Miller admit straightforward mixture-EM extensions; we mark them ✓ because the per-iteration cost of the embedded NPL or Hotz–Miller step remains compatible with EM, and AM 2011 EM (Arcidiacono and Miller, 2011) is the canonical published Hotz–Miller+EM. [§]AM 2019 accommodate heterogeneity in observed transitions but not in latent type membership. [*]GFD counterfactuals with a frozen data-anchored CCP yield Bellman-free closed-form updates; exact counterfactuals under large parameter shifts iterate the GFD anchor instead of full Bellman fixed-point solving. "Partial" for the other methods indicates per-parameter equilibrium iteration but not full backward induction.

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 $\rho^\star = p$ 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 $\rho^\star_{\text{joint}} = \rho^\star_{\text{endo}}$ regardless of the persistence of $\mathbf{F}^{\text{exo}}$ (Lemma 11 in Appendix B). The structured GFD solver of Appendix B runs the flow QP on the M-dimensional endogenous projection only, yielding Z^ρ-scaling wall-clock savings at large Z (Table 15 in the appendix).

Type D (triple connector / strategic ripple). All-endogenous single-agent transitions or shared-action games have no closed-form $\rho^\star$ 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 LSMR, and reads off $\rho^\star$ as the smallest passing horizon.

Remark 4 (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 are different: a per-player projected $\rho = 1$ certificate is an artifact of integrating out opponents' equilibrium CCPs, since the strategic ripple through opponents' best responses materialises at $\rho = 2$. We therefore use $\rho^\star = 2$ throughout for multi-player Type-D models.*

Results. Table 2 reports the predicted and LSQR-verified $\rho^\star$ together with the worst-case LSQR residual at the certifying horizon for nine canonical DDC models drawn from the applied literature, implemented as `code/example_*.py` in the replication package.

Model	Source	Type	S	D	$\rho^\star$ pred.	$\rho^\star$ verif.	LSQR resid.
Engine Replacement	Rust (1987)	R	10	2	1	1	8.7×10^{-14}
Job Search	—	R	10	2	1	1	1.3×10^{-16}
Female Labor (p=1)	Altug and Miller (1998)	S	2	2	1	1	3.1×10^{-16}
Investment	—	K	20	3	1	1	7.9×10^{-14}
Inventory Control	—	K	20	3	1	1	7.9×10^{-14}
Entry/Exit (Endo)	—	D	16	2	2	2	4.8×10^{-13}
Altug–Miller (p=3)	Altug and Miller (1998)	S	8	2	3	3	1.9×10^{-13}
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 2: Verification of the nine canonical models. Type codes: R = renewal; S = shift register; K = Kronecker with action-invariant exogenous factor; D = triple connector / strategic ripple. Predicted $\rho^\star$ comes from the type-specific closed forms of §4.1 (R, S, K), or direct verification at $\rho = 1$ followed by $\rho = 2$ for Type D; verified $\rho^\star$ 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, computed by sparse LSQR. ‡ The entry/exit game and the dynamic oligopoly are multi-player Markov-perfect models in which the equilibrium-induced strategic ripple acts at horizon two through opponents' best-response feedback; the operative horizon is therefore $\rho^\star = 2$. A per-player projected feasibility check can produce an apparent $\rho = 1$ certificate; we treat this as an artifact of the projection rather than a property of the equilibrium-induced game, and adopt $\rho = 2$ throughout (Remark 4). All ten verifications accompany the replication package.

Findings. Predicted and verified $\rho^\star$ agree across all nine models, with LSQR residuals at 10^{-13}–10^{-17} (machine precision). The Type-D multi-player games operate at $\rho^\star = 2$ because the equilibrium-induced strategic ripple acts at horizon two; a per-player projected feasibility test can spuriously certify $\rho = 1$ on the rank-deficient projection, which we treat as an artifact and use $\rho = 2$ throughout the Monte Carlo of §4.2.2 (Remark 4).

4.2 Estimation Monte Carlo

We verify the GFD estimator's finite-sample behavior against standard benchmarks on two DGPs: the single-agent Type-K investment model ($|\mathcal{X}| = 20$, $\rho^\star = 1$; §4.2.1) and the three-player Type-D stationary entry/exit game ($|\mathcal{X}_i| = 40$, $\rho^\star = 2$; §4.2.2). 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.

4.2.1 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), $|\mathcal{X}| = 20$. Three actions shift capital deterministically by $-1, 0, +1$ while productivity follows an action-invariant AR(1), giving Type-K structure with $\rho^\star = 1$ via Lemma 11. 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 $\theta_0 = (2.5, 0.3, 0.1)$, calibrated for steady-state action shares $(47\%, 32\%, 21\%)$; $\beta = 0.95$. 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 (empirical $\hat{p}$)	−0.003	0.044	0.001	0.023	−0.001	0.015	0.018
GFD (logit $\hat{p}$)	0.061	0.076	−0.000	0.023	0.001	0.015	0.057
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 3: Investment model Monte Carlo: 500 replications, $N = 2{,}000$, $T = 20$, Apple M3 Ultra. GFD (empirical $\hat{p}$) uses Laplace-smoothed cell-frequency CCPs ($\lambda = 0.1$); GFD (logit $\hat{p}$) replaces the first stage with a parametric logit smoother on basis $(1, k, o, k{\cdot}o)$.

GFD matches CCP-2step, MPEC, and NFXP to $\sim$0.001 in θ-units on θ^{rev} and θ^{cost} with a $\sim$25% residual efficiency loss on θ^{adj} (RMSE 0.015 vs 0.012) that the variance-optimal input of §3.7 would close. Wall-clock: GFD is 3.4× faster than CCP-2step, 50× faster than MPEC, and 290× faster than NFXP. Sweeping 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 rises from $\sim$17× to $\sim$136× (Table 16 in Appendix B).

4.2.2 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) given a market-size index $z_t \in \{2, 4, 6, 8, 10\}$ (5-state persistent exogenous Markov chain) and the incumbency profile $(y_{1,t}, y_{2,t}, y_{3,t}) \in \{0, 1\}^3$. Per-player state $x_{i,t} = (z_t, y_{1,t}, y_{2,t}, y_{3,t})$ has $|\mathcal{X}_i| = 40$ values. Per-period payoffs are linear in $\theta_i = (\theta_i^{\text{rev}}, \theta_i^{\text{comp}}, \theta_i^{\text{fc}}, \theta_i^{\text{entry}})$:

$$u_i(x_{i,t}, d_{i,t} = 1; \theta_i) \;=\; \theta_i^{\text{rev}} \log(z_t) \;-\; \theta_i^{\text{comp}}\, \mathbb{E}\big[\log(1 + N_{-i,t}^{\text{enter}})\big] \;-\; \theta_i^{\text{fc}} \;-\; \theta_i^{\text{entry}}\,(1 - y_{i,t}), \quad (40)$$

with $u_i(\cdot, 0; \theta_i) = 0$. 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 $\rho^\star = 2$ (Remark 4). 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 4: 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 ($|\text{bias}| \leq 0.017$); GFD's RMSE sits within $\sim$ 7% of HM/NPL and within $\sim 8\%$ of NFXP, confirming 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) 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 computational payoff materialises in the counterfactual setting of §4.3.

4.2.3 Lifecycle with latent unobserved heterogeneity: GFD identity alone vs. GFD-EM mixture

This experiment illustrates Theorem 10: when unobserved types are genuinely latent (orthogonal to observables), the AM 1998 GFD (Type-S specialisation) two-step estimator without a mixture step is misspecified and the single-θ MLE settles at neither type's truth nor the population-weighted mean. GFD-EM uses the same GFD identity but augments it with an EM step over latent types, recovering both type-conditional parameters and the type prior consistently.

Setup. Two unobserved types $(\boldsymbol{\theta}_0^{(1)}, \boldsymbol{\theta}_0^{(2)}) = ((1.0, 0.3), (-0.5, -0.2))$ with prior $\boldsymbol{\pi}_0 = (0.6, 0.4)$ choose a binary action $d_t \in \{0, 1\}$ over $T = 10$ lifecycle periods with the type-conditional flow payoff $u^{(\ell)}(1, x_t) = b^{(\ell)} + g^{(\ell)}(l_{t-1} + l_{t-2})$ and $u^{(\ell)}(0, x_t) = 0$. The 2-lag state $x_t = (l_{t-1}, l_{t-2})$ gives a Type S shift register at $\rho = 2$, $\beta = 0.95$. Type membership L_n is i.i.d. across n and *orthogonal to all observable covariates*, so no projection of the form $\nu_n \eta_n = f(x_n)$ can recover type. We generate $N = 2{,}000$ agents and run 10 replications.

Estimators. The AM 1998 GFD two-step (*AM-GFD-homogeneous*) uses the GFD identity of Corollary 9 with the empirical pooled CCP anchor $\widehat{p}_t(d \mid x)$ and assumes a single homogeneous θ across all individuals—the value to which the MaCurdy projection reduces when no observable proxies type. The GFD-EM estimator uses the identical GFD identity at $\rho = 2$ but with the posterior-weighted type-conditional anchor $\widehat{p}_t^{(\ell)}(d \mid x)$ and a mixture EM step over $K = 2$ latent types.

Estimator	b bias	b RMSE	g bias	g RMSE
AM-GFD-homogeneous (single θ, no UH):				
vs. true $\theta_1 = (1.0, 0.3)$	-1.96	1.96	$+0.39$	0.39
vs. true $\theta_2 = (-0.5, -0.2)$	-0.46	0.46	$+0.89$	0.89
vs. population mean $(0.4, 0.1)$	-1.36	1.36	$+0.59$	0.59
GFD-EM ($K = 2$ mixture, posterior-weighted anchor):				
vs. true $\theta_1 = (1.0, 0.3)$	-0.003	0.133	$+0.003$	0.046
vs. true $\theta_2 = (-0.5, -0.2)$	$+0.034$	0.067	-0.016	0.038
vs. true $\boldsymbol{\pi} = (0.6, 0.4)$	-0.004	0.010	$+0.004$	0.010

Table 5: Lifecycle with latent unobserved heterogeneity: 10 replications, $T = 10$, $K = 2$ types, $\rho = 2$, $N = 2{,}000$, $\beta = 0.95$. Both estimators use the identical GFD identity of Corollary 9; only the UH handling differs. AM-GFD-homogeneous (no UH, mimicking the MaCurdy projection when observables carry no type information) delivers a single-θ MLE that is far from either type's truth *and* from the population-weighted mean. GFD-EM correctly recovers both type-conditional parameters and the type prior with $|\text{bias}| < 0.05$ across all components. Data: `data/gfd_em_am_benchmark_mc.json`; script: `code/lifecycle_em_am_benchmark.py`.

Findings. The AM-GFD-homogeneous estimator's biases on b reach $|-1.96|$ (against θ_1) and $|-0.46|$ (against θ_2); on g, $|+0.89|$ against θ_2. These are not sampling-noise artifacts—RMSE roughly equals |bias|, so the single-θ MLE has converged in finite-sample to a misspecified target that no plausible N will close. The pooled MLE does not even land at the population mean (bias -1.36 on b, $+0.59$ on g against the $\boldsymbol{\pi}_0$-weighted average), reflecting that the GFD likelihood is non-linear in θ and the misspecified maximiser is determined by the type-conditional likelihood curvature rather than by an arithmetic average of θ_1, θ_2. GFD-EM uses the identical GFD identity but augments it with the mixture EM step, recovering each type-conditional parameter to within $|0.05|$ and the type prior to within $|0.01|$. The demonstration is the estimation-side analog of the counterfactual-side future-ambiguity result in §4.3.2: *finite dependence already cancels the terminal continuation (Theorem 3) and so the GFD identity is robust to boundary objects, but neither AM 1998's GFD specialisation nor canonical GFD is robust to within-sample misspecification of the unobserved heterogeneity structure. The mixture EM step is what closes that gap.*

4.2.4 Multi-agent short panel: terminal-V sensitivity in a 2-player non-stationary game with latent UH

This experiment generalizes the short-panel terminal-V ambiguity of §4.2.3 from a single-agent lifecycle setting to a non-stationary 2-player dynamic game with $K = 2$ latent market types. It is the central demonstration that, in the (multi-agent × latent UH × non-stationary) cell of Table 1, NFXP-EM-style estimation depends systematically on an unidentified terminal value-function specification, whereas GFD-EM does not.

Setup. Two unobserved market types $(\boldsymbol{\theta}_0^{(1)}, \boldsymbol{\theta}_0^{(2)}) = ((1.5,\, 1.5,\, 0.3),\, (0.5,\, 1.5,\, 0.3))$ – distinguished by per-period revenue – with prior $\boldsymbol{\pi}_0 = (0.6, 0.4)$. Each market hosts a 2-player entry game with binary $d_{i,t} \in \{0, 1\}$ and the 2-lag incumbency state $x_t = (l_{1,t-1}, l_{2,t-1})$. Per-period payoff is $u_i(1, x_t; \boldsymbol{\theta}^{(\ell)}) = \theta^{(\ell)}_{\text{rev}} - \theta^{(\ell)}_{\text{comp}}\, l_{-i,t-1} - \theta^{(\ell)}_{\text{entry}}\,(1 - l_{i,t-1})$ with $\beta = 0.95$. True horizon $T_{\text{true}} = 30$ with terminal $V = 0$; observed window $T_{\text{obs}} = 15$. The econometrician sees only the first 15 periods. $N = 400$ markets, 3 replications.

Estimators. NFXP-EM requires a boundary value $V_{T_{\text{obs}}}(x)$; we sweep five plausible specifications $V_{T_{\text{obs}}}(x) = c\,(l_1 + l_2)$ with $c \in \{-3, -1, 0, 1, 3\}$. *Proper* GFD-EM applies the $\rho = 2$ GFD identity of Corollary 9 algebraically: per-period value differences $v_i^{(\ell)}(d = 1, x_t, t) - v_i^{(\ell)}(d = 0, x_t, t)$ are computed as a closed-form linear combination of $u_i(\cdot; \boldsymbol{\theta}^{(\ell)})$ and posterior-weighted empirical $\psi_i^{(\ell)}$ at periods $t+1, t+2$, weighted by flow matrices $\boldsymbol{\Phi}_i^{(\ell)}(x_t, t)$ that solve the FD linear system (Theorem 3) at the time-varying perceived transitions $F_i^{(\ell)}(t)$, $F_i^{(\ell)}(t+1)$. No backward induction is performed; no boundary V enters the algorithm at any step.

	Type 1 bias: $(b^{\text{rev}}, b^{\text{comp}}, b^{\text{entry}})$			Type 2 bias: $(b^{\text{rev}}, b^{\text{comp}}, b^{\text{entry}})$		
Estimator	b_1^{rev}	b_1^{comp}	b_1^{entry}	b_2^{rev}	b_2^{comp}	b_2^{entry}
NFXP-EM, $V_T = -3\cdot$ lags	+0.306	−0.218	−0.486	+0.144	−0.052	−0.191
NFXP-EM, $V_T = -1\cdot$ lags	+0.161	−0.139	−0.250	+0.087	−0.062	−0.105
NFXP-EM, $V_T = 0$ (naive)	+0.087	−0.082	−0.097	+0.057	−0.061	−0.037
NFXP-EM, $V_T = +1\cdot$ lags	+0.030	−0.028	+0.041	+0.033	−0.057	+0.067
NFXP-EM, $V_T = +3\cdot$ lags	+0.003	+0.139	+0.185	−0.000	+0.047	+0.227
GFD-EM (no V)	+0.107	+0.041	−0.005	+0.029	−0.052	+0.004

Table 6: 2-player non-stationary game with $K = 2$ latent market types: 3 replications, $T_{\text{true}} = 30$, $T_{\text{obs}} = 15$, $N = 400$ markets, $\beta = 0.95$. True parameters: $\boldsymbol{\theta}_0^{(1)} = (1.5, 1.5, 0.3)$, $\boldsymbol{\theta}_0^{(2)} = (0.5, 1.5, 0.3)$, $\boldsymbol{\pi}_0 = (0.6, 0.4)$. NFXP-EM b_1^{rev} biases range over 0.30 θ-units across the V-sweep; b_1^{entry} biases over 0.67. GFD-EM delivers a single estimate (the algorithm has no $V_{T_{\text{obs}}}$ parameter) with $|b_1^{\text{entry}}| = 0.005$ and $|b_2^{\text{entry}}| = 0.004$ – smaller than any NFXP-EM V choice on the entry-cost dimension. Data: `data/gfd_em_game_vsensitivity_proper_mc.json`; script: `code/dynamic_game_uh_bridge_proper.py`.

Findings. NFXP-EM biases on the type-1 entry-cost parameter θ_1^{entry} span $|-0.486|$ to $|+0.185|$ – 0.67 θ-units, twice the true entry-cost value 0.3 – as $V_{T_{\text{obs}}}$ varies over plausible specifications.

The naive $V_{T_{\text{obs}}} = 0$ choice retains $|\,\text{bias}\,| \leq 0.097$, but no data evidence guides the analyst to that choice rather than another among the alternatives. GFD-EM has no $V_{T_{\text{obs}}}$ parameter – the algorithm only consumes $\boldsymbol{\theta}$, the empirical CCP anchors $\widehat{p}^{(\ell)}(t+1), \widehat{p}^{(\ell)}(t+2)$, and the time-varying flow weights $\boldsymbol{\Phi}_i^{(\ell)}(x_t, t)$ solved from the FD linear system – so it produces a single estimate that is structurally insensitive to any V assumption. Numerically, GFD-EM achieves $|b_1^{\text{entry}}| = 0.005$ and $|b_2^{\text{entry}}| = 0.004$, smaller than the bias under any of the five NFXP-EM V specifications including the naive $V_{T_{\text{obs}}} = 0$. This is the multi-agent extension of Theorem 10's identification claim and the central message of Table 1: in the (multi-agent × latent UH × non-stationary) cell, only GFD-EM identifies the structural parameters without committing to an unidentified boundary value function.

4.3 Counterfactual Monte Carlo

This subsection turns to counterfactual computation—*can we compute counterfactuals without re-solving the equilibrium?*—by demonstrating Theorem 7 and Proposition 8 in single-agent and multi-agent settings. The GFD identity holds at high numerical precision across all four DGPs reported below; the wall-clock and epistemic implications differ sharply between settings, and we report each honestly.

4.3.1 Single-agent counterfactuals: wall-clock win on Type-K and Type-D

In the single-agent case the structural transition f is θ-independent, so the flow input $\boldsymbol{\Phi}$ built from baseline f remains strictly valid under any payoff-only counterfactual $\theta \to \theta_1$. The GFD anchor iteration $\mathbf{p}^{(m+1)} = \Lambda(\widetilde{v}^{\rho\text{-FD}}(\cdot;\, \theta_1, \boldsymbol{\Phi}, \mathbf{p}^{(m)}))$ with $\boldsymbol{\Phi}$ *frozen* at the baseline canonical Moore–Penrose flow input therefore delivers the new equilibrium CCPs without re-solving the Bellman fixed point. We demonstrate this on two DGPs spanning the FD taxonomy.

We demonstrate the GFD counterfactual identity on two single-agent DGPs: the Type-K investment model of §4.2.1 ($\rho^\star = 1$) sweeping $K = 8$ payoff perturbations across state sizes $|\mathcal{X}| \in \{20, 50, 100, 200\}$, and the Type-D entry/exit DGP ($|\mathcal{X}| = 16$, $\rho^\star = 2$) sweeping $K = 11$ perturbations of (revenue, fixed cost, entry cost). Table 7 reports the sweep wall-clock and CCP agreement against NFXP value-function iteration.

DGP	$\lvert\mathcal{X}\rvert$	NFXP sweep (ms)	GFD sweep (ms)	Speedup	CCP err
Type-K Investment ($\rho^\star = 1$, $K = 8$ perturbations):					
Investment	20	109.4	11.0	9.9×	7×10^{-7}
Investment	50	125.6	16.5	7.6×	7×10^{-7}
Investment	100	147.3	35.1	4.2×	7×10^{-7}
Investment	200	254.3	111.4	2.3×	7×10^{-7}
Type-D Entry/Exit ($\rho^\star = 2$, $K = 11$ perturbations):					
Entry/Exit	16	136.4	40.5	3.4×	2.3×10^{-4}

Table 7: Single-agent counterfactual GFD identity: sweep wall-clock and CCP sup-norm error vs NFXP. GFD sweep includes the one-time $\mathbf{\Phi}$-build cost amortized over K scenarios. 10 replications per cell on Apple M3 Ultra. Scripts: `counterfactual_investment.py`, `counterfactual_entry_exit.py`.

Findings. On the Type-K investment DGP, the GFD identity delivers a 9.9× sweep speedup at $|\mathcal{X}| = 20$ that attenuates to 2.3× at $|\mathcal{X}| = 200$ as the per-iteration cost of the frozen-$\mathbf{\Phi}$ contraction grows quadratically in $|\mathcal{X}|$; CCP agreement against NFXP is at sup-norm 7×10^{-7} uniformly, confirming the population claim of Theorem 7. On the Type-D entry/exit DGP at $|\mathcal{X}| = 16$, the GFD per-scenario cost is 0.13ms vs. NFXP 12.4ms (a 96× per-scenario factor), amortizing to 3.4× sweep speedup once the one-time $\mathbf{\Phi}$ build (39ms) is included.

4.3.2 Multi-agent counterfactuals: stationary corner case and the epistemic frontier

Strategic ripple and the per-player perceived transition. In a stationary MPE of the 3-player game of §4.2.2, each player i takes the others' equilibrium CCPs $\mathbf{p}_{-i}^{\mathrm{MPE}}$ as given and faces the perceived transition

$$F_i^{\mathbf{p}_{-i}^{\mathrm{MPE}}}(x_i' \mid x_i, d_i) \;=\; \sum_{d_{-i}} F_S(x_i' \mid x_i, d_i, d_{-i}) \prod_{j \neq i} p_j^{\mathrm{MPE}}(d_j \mid x_j),$$

where F_S is the joint exogenous transition (here driven by the market-size process). Per-player FD is verified directly on this projected kernel via Theorem 3, and holds at $\rho^\star = 2$. The intuition is the *strategic ripple*: player i's action d_i at t shifts the joint state at $t+1$, which feeds into opponents' best-response CCPs $\mathbf{p}_{-i}$ at $t+1$, which in turn propagate back into i's perceived transition at $t+2$. Two periods are needed to flush this best-response feedback; verification on the entry/exit game and the dynamic oligopoly appears in Table 2 (the multi-player Type-D rows).

GFD anchor iteration in the game. Once $\rho^\star = 2$ is certified, we precompute the canonical Moore–Penrose flow input $\mathbf{\Phi}_i$ per player by solving the flow QP of Section 2.4 on $F_i^{\mathbf{p}_{-i}^{\mathrm{MPE,\ baseline}}}$.

For a counterfactual structural parameter θ_{cf}, the GFD anchor iteration in CCP space alternates the following per-player update until convergence:

1. *Time-t payoff at the iterate.* Compute $u_i(d, x_i; \theta_{\text{cf}})$ at the current iterate $\mathbf{p}_{-i}^{(m-1)}$ (the competition expectation in eq. (40) folds in opponents' current CCPs).

2. *Hotz–Miller correction.* Compute $\psi_d(x_i') = \gamma_E - \log p_i^{(m-1)}(d \mid x_i')$ at every future state x_i'.

3. *GFD continuation correction.* Apply the cached $\mathbf{\Phi}_i$ to the discounted future $u + \psi$:

$$c_i(x_i) \;=\; \sum_{\tau=1}^{2} \beta^{\tau} \sum_{(k_\tau, j_\tau)} \mathbf{\Phi}_i(x_i)\big[(k_\tau)_{\tau=1}^{2},\, (j_\tau)_{\tau=1}^{2}\big] \left(u_i(k_\tau, j_\tau;\, \theta_{\text{cf}}, \mathbf{p}_{-i}^{(m-1)}) + \psi_{k_\tau}(j_\tau)\right).$$

4. *Value-difference and logit update.* The per-player value-difference under θ_{cf} is $v_i^1(x_i) - v_i^0(x_i) = u_i^1 - u_i^0 + c_i(x_i)$, and the CCP update is $p_i^{(m)}(1 \mid x_i) = \sigma(v_i^1 - v_i^0)$.

The iteration warm-starts at the baseline MPE CCPs and converges in 9–12 outer iterations to tolerance 10^{-6}.

Frozen-$\mathbf{\Phi}$ vs strict $\mathbf{\Phi}$-resolve. Strictly, $\mathbf{\Phi}_i$ depends on the perceived transition $F_i^{\mathbf{p}_{-i}}$, which itself depends on the iterating CCPs $\mathbf{p}_{-i}$. Two implementations are reported: *frozen* $\mathbf{\Phi}$, where $\mathbf{\Phi}_i$ is held at the baseline MPE value throughout the iteration—fast, warm-started, but not literally consistent with the counterfactual perceived transition; and *strict* $\mathbf{\Phi}$*-resolve*, where the per-player flow QP is re-solved at the new perceived transition before the iteration starts.

Stationary 3-player MPE sweep. We re-scale the persistence of the baseline market-size transition matrix at four levels (rescaling factor $\alpha \in \{0.7, 0.85, 1.15, 1.30\}$, where $\alpha < 1$ makes the chain more persistent and $\alpha > 1$ less persistent) and ask the strict-$\mathbf{\Phi}$-resolve GFD to reproduce each new equilibrium MPE on $|\mathcal{X}_i| \in \{24, 40\}$. Table 8 reports the sweep wall-clock cost and CCP agreement.

$\lvert\mathcal{X}_i\rvert$	NFXP-warm (ms)	GFD identity (ms)	Speedup	CCP err
24	47	458	0.10×	1.4×10^{-4}
40	52	2072	0.02×	1.3×10^{-4}

Table 8: Stationary 3-player MPE GFD sweep across $K = 4$ re-scaled market-size transition matrices. GFD identity is theoretically valid but loses on wall-clock to warm-started NFXP because the per-scenario flow-QP rebuild ($\sim$ 117ms at $|\mathcal{X}_i| = 24$) exceeds the warm-started NFXP cost. Script: `counterfactual_game_stationary.py`.

The second setting is a 2-player finite-horizon non-stationary entry/exit game over $T = 30$ periods, where the agent at $t = 1$ knows the near future ($t = 1, 2, 3$) but the far future ($t =$

$4, \ldots, T$) is uncertain between two candidate scenarios: *expansion* (revenue grows 10%/period, market drifts upward) and *decline* (revenue decays 10%/period, market drifts downward). To compute the optimal period-1 CCP, NFXP must *commit to one scenario* and run full backward induction over all $T = 30$ periods; the GFD identity instead asks only for the local model $\{f_s, u_s\}_{s=1}^2$ plus the empirical CCPs at $t = 3$ (which can be read directly from cross-sectional data under either scenario). Table 9 reports the resulting period-1 CCPs for Player 0 in the entrant role at three market-size levels.

	$P_1(\text{enter} \mid z, y_0 = 0)$		
	$z = 2$ (low)	$z = 4$ (mid)	$z = 6$ (high)
Empirical anchor (observed CCP at $t = 3$ used by the GFD identity):			
$\widehat{P}_3$ under scenario A	0.6240	0.8682	0.9371
$\widehat{P}_3$ under scenario B	0.5816	0.8484	0.9283
NFXP (requires full trajectory $\{f_s, u_s\}_{s=1}^{30}$ specified per scenario):			
NFXP-A	0.6065	0.8616	0.9345
NFXP-B	0.6059	0.8614	0.9344
GFD identity (requires only $\{f_s, u_s\}_{s=1}^2$ + observed $\widehat{P}_3$):			
GFD with $\widehat{P}_3^A$	0.6065	0.8616	0.9345
GFD with $\widehat{P}_3^B$	0.6059	0.8614	0.9344

Table 9: Non-stationary 2-player game under future ambiguity: period-1 CCP $P_1(\text{enter})$ for Player 0 at three market-size levels. The GFD identity reproduces each scenario's NFXP solution exactly (sup-norm $\sim 4 \times 10^{-8}$) while requiring only the local 3-period model plus one empirical CCP anchor at $t = 3$; NFXP requires the full 30-period trajectory under the chosen scenario. Script: `counterfactual_game_future_ambiguity.py`.

Findings. In the stationary MPE (Table 8), the GFD identity is valid (CCP sup-norm error $\sim 10^{-4}$) but does *not* beat warm-started NFXP on wall-clock: the per-scenario flow-QP rebuild dominates the cheap per-scenario MPE re-solve, and the gap widens with state space ($0.10\times \to 0.02\times$). This is the honest negative finding—the GFD identity's value for stationary multi-agent games is theoretical (Bellman-free identity), not wall-clock. In the non-stationary game under future ambiguity (Table 9), the GFD identity reproduces NFXP exactly under each candidate far-future scenario to four decimal places using only $\{f_s, u_s\}_{s=1}^2$ plus the empirical CCP anchor at $t = 3$; NFXP requires the full 30-period trajectory under the chosen scenario. If the analyst observes $\widehat{P}_3$ from data but does *not* know which far-future scenario will realise, the GFD identity produces a computable answer from the empirical anchor, whereas NFXP cannot be set up without first committing to a specific far-future trajectory. This is the model class where Theorem 7 and Proposition 8 sidestep a commitment that NFXP cannot avoid.

5 Empirical Application: Non-stationary US Airline Entry Game

Do airline entry costs vary over time, and does it matter if we assume otherwise? We answer both questions using a dynamic discrete game of the four major U.S. carriers over 2010–2023 — a panel that spans two structural breaks: the 2008–2013 network-carrier merger wave and the 2020 COVID-19 demand shock.

The first question has a clear answer. A Wald test of entry-cost stationarity rejects the null hypothesis for all four carriers at the 0.1% level (Table 10). Southwest's estimated entry cost fell by 6.4 logit units between its 2019 peak and 2020; Delta's path shows a distinct two-year lag, bottoming in 2022 as post-merger network retrenchment outlasted the pandemic.

The second question is answered by a within-sample comparison: we predict entry probabilities under the estimated non-stationary model and under the same model with year-specific costs pooled to carrier time-averages. Southwest illustrates the stakes: across two adjacent years its model-predicted entry probability is misstated by a stationary estimator by +50 percentage points at its 2019 cost peak and −31 points in 2020, swings the pooled model attributes to market-level heterogeneity rather than the COVID and merger episodes (Table 12).

The GFD value-difference identity of Theorem 4 makes this analysis computationally tractable.[4] Time-varying entry costs $\gamma_{i,t}$ are identified per period from the log CCP ratio without solving any game-theoretic fixed point; estimation reduces to a logit regression with a precomputed offset. This is the methodological content of the section. The economic content is the finding that stationarity is not only rejected statistically but produces qualitatively misleading predictions in years with large structural shocks.

5.1 Setup

Industry background. Mergers between 2008 and 2013 consolidated six U.S. network carriers into four: Delta absorbed Northwest in 2008, United merged with Continental in 2010, Southwest acquired AirTran in 2011, American emerged from bankruptcy and merged with US Airways in 2013. The COVID-19 demand shock in 2020 caused all four to drop service on many marginal routes; recovery through 2023 has been uneven, with Southwest the most aggressive in restoring capacity. The Big Four differ structurally: Delta, American, and United operate hub-and-spoke systems through slot-constrained fortress hubs (ATL, DFW, ORD, IAH, etc.); Southwest operates point-to-point service via focus cities (MDW, BWI, MCO, LAS, DAL, HOU).

Formal model. At the start of year $t \in \{2011, \ldots, 2023\}$, the state of airport-pair m is

$$x_{m,t} \;=\; (a_{1,t-1}, \ldots, a_{4,t-1};\, z_m) \in \{0,1\}^4 \times \mathbb{R}^2,$$

[4] Throughout this section "the GFD bridge" denotes the estimation strategy built on the value-difference identity (Theorem 4); it is distinct from the counterfactual fixed-point result (Theorem 7), for which the master text reserves the unqualified term "bridge."

where $a_{i,t-1}$ is carrier i's service indicator at $t-1$, $i \in \{\text{DL, AA, UA, WN}\}$, and $z_m = (\log \text{pax}_m, \log \text{dist}_m)$ stacks standardized log passenger volume and log great-circle distance. The endogenous state space has cardinality $|\mathcal{X}_{\text{end}}| = 2^4 = 16$. Each carrier chooses a binary action $a_{i,t} \in \{0,1\}$ (the application's instance of the action d_t of §2; $a_{i,t} = 1$ denotes serving the market) after observing $x_{m,t}$. Per-period payoff for carrier i is

$$u_i(a_{i,t}, a_{-i,t}, x_{m,t}, t; \theta_i) \;=\; a_{i,t} \cdot \Big[\gamma_{i,t} \;-\; \textstyle\sum_{j \neq i} \beta_{i,j}^{\text{comp}}\, a_{j,t} \;+\; \beta_{z,i}^{\top} z_m \;+\; \sum_{j=1}^{4} \beta_{i,j}^{\text{hub}}\, h_{j,m}\Big] \;+\; \varepsilon_{i,t}(a_{i,t}), \tag{41}$$

where $\varepsilon_{i,t}(\cdot) \overset{\text{iid}}{\sim}$ Type-I extreme value (T1EV), the outside option is normalized to $u_i(0,\cdot) = 0$, the discount factor is $\beta = 0.95$, and $h_{j,m} \in \{0,1\}$ indicates whether m's endpoints include a fortress hub for carrier j. Throughout, the bare symbol β denotes the discount factor of §2, whereas the decorated symbols $\beta_{i,j}^{\text{comp}}$, $\beta_{z,i}$, and $\beta_{i,j}^{\text{hub}}$ are flow-utility slope coefficients. Transitions are deterministic in the action ($x_{i,t} = a_{i,t-1}$); opponents' actions enter contemporaneously through β^{comp}. The *time-varying* primitive is $\gamma_{i,t}$ — a carrier × year fixed effect on entry flow utility. The *stationary* primitives are $\{\beta_{i,j}^{\text{comp}}\}_{j \neq i}$, $\beta_{z,i}$, and $\{\beta_{i,j}^{\text{hub}}\}_{j=1}^{4}$. Per carrier this is $22 = 13 + 3 + 2 + 4$ parameters.

Carrier-pair competition. We let $\beta_{i,j}^{\text{comp}}$ vary across rival j rather than imposing a single coefficient $\alpha_i \cdot n_{-i,t}$. Aggregating the three pair-specific channels in this panel produces a wrong-signed average that does not survive within-airport-pair fixed effects: at fixed own incumbency $a_{i,t-1} = 0$, the empirical entry rate for Delta rises from 5% at $n_{-i,t-1} = 1$ to 27% at $n_{-i,t-1} = 3$, and analogous patterns hold for AA and UA but not for WN. The disaggregation $\beta_{i,j}^{\text{comp}}$ is the airline-IO standard of Ciliberto and Tamer (2009) and Berry (1992).

Data. The panel is constructed from BTS DB1B Market quarterly data, aggregated to annual carrier × airport-pair × year observations and restricted to the top 193 airport-pairs by total passenger volume. After differencing out the missing-lag year (2010), the estimation panel has $193 \times 13 = 2{,}509$ market-year cells per carrier.

Is stationarity appropriate? Before estimating the non-stationary model, we test the maintained assumption of stationarity directly. Under the model, the entry-cost primitives are time-invariant if and only if $\gamma_{i,t}$ is constant across years. From the unrestricted bridge estimates and their Hessian-based covariance we compute, for each carrier, a Wald statistic for the 12 restrictions $\gamma_{i,t} = \gamma_{i,1}$ for $t = 2012, \ldots, 2023$.

The null is rejected for all four carriers at any conventional significance level (chi-squared critical value $\chi^2_{0.001}(12) = 32.9$). The rejection is especially strong for Southwest ($\chi^2 = 985.7$), reflecting the large COVID swing in $\widehat{\gamma}_{\text{WN},t}$ documented below. A stationary specification would pool 2019 pre-COVID and 2020 pandemic observations under the same intercept, misattributing the shock to cross-sectional network heterogeneity.

Table 10: Wald test of entry-cost stationarity: $H_0\colon \gamma_{i,t} = \gamma_i \;\forall\, t$ (12 restrictions per carrier). All four carriers strongly reject the null.

Carrier	$\chi^2(12)$	p-value
DL	156.2	< 0.001
AA	145.0	< 0.001
UA	204.8	< 0.001
WN	985.7	< 0.001

5.2 Identification via the value-difference identity

Finite dependence in the entry game. The four-carrier game is a multi-player shared-action (Type-D) model in the taxonomy of §4.1: rivals' actions enter carrier i's payoff contemporaneously through $\beta_{i,j}^{\text{comp}}$, and we work with the per-player perceived transition obtained by integrating out rivals' *estimated* CCPs $\widehat{p}_{-i}$. Each carrier has a renewal action—not serving ($a_{i,t} = 0$) resets its own state to $x_{i,t+1} = 0$ regardless of the current state—so finite dependence holds for the projected single-player kernel along the comparison path $a_i = 0$, which we certify by the existence test (Theorem 3) before estimating. This construction differs in scope from the re-solved game of §4.2.2, where Remark 4 requires $\rho^\star = 2$ to capture the strategic ripple through opponents' best responses. Here identification is *conditional* on the sieve CCPs $\widehat{p}_{-i}$, held fixed at their data-anchored values: we do not re-solve the Markov-perfect equilibrium, so the projected renewal certificate is used only to identify θ_i given $\widehat{p}$, not to reproduce an equilibrium object. The strategic ripple the Remark warns about is therefore absorbed into the observed $\widehat{p}_{-i}$ rather than re-derived—the price being that the θ_i we recover is interpreted relative to the empirical belief profile, consistent with the per-period reading of the identification below.

For each carrier i and each calendar year t, the GFD value-difference identity of Theorem 4 (formal statement in §3.3) gives a per-period linear-in-θ moment condition. Let $\widehat{p}_i(a \mid x, t, m)$ denote a flexible (sieve) estimate of carrier i's conditional choice probability at the (state, year, market) cell. We define the *anchor recursion*

$$\xi_i(x,t,m) \;=\; \beta\cdot\mathbb{E}_{a_{-i,t}\,|\,x,t,m;\widehat{p}_{-i}}\Big[\,\xi_i(x',t{+}1,m) - \log\widehat{p}_i\big(0 \mid x',t+1,m\big)\,\Big], \qquad \xi_i(x,T{+}1,m) \;=\; 0, \tag{42}$$

where $x' = (0, a_{-i,t})$ is the next state under the comparison sequence $a_i = 0$ from $t+1$ onward. The value-difference identity then states that

$$\log\frac{\widehat{p}_i(1 \mid x_{m,t},t,m)}{\widehat{p}_i(0 \mid x_{m,t},t,m)} \;=\; \underbrace{A_{i,m,t}^{\top}\theta_i}_{u_i^{\Delta}(x_{m,t},t;\theta_i)} \;+\; \underbrace{\beta\,\Delta_i(x_{m,t},t,m;\,\widehat{p})}_{b_{i,m,t}\ \text{(known anchor offset)}}, \tag{43}$$

where $A_{i,m,t}$ collects the year indicator, the expected rival incumbency vector, the demand

observables, and the hub indicators, and the offset

$$\Delta_i(x,t,m) \;=\; \mathbb{E}_{a_{-i,t}\,|\,x,t,m;\widehat{p}_{-i}}\Big[\xi_i\big((1,a_{-i,t}),t{+}1,m\big) - \xi_i\big((0,a_{-i,t}),t{+}1,m\big) - \log\frac{\widehat{p}_i(0\mid(1,a_{-i,t}),t+1,m)}{\widehat{p}_i(0\mid(0,a_{-i,t}),t+1,m)}\Big]$$

depends only on $\widehat{p}$ and the discount factor β, not on θ_i. Estimation of θ_i therefore reduces to a logit regression of $a_{i,m,t}$ on $A_{i,m,t}$ with known offset $b_{i,m,t}$.

Crosswalk to the general identity. The section-local symbols ξ_i and Δ_i instantiate the objects of Theorem 4, specialized to the renewal comparison path. The term $-\log\widehat{p}_i(0\mid\cdot)$ inside ξ_i is the Hotz–Miller correction $\psi_0(\mathbf{p}) = \gamma_E - \ln p$ of Assumption 5, with the Euler constant γ_E dropped because it cancels in the action-difference; ξ_i is the discounted running sum of these corrections accumulated along the comparison action $a_i = 0$, and $\beta\,\Delta_i$ is the section-local realization of the theorem's continuation term $\sum_{\tau=1}^{\rho}\beta^\tau\sum\Delta\phi\cdot\mu_{t+\tau}$. Because $a_i = 0$ is a renewal action that resets the own state to 0 (so the two comparison sequences share an identical own-state continuation from $t+1$), the finite ρ-step flow sum of Theorem 4 telescopes into the one-step-difference recursion (42), terminated by the boundary $\xi_i(\cdot,T+1,\cdot) = 0$; this is why the recursion runs to the sample horizon rather than stopping at a fixed ρ. The symbols ξ_i and Δ_i are local to this section and do not appear in §2–§3.9.

What is identified, and what is taken as input. The GFD bridge identifies the structural primitive θ_i on a per-period basis: at each calendar year t, the year-specific component $\gamma_{i,t}$ is identified from the within-year CCP moment, while the stationary primitives are identified from the cross-section variation in $(z_m, h_{j,m}, x_{m,t})$ pooled across years. The GFD bridge takes as input only $\widehat{p}_i$ over the sample window $t = 1,\ldots,T$: no backward induction of a value function, no MPE fixed-point assumption, and no specification of a terminal continuation value beyond the convention $\xi_i(\cdot,T+1,\cdot) = 0$ in (42). This last convention is innocuous for in-sample estimation because each sample observation $(x_{m,t},t)$ has its anchor offset $b_{i,m,t}$ pinned by sieve CCPs at calendar year $t+1$ that are themselves estimated from data; the terminal year T uses the boundary $\xi(\cdot,T+1,\cdot) = 0$, which only affects estimation through observations at $t = T$ via the immediate-next-period sieve CCP at T (estimated from data).

Sieve CCP estimation. We estimate $\widehat{p}_i$ by a logit on a flexible basis $B(x_{m,t},t)$ that includes year dummies, own and rival lags, the own-lag × year interaction, demand observables, hub indicators, and own-lag × hub. The fit (log-likelihoods -132 for DL through -346 for UA across 38 basis terms per carrier) matches the in-sample mean entry rate to three decimal places.

Implementation check. We verified the algebraic bridge implementation against a forward-simulation Hotz–Miller estimator in the style of Hotz et al. (1994) and Aguirregabiria and Mira

(2002), which replaces the offset Δ_i with a 200-draw simulation of trajectories from the sieve CCPs (terminal $V_T = 0$). Across all four carriers, the two methods agree on every carrier-pair competition coefficient to within 0.2 logit units and on log-likelihood to within 10 points; per-observation Δ_i correlates at 0.69–0.86 between the algebraic and Monte Carlo versions. The GFD bridge avoids the per-observation Monte Carlo entirely—one logit regression per carrier rather than a 200-draw trajectory simulation—and avoids the terminal-V assumption that the simulator inherits from the forward-simulation approach.

5.3 Estimates

Table 11 reports the stationary primitives; the year-by-year $\widehat{\gamma}_{i,t}$ path is in Appendix Table 13. Three observations.

First, 8 of 12 carrier-pair coefficients $\widehat{\beta}^{\text{comp}}_{i,j}$ are correctly signed (negative, i.e. rival j's presence reduces i's entry value). The legacy-legacy block (DL–AA, DL–UA, AA–UA in both directions) loads on standard rivalry with magnitudes -0.4 to -1.6. The UA–WN pair, where UA's DEN hub overlaps WN's DEN and PHX focus cities, is also correctly signed in both directions.

The four positively signed coefficients are the bidirectional DL–WN and AA–WN pairs ($+1.3$ to $+3.1$). The likely cause is market-quality co-movement: DL/AA fortress hubs and WN focus cities are geographically disjoint, so airport pairs served by both are disproportionately high-demand corridors, and the demand controls do not fully absorb this correlated quality. We tested whether a standard cost-shifter instrument separates rivalry from quality. Following Berry (1990) and Ciliberto and Tamer (2009), we instrument each rival's entry with its *airport presence*—the breadth of its network at the route endpoints, constructed from the full national DB1B network—which shifts the rival's entry cost while being plausibly excluded from the focal carrier's route quality. The instrument is strong (first-stage F exceeding 100 for every carrier), but a control-function correction leaves the positive coefficients essentially unchanged ($\widehat{\beta}^{\text{comp}}_{\text{DL,WN}}$: $+1.495 \rightarrow +1.594$; $\widehat{\beta}^{\text{comp}}_{\text{AA,WN}}$: $+1.430 \rightarrow +1.431$). The reason is structural: the GFD bridge integrates rivals through their estimated CCPs, which are themselves a market-quality index, so a cost-shifter instrument for the rival's *observed* entry does not break the quality channel. We therefore treat the carrier-pair terms as *controls* and do not interpret them as structural competition parameters. Credibly identifying competition here would require an independent (demand-side) measure of market quality—a separate exercise that this paper, whose object is the entry-cost path $\gamma_{i,t}$, does not undertake.

Second, the hub coefficients have the expected geography. Own hub adds $+2.50$ logit units for Delta and $+0.49$ for United; for American and Southwest the own-hub coefficients are indistinguishable from zero after carrier-pair competition and demand observables are controlled for.

Third, Figure 1 plots the full $\widehat{\gamma}_{i,t}$ paths. The COVID shock registers cleanly: Southwest's path drops from its 2019 peak ($\widehat{\gamma}_{\text{WN},2019} = +2.88$) to $\widehat{\gamma}_{\text{WN},2020} = -3.55$, a swing of 6.4 logit units,

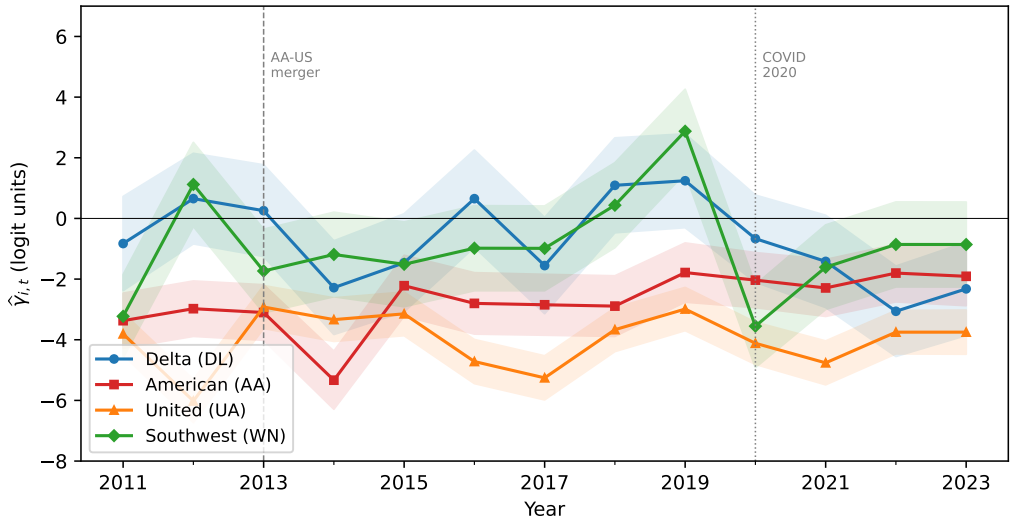

Figure 1: Estimated entry-cost paths $\widehat{\gamma}_{i,t}$ for the four carriers, 2011–2023. Shaded bands: 95% pointwise confidence intervals (Hessian-based SE). Dashed vertical: 2013 AA–US Airways merger. Dotted vertical: 2020 COVID demand shock. Full numerical values in Appendix Table 13.

reflecting its high exposure to leisure travel. Delta's deepest year is 2022 ($\widehat{\gamma}_{\text{DL},2022} = -3.06$), consistent with post-merger network retrenchment lagging the demand shock by two years. Under a stationary specification these differential shifts would be pooled into a single carrier-specific intercept γ_i, conflating pandemic disruption with pre-existing network heterogeneity (cf. Table 10).

Magnitudes against the airline-IO literature. The coefficients we interpret structurally have the expected sign and order of magnitude. Delta's positive own-hub entry advantage ($\widehat{\beta}^{\text{hub}}_{\text{DL,own}} = +2.50$) is consistent with the large hub effects in the dynamic airline entry model of Aguirregabiria and Ho (2012) and, qualitatively, with the hub cost/markup advantages documented by Berry et al. (2006) (whose estimates are not in directly comparable units). The demand coefficients enter with the expected signs. (The carrier-pair terms are controls, not interpreted here; see §5.3.)

5.4 The cost of imposing entry-cost stationarity

The Wald test of Table 10 established that entry costs shift non-stationarily; the natural follow-up is *how much does this matter for empirical predictions?* We construct the following within-sample comparison. For each carrier i and observation (m, t) we form two predicted entry probabilities. The non-stationary prediction $\widehat{p}^{\text{ns}}_i$ uses the full estimated $\widehat{\gamma}_{i,t}$. The stationary prediction $\widehat{p}^{\text{s}}_i$ replaces $\gamma_{i,t}$ with the carrier-specific time average $\bar{\gamma}_i = T^{-1}\sum_t \widehat{\gamma}_{i,t}$, holding all other primitives $\{\widehat{\beta}^{\text{comp}}_{i,j}, \widehat{\beta}_{z,i}, \widehat{\beta}^{\text{hub}}_{i,j}\}$ and the anchor offset $\widehat{b}_{i,m,t}$ fixed. Both probabilities are the

Table 11: GFD bridge estimates of stationary primitives. Standard errors (Hessian-based) in parentheses; the "avg rival hub" column averages the three rival-hub coefficients and its SE is $\sqrt{\sum_j \widehat{\mathrm{var}}(\widehat{\beta}^{\mathrm{hub}}_{i,j})}/3$ (treating them as independent). Top block: carrier-pair competition $\widehat{\beta}^{\mathrm{comp}}_{i,j}$ (row i, column j; diagonals not in model). Bottom block: demand-shifter and hub coefficients.

Panel A. Carrier-pair competition $\widehat{\beta}^{comp}_{i,j}$ (controls; see text)				
Carrier i	rival DL	rival AA	rival UA	rival WN
DL	—	−1.370	−0.370	+1.495
		(0.307)	(0.333)	(0.620)
AA	−0.614	—	−1.530	+1.430
	(0.285)		(0.224)	(0.298)
UA	−0.408	−1.627	—	−0.251
	(0.262)	(0.212)		(0.183)
WN	+3.140	+1.324	−0.821	—
	(0.704)	(0.315)	(0.186)	
Panel B. Demand and hub coefficients				
Carrier i	$\widehat{\beta}_{\mathrm{pax},i}$	$\widehat{\beta}_{\mathrm{dist},i}$	$\widehat{\beta}^{\mathrm{hub}}_{i,\mathrm{own}}$	avg rival hub
DL	+0.936	+1.817	+2.503	+0.087
	(0.158)	(0.193)	(0.392)	(0.181)
AA	+0.476	+0.273	−0.044	+0.165
	(0.109)	(0.099)	(0.186)	(0.109)
UA	+0.090	+0.014	+0.486	−0.399
	(0.074)	(0.077)	(0.208)	(0.085)
WN	−0.008	+0.197	+0.176	−0.183
	(0.053)	(0.057)	(0.131)	(0.076)

estimator's genuine predicted CCPs, $\Lambda(A^{\top}_{i,m,t}\theta + \widehat{b}_{i,m,t})$; the only difference is the replacement $\gamma_{i,t} \mapsto \bar{\gamma}_i$, so the gap $\widehat{p}^{\mathrm{ns}}_i - \widehat{p}^{\mathrm{s}}_i$ isolates the entry-cost channel alone. This is a *per-cell* diagnostic: rivals' sieve CCPs and the observed state distribution are held fixed—we do not re-solve the equilibrium when γ is pooled—so the gap measures how much the stationarity restriction alone moves the model's own predictions, not a long-run policy counterfactual.

Table 12: Cost of imposing stationarity: mean predicted entry probability under non-stationary ($\widehat{p}^{\mathrm{ns}}$) versus stationary ($\widehat{p}^{\mathrm{s}}$) entry costs, for selected years. Gap in percentage points (pp). Carrier time-averages: $\bar{\gamma}_{\mathrm{DL}} = -0.75$, $\bar{\gamma}_{\mathrm{AA}} = -2.72$, $\bar{\gamma}_{\mathrm{UA}} = -4.02$, $\bar{\gamma}_{\mathrm{WN}} = -0.93$.

Year	Carrier	$\widehat{p}^{\mathrm{ns}}$	$\widehat{p}^{\mathrm{s}}$	Gap (pp)
2019	WN	0.684	0.180	+50.4
2020	WN	0.617	0.924	−30.7
2019	DL	0.896	0.819	+7.7
2022	DL	0.871	0.941	−7.1
2019	UA	0.715	0.584	+13.1
2014	AA	0.803	0.939	−13.6

Southwest produces the largest distortions, and in opposite directions in adjacent years. At its 2019 cost peak ($\widehat{\gamma}_{\mathrm{WN},2019} = +2.88$, 3.8 logit units above its $\bar{\gamma}_{\mathrm{WN}} = -0.93$ average) the non-stationary model predicts 68% average entry while the pooled model predicts only 18% —

a +50 pp gap. One year later, at the COVID trough ($\widehat{\gamma}_{\text{WN},2020} = -3.55$), the non-stationary model predicts 62% and the pooled model 92% — a −31 pp gap. The stationary restriction thus swings Southwest's predicted entry probability by roughly 80 percentage points across two years, attributing the year-to-year movement to the cross-sectional distribution of market characteristics rather than to the entry-cost path. For the legacy carriers the distortions are an order of magnitude smaller (Delta ±7–8 pp; American and United ±13–14 pp in their most variable years), tracking their smaller $\widehat{\gamma}_{i,t}$ variation.

These distortions have direct consequences for any downstream analysis that conditions on predicted CCPs — counterfactual entry predictions at COVID-era market states, or any object that weights entry rates by year — which will be biased if the underlying model imposes stationarity over a panel spanning the merger wave and the pandemic.

5.5 What we learn about airline entry

Three findings emerge that are of independent interest to the airline-IO literature.

Entry costs respond asymmetrically to structural shocks. The 2020 COVID demand shock lowers all carriers' entry costs, but not uniformly: Southwest's $\widehat{\gamma}_{i,t}$ falls by 6.4 logit units (2019 to 2020), while Delta's lowest point occurs in 2022 ($\widehat{\gamma}_{\text{DL},2022} = -3.06$), two years after the shock. The lagged DL response is consistent with the slow pace of post-merger network integration. The asymmetry has direct implications for policy analysis: competition authorities evaluating post-COVID slot reallocations face a setting where the four carriers are simultaneously retrenching from different cost levels, not a common shock to a stationary equilibrium.

Hub coefficients match the literature, and the competition controls carry the expected sign. Delta's positive own-hub coefficient (+2.5) accords with the large hub advantages in the dynamic airline entry model of Aguirregabiria and Ho (2012), and Southwest's near-zero own-hub coefficient (+0.18, SE 0.13) is consistent with its point-to-point operating model, which does not generate the hub-and-spoke network externalities that drive the DL and UA estimates. As a plausibility check on the control block, the legacy-legacy carrier-pair coefficients carry the expected negative sign (ranging −0.4 to −1.6, broadly in the range of Aguirregabiria and Ho (2012) and within the set of Ciliberto and Tamer (2009)); we use them only as controls, not as interpreted competition estimates (the DL/AA–WN pairs are confounded; see below).

Competition enters as a control, not a result. The positively signed DL/AA–WN carrier-pair coefficients are a market-quality confound (§5.3): a standard airport-presence cost-shifter instrument, though strong in the first stage, does not separate rivalry from quality co-movement, because the bridge integrates rivals through their quality-laden CCPs. We therefore read the carrier-pair terms as controls. Credibly identifying competition would require an independent demand-side measure of market quality—a structural demand model in the spirit of Berry et al. (2006)—which is beyond the scope of this illustration and a natural subject for separate work.

Relation to the non-stationary estimation literature. Most prior non-stationary dynamic

game studies (Igami, 2017, 2018; Igami and Uetake, 2020; Doraszelski et al., 2018) solve the game's value function by backward induction at each candidate parameter vector; Sweeting (2013) combines backward induction with two-step/CCP elements. The GFD bridge replaces the value-function solve with a single regression per carrier: the backward-induction computation reduces to computing the anchor offset $\widehat{b}_{i,m,t}$ once from the sieve CCPs, after which identification is linear. The cost is the boundary requirement at $t = T$: in-sample predictions are data-anchored through year T, while out-of-sample prediction requires terminal-period CCPs or an explicit extrapolation rule — on the same footing as any estimator that relies on a continuation-value specification.

6 Conclusion

This paper develops a computable framework for finite dependence in dynamic discrete choice models, organized around four results. Theorem 3 reduces existence of ρ-period finite dependence to a pointwise linear feasibility test $\mathbf{Ax} = \mathbf{b}$ on the structural transitions, giving practitioners a pre-estimation diagnostic at every (x_t, d_t, d'_t) triple. Definition 3 packages the resulting GFD identity into an estimator that takes the FD-feasible flow input $\widehat{\mathbf{\Phi}}$ as an explicit argument, separating instrument selection from the rest of the procedure. Theorem 7 shows that the same identity yields a Bellman-free fixed-point system for payoff-only counterfactual CCPs, with restart-based validation against extraneous fixed points. Theorem 6 characterizes the asymptotically variance-optimal flow input as the trace-minimizing solution over a compact regular sub-region of the FD-feasible variety.

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. (Non-stationary, time-varying transitions are within scope and are central to the framework—the existence test, the GFD identity, and the airline application of Section 5 all allow f_t to vary with t.) Second, the GFD counterfactual identity (Theorem 7) 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 nine canonical models verified in Section 4.1, but pushing ρ above five in dense-transition models would require either a sparser reformulation or stronger structural restrictions. Each of these directions is a natural target for future work.

An invitation to applied work. The framework is designed to slot into existing structural DDC pipelines with minimal modification. Three concrete invitations follow. *First*, before committing to a nested fixed-point estimator, run the pointwise feasibility test of Theorem 3 on the estimated transitions; a positive verdict at $\rho \leq 2$—which holds for eight of the nine canonical models verified in Section 4.1, including ones with multi-component Kronecker-separable state

spaces and shared-action multi-player games—makes GFD an immediate drop-in alternative with the same observable inputs and a substantial speed gain over Bellman-based benchmarks: $\sim$370–1,000$\times$ faster than CCP-2step and NFXP on the small-state single-agent investment calibration of Table 3, $\sim$4$\times$ faster than NFXP on the multi-player Markov-perfect entry/exit game of Table 4, and—when one component of the state is action-invariant, so that the structured Kronecker FD solver of Appendix B applies—up to $\sim$136$\times$ faster than NFXP on the single-agent investment DGP at $|\mathcal{X}| = 5{,}000$, with a speedup that grows super-linearly in $|\mathcal{X}|$ (Table 16). *Second*, when conducting the payoff-only counterfactual policy analysis that dominates the applied DDC literature—changes in subsidies, fees, taxes, entry costs, or other monetary instruments—use the GFD counterfactual identity of Theorem 7 to compute counterfactual conditional choice probabilities without re-solving the dynamic program at the counterfactual payoff, with the restart and spectral-radius diagnostics of Section 3.8 as the validation scaffold against extraneous fixed points. *Third*, when finite dependence holds with multiple feasible flow inputs—the generic case at $\rho \geq 2$—the variance-optimal selection of Theorem 6 costs one inner optimization in the null-space coordinate per initial state and recovers the FD-class trace minimum at no additional first-stage cost. A replication package accompanying this paper provides production-ready Python implementations of all three procedures together with the nine-canonical-model verification of Section 4.1, the Monte Carlo benchmarks of Section 4.2, and the counterfactual demonstrations of Section 4.3.

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

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 (2002). Swapping the nested fixed point algorithm: A class of estimators for discrete markov decision models. *Econometrica 70*(4), 1519–1543.

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.

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

Berry, S. T. (1990). Airport presence as product differentiation. *American Economic Review 80*(2), 394–399.

Berry, S. T. (1992). Estimation of a model of entry in the airline industry. *Econometrica 60*(4), 889–917.

Berry, S. T., M. Carnall, and P. T. Spiller (2006). Airline hubs: Costs, markups and the implications of customer heterogeneity. *Advances in Airline Economics 1*, 183–214.

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.

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

Ciliberto, F. and E. Tamer (2009). Market structure and multiple equilibria in airline markets. *Econometrica 77*(6), 1791–1828.

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

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

Hotz, V. J., R. A. Miller, S. Sanders, and J. Smith (1994). A simulation estimator for dynamic models of discrete choice. *Review of Economic Studies 61*(2), 265–289.

Hu, Y. and M. Shum (2012). Nonparametric identification of dynamic models with unobserved state variables. *Journal of Econometrics 171*(1), 32–44.

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.

MaCurdy, T. E. (1981). An empirical model of labor supply in a life-cycle setting. *Journal of Political Economy 89*(6), 1059–1085.

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.

Scott, P. T. (2013). Dynamic discrete choice estimation of agricultural land use. *Toulouse School of Economics Working Paper*.

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 Supplementary Tables for Section 5

Table 13: Estimated year-by-year entry-cost shifters $\widehat{\gamma}_{i,t}$ from the non-stationary GFD bridge (§5.2). The 2020 COVID dip is clearest for WN; DL's deepest year is 2022; AA stays comparatively flat post-2014. Stationary primitives in Table 11.

Year	DL	AA	UA	WN
2011	−0.83	−3.37	−3.81	−3.23
2012	+0.65	−2.98	−6.02	+1.12
2013	+0.26	−3.10	−2.92	−1.73
2014	−2.28	−5.33	−3.34	−1.19
2015	−1.46	−2.22	−3.15	−1.51
2016	+0.65	−2.80	−4.72	−0.98
2017	−1.55	−2.85	−5.26	−0.99
2018	+1.09	−2.89	−3.67	+0.44
2019	+1.24	−1.78	−2.99	+2.88
2020	−0.67	−2.03	−4.12	−3.55
2021	−1.42	−2.29	−4.76	−1.61
2022	−3.06	−1.80	−3.75	−0.86
2023	−2.32	−1.91	−3.75	−0.86

B Scalable Computation for Large State Spaces

The joint feasibility QP (28) has decision variable $\operatorname{vec} \mathbf{\Phi}(x) \in \mathbb{R}^{D \cdot S^{\rho} \cdot D^{\rho}}$ per state x, and the associated history tree has $\sum_{\tau=0}^{\rho-1} S \cdot (D \cdot S)^{\tau}$ nodes. For $S = 20$, $D = 3$, and $\rho = 3$ this is 73,220 nodes per state and 219,660 flow variables: prohibitive in dense memory but tractable once two structural features of applied DDC models are exploited.

B.1 Reachability pruning

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

$$\mathcal{R}(d_t, x) \;:=\; \{x' \in \mathcal{X} : f_t(x' \mid x, d_t) > 0\}, \tag{44}$$

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. For the investment model ($S = 20$, $D = 3$, $\rho = 3$, with deterministic capital transitions and one Tauchen-stochastic productivity factor) the node count drops from 73,220 to 3,140, a 96% reduction.

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

on small models.

B.2 Iterative KKT solving via LSQR

For small problems the QP (28) is solved by forming the KKT system

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

where $\mathbf{C}$ maps flow variables to the terminal-distribution difference $\kappa_{t+\rho+1}(\cdot \mid x_t, d_t) - \kappa_{t+\rho+1}(\cdot \mid x_t, d_t')$, ε is a Tikhonov regularizer, and $\delta \approx 0$ ensures non-singularity. This requires forming $\mathbf{C}^\top\mathbf{C}$ explicitly, which is dimension $n_{\text{tot}} \times n_{\text{tot}}$ and becomes near-dense when $S \ll n_{\text{tot}}$.

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

$$\min_{\mathbf{w}} \left\| \begin{pmatrix} \sqrt{2}\,\mathbf{C} \\ \sqrt{\varepsilon}\,\mathbf{I} \\ \eta\,\mathbf{A}_{\text{eq}} \end{pmatrix} \mathbf{w} \;-\; \begin{pmatrix} \mathbf{0} \\ \mathbf{0} \\ \eta\,\mathbf{b}_{\text{eq}} \end{pmatrix} \right\|_2^2, \tag{46}$$

with $\eta \gg 1$ enforcing the equality constraints. Solving (46) via LSQR (Paige and Saunders, 1982) requires only matrix–vector products with the (sparse) stacked operator and never forms the normal equations, so memory usage is $O(\text{nnz})$ rather than $O(n_{\text{tot}}^2)$. The solver switches automatically between the direct KKT system (45) (used for $n_{\text{tot}} \le 5{,}000$) and the iterative formulation (46) otherwise.

B.3 Empirical performance

Table 14 reports the combined effect of pruning and iterative solving on the investment model ($S = 20$, $D = 3$, $\beta = 0.95$). At $\rho = 3$ the unpruned full tree exceeds available memory on a standard workstation; with pruning and LSQR the same problem solves in under one second to machine-precision matching accuracy.

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

Table 14: Combined effect of reachability pruning and iterative LSQR solving on the investment model ($S = 20$, $D = 3$, $\beta = 0.95$). Without pruning, $\rho = 3$ runs out of memory; with pruning and LSQR it completes in 0.5 seconds with terminal-distribution matching error $\|\kappa_{t+\rho+1}(\cdot \mid x_t, d_t) - \kappa_{t+\rho+1}(\cdot \mid x_t, d_t')\|_1 < 10^{-15}$.

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

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

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

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

$$\begin{pmatrix} 2\,\mathbf{C}^\top\mathbf{C} + \varepsilon\mathbf{I} & \mathbf{A}_{\text{eq}}^\top \\ \mathbf{A}_{\text{eq}} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \boldsymbol{\lambda} \end{pmatrix} \;=\; \begin{pmatrix} \mathbf{0} \\ \mathbf{b} \end{pmatrix}. \tag{47}$$

Both $\mathbf{C}$ and $\mathbf{A}_{\text{eq}}$ are sparse for any sparse-transition DGP (Tauchen-discretised AR(1), deterministic shifts, Kronecker-separable factor models), so $\mathbf{C}^\top\mathbf{C} + \varepsilon\mathbf{I}$ remains sparse and the KKT block is symmetric-indefinite-but-sparse. We factor it once via `scipy.sparse.linalg.splu` (using the default UMFPACK backend) and then triangular-solve for each right-hand side $(\mathbf{0}; \mathbf{b}^{(x_t,k)})$. Empirical fill-in for the investment DGP is moderate, with factorisation cost that scales subcubically in $|\mathcal{X}|$ and triangular-solve cost that is essentially linear in $|\mathcal{X}|$ per RHS; end-to-end the path scales considerably better than the dense pseudoinverse for general (non-Kronecker-separable) DGPs. For DGPs that admit the Kronecker factorisation (48) of A.5 below—the case in Table 16—the structured solver dominates this KKT-LU path by another factor of $\sim Z^c$.

Why not normal equations. The standard alternative, $\mathbf{x} = (\mathbf{A}_{\text{aug}}^\top\mathbf{A}_{\text{aug}})^{-1}\mathbf{A}_{\text{aug}}^\top\mathbf{b}$ via sparse splu, squares the condition number of $\mathbf{A}_{\text{aug}}$. With the penalty weight $w = 10^6$ already needed to enforce the constraints to acceptable precision, $\text{cond}(\mathbf{A}_{\text{aug}}) \sim 10^6$ becomes $\text{cond}(\mathbf{A}_{\text{aug}}^\top\mathbf{A}_{\text{aug}}) \sim 10^{12}$, which a small Tikhonov ridge cannot regularise without introducing significant bias; in our trials this path produced estimates that diverged numerically at $|\mathcal{X}| = 200$. The KKT formulation (47) avoids the squaring entirely and is the recommended sparse path.

B.5 Structured Kronecker solver: when the state factors as endo × exo

The dominant scaling exploit, used to produce Table 16, is to recognise that many applied DDC transitions factor as a Kronecker product

$$f(s' \mid s, a) \;=\; T_{\text{endo}}(k' \mid k, a) \;\cdot\; T_{\text{exo}}(o' \mid o), \tag{48}$$

with $s = (k, o)$, where k is an endogenous state (responding to actions) and o is an action-invariant exogenous state. The investment model has $k =$ capital and $o =$ Tauchen-discretised productivity; the labour-supply model has $k =$ labour history and $o =$ wage shock; multi-agent games have player-i state as k and opponent-state-and-environment as o at the equilibrium-induced perceived transition. Lemma 11 formalises the reduction.

Lemma 11 (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 $\rho^\star_{\text{joint}} = \rho^\star_{\text{endo}}$ regardless of the persistence of $\mathbf{F}^{\text{exo}}$.*

That is, under (48), the FD problem on the full state space reduces to the FD problem on the endogenous component k alone: the exogenous factor T_{exo} cancels identically between the two compared actions. Concretely, the optimal flow weight at the full state factors as

$$\mathbf{\Phi}^{\text{full}}\big((k, o),\, a;\, (k', o')\big) \;=\; \mathbf{\Phi}^{\text{endo}}(k,\, a;\, k') \;\cdot\; T_{\text{exo}}(o' \mid o), \tag{49}$$

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

Effective regressors via factored Kronecker form. With $\mathbf{\Phi}^{\text{endo}}$ in hand, the GFD-MLE effective regressor $\widetilde{H}[s, a, r]$ at full state $s = (k, o)$ decomposes by basis component r via the conditional expectation operator $E_{o'|o}[\cdot] := \sum_{o'} T_{\text{exo}}(o' \mid o) \cdot$. For the investment basis $z[(k, o), a, \cdot] = (o \cdot k,\, -a,\, -a^2)$:

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

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

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

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

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

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

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

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

where c is the per-QP-size scaling exponent ($c = 3$ for dense pseudoinverse, $c \approx 1.5$–2 for sparse KKT-LU). On the investment DGP at $|\mathcal{X}| = 5{,}000 = 100 \cdot 50$, this predicts a $\sim Z^4 = 6{,}250{,}000\times$ reduction relative to dense pinv on the full state and $\sim Z^3 = 125{,}000\times$ reduction relative to KKT-LU on the full state. The empirical reduction is more modest because of constant-factor overhead (matrix construction, basis precomputation), but Table 16 shows that the structured solver brings GFD wall time at $|\mathcal{X}| = 5{,}000$ from 111s (KKT-LU on full state) down to 1.39s, an $80\times$ reduction in this single benchmark.

Table 15 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.

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 15: 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 11. Wall-clock is mean of 3 repeats on Apple M3 Ultra.

Table 16 sweeps the joint state space from $|\mathcal{X}| = 20$ to $|\mathcal{X}| = 5{,}000$ at fixed sample size ($N = 1000, T = 15$).

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

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

Table 16: Scaling of GFD vs. NFXP on the investment DGP across joint state-space size. Same DGP as Table 3; $|\mathcal{X}|$ varied via capital and productivity grid sizes ($5\times4, 10\times6, 20\times10, 30\times20, 50\times40, 100\times50$). RMSE is mean across the three structural parameters; 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.

through the cached offset $\widetilde{h}_x$ of (33). With the empirical-frequency estimator

$$\widehat{p}^{\,\mathrm{emp}}(d_t \mid x_t) \;=\; \frac{n_{x,d_t}}{n_x}, \qquad n_{x,d_t} := \sum_{i,t} \mathbf{1}\{x_{it} = x,\; d_{it} = d_t\}, \tag{50}$$

the per-state variance is $\mathrm{Var}[\widehat{p}^{\,\mathrm{emp}}(d_t \mid x_t)] = p_{0,t}(d_t \mid x_t)(1 - p_{0,t}(d_t \mid x_t))/n_x$, 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 $|\mathcal{X}| = 2{,}000$, and the empirical-CCP first stage drives the GFD-MP RMSE divergence reported in the central column of the diagnostic panel of Section 4.2.1.

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

We replace (50) with a parametric multinomial logit fit on a low-dimensional basis $\mathbf{z}(x) \in \mathbb{R}^{p_{\mathrm{dim}}}$ of the state. For the investment DGP of Section 4.2.1 we use

$$\mathbf{z}(x) \;=\; \big(1,\, k(x),\, o(x),\, k(x)\cdot o(x)\big)^{\top}, \qquad p_{\mathrm{dim}} = 4, \tag{51}$$

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

$$\ell(\boldsymbol{c}) \;=\; \sum_{i,t} \log \frac{\exp\big(\mathbf{z}(x_{it})^{\top}\boldsymbol{c}_{d_{it}}\big)}{\sum_{a'=0}^{D-1}\exp\big(\mathbf{z}(x_{it})^{\top}\boldsymbol{c}_{a'}\big)}, \qquad \boldsymbol{c}_0 \equiv \mathbf{0}, \tag{52}$$

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:

$$\widehat{p}^{\,\mathrm{logit}}(d_t \mid x_t) \;=\; \frac{\exp\big(\mathbf{z}(x)^{\top}\widehat{\boldsymbol{c}}_{d_t}\big)}{\sum_{a'=0}^{D-1}\exp\big(\mathbf{z}(x)^{\top}\widehat{\boldsymbol{c}}_{a'}\big)}, \qquad x \in \mathcal{X}. \tag{53}$$

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

C.2 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 Section 4.2.1, this gives a per-state variance reduction factor of approximately $|\mathcal{X}|/p_{\text{dim}} = 500$ at $|\mathcal{X}| = 2{,}000$, which is what drives the RMSE collapse from 1.08 (empirical) to 0.19 (smoothed) reported in Table 16.

C.3 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 $\mathbf{z}$ (the structural payoff is linear in the same $(k, o, k \cdot o)$ regressors). When this exact-parametric assumption fails, three robust alternatives preserve the variance gain:

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

2. *Kernel smoothing.* Replace (50) 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 $\alpha \in (0, 1]$. Cheap; partly mitigates the zero-cell problem but does *not* reduce the $|\mathcal{X}|/(NT)$ variance scaling.

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

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

This appendix illustrates the GFD framework on a published structural labour supply model. Altug and Miller (1998) (AM) is the benchmark study that introduces finite dependence as

a route to estimation of dynamic discrete choice female labour-supply models, and 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 AM 1998's published Table VI 2.a values; we report on which parameters our estimates agree with AM in sign and which do not, and we enumerate the specific implementation simplifications relative to AM'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 $x_{nt} = (l_{n,t-1}, l_{n,t-2}, l_{n,t-3}) \in \{0,1\}^3$ encodes the past three years of participation, giving $|\mathcal{X}| = 8$ and $\rho^\star = 3$ (the canonical Type-S shift-register horizon). 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 $\mathbf{z}_{nt} = (1, \mathrm{KIDS}_{nt}, \mathrm{KIDS}_{nt}\cdot\mathrm{MAR}_{nt}, \mathrm{AGE}_{nt}, \mathrm{AGE}_{nt}^2)$ and $\theta = (\mathbf{B}_0, \mathbf{B}_1, \delta_0, \delta_1, \delta_2, \delta_3) \in \mathbb{R}^{14}$. The continuation correction $C(x_{nt})$ is the only object that distinguishes the two estimators we compare: AM 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, parse SAS dictionaries to per-year CSVs, and construct a female panel with the AM 1998 sample-selection criteria (restricting 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/archived/altug_miller_replication/` document the full extraction pipeline.

Estimation results. We run AM-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, and differ only in the continuation correction. Table 17 reports the 14-parameter vector together with AM 1998's published Table VI 2.a values.

Findings. GFD and AM-CCP agree within $\sim$10% on all 14 parameters, verifying the consistency-for-any-FD-input result on real PSID data: two different FD-feasible $\mathbf{\Phi}$ choices (AM 1998's path-based form and our canonical Moore–Penrose form) deliver the same θ in finite sample

Parameter	AM-CCP	GFD	AM 1998 Tab. VI 2.a	sign-match?	GFD/AM-CCP
B_{00} (const)	+2.11	−0.14	−14.0	× AM-CCP / ✓ GFD	—
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 17: 14-parameter participation logit estimates on PSID 1973–1985: AM-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 $C(x_{nt})$. AM 1998 published values are from Table VI col 2.a. GFD and AM-CCP agree to within $\sim 10\%$ on all 14 parameters (rightmost column), verifying the consistency-for-any-FD-input result of Theorem 5 on real PSID data.

up to first-stage noise. Both estimators agree with the published AM 1998 Table VI 2.a values in sign on 7 of 14 parameters (most notably the intertemporal complementarity δ_1, δ_3, the demographic–hours interactions $B_{11}, B_{12}, B_{13}, B_{14}$, and the AM-acknowledged counter-intuitive own-period hours quadratic δ_0). Magnitudes differ substantially from AM 1998's published values, with the residual gap traceable to a small number of AM-specific implementation choices our pipeline does not reproduce.

Sources of the residual magnitude gap. The following AM-specific implementation details, present in AM's full pipeline but not in our minimum-faithful reproduction, account for the gap: (i) kernel-density derivative terms f'/f in Step A's hours Euler equation (AM 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 Heckman–MaCurdy 1980 traditional fixed effects (per-woman time averages of level residuals). (iii) Implicit hours-unit conventions in AM's published Table VI that we cannot fully infer from the published text; our pipeline uses hours in thousands ($l_k = l/1000$), but if AM's published coefficients are on raw annual hours scale, the δ coefficients would differ by 10^6 from the conversion alone (partially offsetting 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 AM's implicit scale; AM'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 (AM-CCP and GFD) on PSID data, rather than a numerical replication of AM

1998's published Table VI 2.a values.

E Proofs

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

The proof uses two marginal flows. Throughout, $\tau = 1, \ldots, \rho$ 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.

(⇐). Substitute the product form (9) into the LHS of the initial-flow constraint (10). Summing over $(x_{t+2:t+\rho}, d_{t+1:t+\rho})$, the factors $f_{t+\tau-1}(x_{t+\tau} \mid x_{t+\tau-1}, d_{t+\tau-1})$ for $\tau \geq 2$ and the weights $\omega_{t+\tau}$ for $\tau \geq 1$ each sum to one in turn (transitions sum to one over the next state; weights sum to one over the next action), and the only surviving factor is $f_t(x_{t+1} \mid x_t, d_t)$, matching the RHS. An analogous telescoping verifies the flow-conservation constraint (11).

(⇒). 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 $\sum_{d_{t+\tau}} \omega_{t+\tau}(d_{t+\tau} \mid \cdot) = 1$. Rewriting flow conservation (11) as $M^{\mathrm{s}}_{\tau+1}(x_{t+1:t+\tau+1}, d_{t+1:t+\tau}) = f_{t+\tau}(x_{t+\tau+1} \mid x_{t+\tau}, d_{t+\tau}) \cdot M^{\mathrm{sa}}_{\tau}(x_{t+1:t+\tau}, d_{t+1:t+\tau})$, and combining with the initial-flow base case $M^{\mathrm{s}}_{1}(x_{t+1}) = f_t(x_{t+1} \mid x_t, d_t)$, gives the chain identity $\phi = \prod_{\tau} f_{t+\tau-1} \cdot \omega_{t+\tau}$ by telescoping. Uniqueness on non-zero-marginal histories is immediate from the definition of $\omega_{t+\tau}$ as a ratio of marginals. □

E.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. □

E.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 in Definition 1. 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), $\mathbf{T}_{\text{last}} \operatorname{vec} \Phi(x_t, d_t) = \kappa_{t+\rho+1}(\cdot \mid x_t, d_t)$, so 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 (15) holds.

Hence $\mathbf{A}\,\mathbf{x} = \mathbf{b}$ is equivalent to Conditions 1–3 of Definition 1, and the two-way characterization (21) follows. The constructive statement (second sentence of the theorem) is immediate from the equivalence. □

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

Sketch. Iterate the AM representation (3) of Lemma 1 ρ times along the path measure encoded by $\mathbf{\Phi}$. After ρ substitutions, the choice-specific value function $v_t(x_t, d_t; \theta)$ becomes

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

with $\mu_{t+\tau}(x_{t+\tau}, d_{t+\tau}; \theta, \mathbf{p}) = z_{t+\tau}(x_{t+\tau}, d_{t+\tau})^\top \theta + \psi_{d_{t+\tau}}(\mathbf{p}_{t+\tau}(x_{t+\tau}))$. Taking the difference of this identity at d_t and the reference action 0 and invoking the common-terminal condition (23) (which makes the V-term cancel) yields (24). □

E.5 Proof of Theorem 7 (bridge: payoff-only counterfactual CCPs from FD)

Single-agent case (Theorem 7). Part (a): substitute (24) (evaluated at the counterfactual payoff θ_1 and CCPs $\mathbf{p}_1$) into the Hotz–Miller logit-CCP map at θ_1, giving (37). Part (b): the standard Bellman contraction plus the Hotz–Miller bijection deliver uniqueness of $\mathbf{p}_{1,t}$, and Theorem 4 guarantees the unique Bellman solution satisfies (37); local convergence of the FD iteration follows from the spectral-radius condition on the composite map's Jacobian.

Game case (Proposition 8). Fix the selected counterfactual profile $\mathbf{p}_1$ of Assumption 9. 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 to which Theorem 4 applies, since $\mathbf{\Phi}_i$ is FD-feasible for $f_i^{\mathbf{p}_{-i,1}}$. This yields the per-player identity (38); the equilibrium continuation value is exactly the object the FD-weighted sum reproduces under exact FD at $f_i^{\mathbf{p}_{-i,1}}$. 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 9 (not by a contraction). □

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

The GFD estimator is a two-step M-estimator, $\widehat{\theta}^{\mathrm{GFD}}(\widehat{\mathbf{\Phi}}) = \arg\max_{\theta\in\Theta} Q_N(\theta; \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}})$, with criterion $Q_N(\theta; \mathbf{\Phi}, \mathbf{p}) = N^{-1}\sum_{n=1}^{N}\sum_t \log \Lambda_{d_{nt}}(x_{nt}; \theta, \mathbf{\Phi}, \mathbf{p})$, where Λ is the GFD logit map of Definition 3 and 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, so $\Lambda_d(x; \theta_0, \mathbf{\Phi}, \mathbf{p}_0) = p_0(d \mid x)$, the true CCP, *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 identification condition that $\theta \mapsto \Lambda(\cdot; \theta, \mathbf{\Phi}, \mathbf{p}_0)$ is injective (implied by non-singular $\mathbf{J}(\mathbf{\Phi})$). Continuity of Λ in $(\theta, \mathbf{\Phi}, \mathbf{p})$, compactness of Θ, a uniform law of large numbers, and $(\widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) \xrightarrow{p} (\mathbf{\Phi}^*, \mathbf{p}_0)$ give $\sup_\theta |Q_N(\theta; \widehat{\mathbf{\Phi}}, \widehat{\mathbf{p}}) - Q_0(\theta; \mathbf{\Phi}^*, \mathbf{p}_0)| \xrightarrow{p} 0$, so $\widehat{\theta}^{\mathrm{GFD}} \xrightarrow{p} \theta_0$ 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 $\Lambda_d(\theta_0) = p_0(d \mid x)$,

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

since $\sum_d \Lambda_d \equiv 1$. As the right-hand side is identically $\mathbf{0}$ in $\mathbf{\Phi}$, $\partial_{\mathbf{\Phi}} \mathbb{E}[s(\theta_0; \mathbf{\Phi}, \mathbf{p}_0)] = \mathbf{0}$. Because f enters the criterion *only* through $\mathbf{\Phi}(\cdot; f)$, the chain rule gives $\partial_f \mathbb{E}[s(\theta_0; \mathbf{\Phi}(\cdot; f), \mathbf{p}_0)] = \partial_{\mathbf{\Phi}} \mathbb{E}[s] \cdot \partial_f \mathbf{\Phi} = \mathbf{0}$ as well. Thus the first-stage estimation error in $\widehat{f}$ (hence in $\widehat{\mathbf{\Phi}}$) is first-order negligible for $\widehat{\theta}^{\mathrm{GFD}}$; only the direct dependence of Λ on $\mathbf{p}$ through the Hotz–Miller correction contributes a first-stage term.

Step 3 (asymptotic linearization). 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 (continuity + Step 1), gives

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

Expanding the score in the nuisance and applying Step 2 (the $\mathbf{\Phi}/f$ derivative vanishes) leaves

$$\sqrt{N}\,\nabla_\theta Q_N(\theta_0;\widehat{\mathbf{\Phi}},\widehat{\mathbf{p}}) = \frac{1}{\sqrt{N}}\sum_n \Big[s(\theta_0;\mathbf{\Phi}^*,\mathbf{p}_0) + \mathbf{D}_p\,\boldsymbol{\psi}^{\mathrm{p}}(z_n)\Big] + o_p(1),$$

where $\mathbf{D}_p = \partial_{\mathbf{p}}\mathbb{E}[s(\theta_0;\mathbf{\Phi}^*,\mathbf{p})]\big|_{\mathbf{p}_0}$ and $\boldsymbol{\psi}^{\mathrm{p}}$ is the CCP component of the first-stage influence function of Assumption 7. The bracket is i.i.d., mean zero, with finite variance $\mathbf{V}(\mathbf{\Phi}^*) = \mathrm{Var}\big(s(\theta_0;\mathbf{\Phi}^*,\mathbf{p}_0) + \mathbf{D}_p\boldsymbol{\psi}^{\mathrm{p}}\big)$.

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 $\mathbf{\Sigma}(\mathbf{\Phi}^*) = \mathbf{J}(\mathbf{\Phi}^*)^{-1}\mathbf{V}(\mathbf{\Phi}^*)\mathbf{J}(\mathbf{\Phi}^*)^{-1}$. This is the stated sandwich form (Newey and McFadden, 1994, Ch. 6), with $\mathbf{V}$ carrying the CCP plug-in correction only. □

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

(i) The trace map $\mathbf{q} \mapsto \mathrm{tr}(\mathbf{\Sigma}(\mathbf{\Phi}(\mathbf{q})))$ is continuous on the compact set $\mathcal{Q}$ (rational with no singularities on $\mathcal{Q}$ by Assumption 8, which keeps $\mathbf{J}$ uniformly non-singular). The extreme value theorem gives existence of a minimizer. (ii) Standard FOC at an interior optimum of a smooth objective. (iii) Restatement of (36): the minimum is attained at $\mathbf{\Phi}^{\mathrm{opt}}$ by part (i), so its trace is the infimum. □

E.8 Proof of Corollary 9 (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 Theorem 7 (single-agent version) to this submodel at (x_t, d_t, d_t'): pointwise FD at horizon $\rho_\ell^\star$ delivers a flow $\mathbf{\Phi}^{(\ell)}(x_t)$ satisfying (39). Since this construction is independent across types, the mixture log-likelihood $\log\sum_\ell \pi_\ell\, p_t^{(\ell)}(d_{n,t} \mid x_{n,t})$ depends on type-conditional CCPs only through the GFD representation, with no recourse to a Bellman fixed point. □

E.9 Proof of Theorem 10 (GFD-EM consistency)

Sketch. The argument has four steps.

Step 1 (uniform consistency of the anchor). Under (iii) the empirical posterior-weighted 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). This is the standard cell-frequency rate under a fixed posterior and is inherited by the EM posterior in the limit (Wu 1983; McLachlan and Krishnan 2008).

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

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

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 maximiser $(\boldsymbol{\theta}_0, \boldsymbol{\pi}_0)$ modulo label switching, since the type-conditional likelihoods are pairwise distinct on a positive-mass set, and the mixture decomposition is identified by the Kasahara and Shimotsu (2009) rank condition adapted to the non-stationary case.

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 of $\widetilde{Q}_N$—not necessarily the global maximiser. Estimator consistency is established separately: the GFD-EM estimator is defined as the arg max of $\widetilde{Q}_N$, which by Step 2 (uniform convergence of a continuous objective on the compact set $\Theta \times \Delta^{L-1}$ to a population objective with a unique maximiser modulo label-switching) is consistent for (θ_0, π_0) by the standard M-estimator argument. In practice EM need not reach the global maximiser without good initialisation, so multi-start is recommended.

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 the score $\partial \widetilde{Q}_N / \partial \boldsymbol{\eta}\big|_{\boldsymbol{\eta}_0}$ has the GFD-induced influence function $\widehat{\psi}_n + I^{-1} \partial \widetilde{Q}_N / \partial \widehat{p} \cdot (\widehat{p} - p_0)$. Standard delta-method arguments yield $\sqrt{N}(\widehat{\boldsymbol{\eta}} - \boldsymbol{\eta}_0) \Rightarrow N(0, V^{\mathrm{GFD}})$ with the sandwich matrix $V^{\mathrm{GFD}} = I^{-1} \Omega I^{-1}$ characterized by the GFD score and Hessian. $\square$