DYNAMIC MARKET DESIGN IN LARGE ECONOMIES

Foundations and Applications to the Design of Dynamic Incentives*

César Barilla† Charlotte Martres‡ Olivier Tercieux§

August 2026

Abstract

Motivated by large centralized career-assignment systems for public servants—including teachers, military officers, police officers, and judges—this paper develops a framework for analyzing and designing dynamic matching markets without transfers, with particular attention to persistent compositional imbalances across institutions, such as the concentration of inexperienced staff in disadvantaged or hard-to-staff locations.

We study an overlapping-generations environment in which a continuum of agents is repeatedly assigned to institutions and current assignments affect both future preferences and future priorities. We first conduct a positive analysis of repeated agent-proposing Deferred Acceptance (DA), a commonly used procedure in these applications. We characterize its equilibrium outcomes through dynamic market-clearing cutoff paths and show that this characterization remains approximately valid in large finite markets. We then study stationary market-clearing cutoffs. Such cutoffs always exist, but familiar static structure need not survive: stationary demand may violate monotonicity and gross substitutes, stationary cutoffs may be robustly multiple, and, under aggregate scarcity, their set is an anti-lattice.

We next turn to normative analysis. Public authorities in our applications often determine both how priority scores are initially assigned and how current assignments affect future scores, frequently using these dynamic incentives to address distributional imbalances. We therefore treat the initial score distribution and the assignment-dependent evolution of scores as design instruments. Our implementation principle shows that any stationary equilibrium allocation generated by a broad class of mechanisms can be reproduced by a truthful stationary equilibrium under repeated DA through an appropriate design of initial scores and score transitions.

Finally, we use this principle to address a central design question: how should dynamic priorities be chosen to pursue distributional objectives while preserving efficiency? We develop a constructive two-step method that first identifies a constrained-efficient allocation through a dynamic pseudo-market and then translates it into an implementable dynamic priority system under DA. A quantitative proof of concept shows that the method is computationally tractable at the scale of major real-world centralized assignment systems.

*We are grateful to Yeon-Koo Che, Qingmin Liu, Alessandro Pavan, Tayfun Sönmez, and Utku Ünver for helpful comments and discussions. We also thank seminar and conference participants at Brown University, Boston College, Columbia University, the plenary session of the 2026 Conference on Mechanism and Institution Design at York University, the Hausdorff Research Institute for Mathematics in Bonn, and the Tinbergen Institute.

†University of Oxford and Nuffield College.

‡Paris School of Economics.

§Paris School of Economics and CNRS.

Contents

1 Introduction **4**

2 The Continuum Model **8**
2.1 Primitives and Matchings 9
2.2 Solution concepts 12
2.2.1 Cutoffs and Dynamic Demand 12
2.2.2 Market-clearing Cutoff Path 14
2.2.3 Stationary Market-clearing Cutoff 14

3 Positive Analysis **15**
3.1 Reported Demand and Deferred Acceptance 15
3.2 The Dynamic Reporting Game 16
3.3 Equilibrium Characterization 17
3.4 Existence and Structural Properties 18
3.5 Finite Economies 23

4 Normative Analysis **26**
4.1 General Mechanisms 26
4.2 The Implementation Principle 27
4.2.1 Result 30
4.2.2 Proof Sketch 31

5 Applications **32**
5.1 Implementation of a Constrained Efficient Allocation via Deferred Acceptance 32
5.2 Quantitative Exploration: The Welfare-Inequality Frontier 35

A Distributional Implications of the Maintained Assumptions **44**

B Proof of Lemma 2.1 **45**

C Proofs for Section 3 **47**
C.1 Stable Cutoff Selectors and the General Reporting Game 47
C.2 Equilibrium Characterization 48
C.3 Selector-independent Characterization 49
C.4 Proof of Theorem 3.1 51
C.5 Finite economies: proof of Theorem 3.2 57

D Proofs for the Implementation Results **59**
D.1 Proof of Theorem 4.1 59
D.2 Nonnegative Score 65

E Additional Results and Examples for Section 4 **67**
E.1 Deferred Acceptance as a Particular Case of our General Mechanisms 67
E.2 The Value of Message Records 68

F Additional results and proofs for Section 5 **70**
F.1 Framework and method to find a constrained efficient allocation 70
F.1.1 Complete Information 70
F.1.2 Incomplete Information 74
F.1.3 Implementation via Deferred Acceptance 77
F.2 Existence proofs 79
F.2.1 Proof of Proposition F.1 79
F.2.2 Proof of Proposition F.2 81
F.3 Constrained efficiency proofs 82
F.3.1 Proof of Theorem F.1 82
F.3.2 Proof of Theorem F.2 83
F.4 Implementation Proofs 88
F.5 Pre-processing of lower bounds on average age 89
F.6 Simulation design and additional simulation results 92
F.6.1 Demand Engine 92
F.6.2 Approximation of low prices 94
F.6.3 Equilibrium solver 94
F.6.4 Certification 95
F.6.5 Additional simulation results and figures 95

1 Introduction

Many centralized matching systems assign the same individuals repeatedly over the course of their careers. A teacher's current school, a police officer's district, or a civil servant's current post has consequences beyond their contemporaneous allocation: today's assignment may change the agent's location, accumulated priority, eligibility, and future set of opportunities. Such systems are often operated through large, annual clearinghouses in which agents submit preferences and positions are allocated according to administratively determined priorities. Prominent examples include recurrent centralized systems for teacher transfers, military officer assignments, police deployments, and civil-service postings (Combe et al., 2022; Sidibe et al., 2021; Thakur, 2021; Elacqua and Rosa, 2025; Davis et al., 2026; Republic of Türkiye Ministry of National Education, 2026a).[1]

These applications share two features that distinguish them from the static environments at the center of the matching literature. First, assignments are linked across time. A current match can alter both the agent's future preferences and the priority with which the agent competes for later positions. For example, a teacher's current placement may induce relocation, housing investment, or schooling choices for their children, thereby changing the value of future locations. The assignment may also generate seniority or points that affect access to later positions. Earlier assignments also shape the cross-sectional composition of older age groups: which agents reach a given career stage, and with which priorities, is endogenous to their previous assignments. The market at date $t+1$ is therefore endogenous to the allocation at date t, so the dynamic problem cannot generally be analyzed as a sequence of independent static markets.

Second, these markets frequently exhibit persistent distributional imbalances. Less experienced or less effective teachers tend to be concentrated in disadvantaged schools, and workers often seek to leave remote or otherwise difficult positions as their careers progress (Hanushek et al., 2004; Jackson, 2009; Combe et al., 2022; Bobba et al., 2021). Monetary compensation may be unavailable, rigid, or costly to adjust, particularly in public-sector systems governed by uniform pay scales and tight administrative budgets. Policymakers often rely instead on designing what we call *dynamic incentives*: accepting an unattractive assignment today may generate additional points, seniority, or priority in a future assignment round. Such rewards are widespread in practice.[2] Their equilibrium effects are subtle. A more generous priority bonus changes not only the incentives of its intended recipients, but also future congestion, accessibility thresholds, and the value of the bonus itself.

This paper develops a framework for analyzing and designing these dynamic matching systems without transfers. We make three contributions. First, we introduce a tractable equilibrium framework for repeated many-to-one matching. Second, we establish an implementation principle showing when a rich dynamic design problem can be reduced to the design of priorities under repeated Deferred Acceptance. Third, we use this principle to characterize constrained-efficient dynamic priority systems and provide a systematic two-step computational method to their design.

Model and solution concepts. We consider an overlapping-generations economy with a finite set

[1] Although our main application is the assignment of public servants, the framework also covers other repeated matching environments, including campus housing (Guillen and Kesten, 2012) and school transitions in which current assignments affect future application sets or priorities (Hahm and Park, 2024).

[2] In French teacher assignment, for example, service in designated disadvantaged schools can generate additional points that improve priority in subsequent transfer rounds; see Combe et al. (2022). Comparable assignment-dependent rewards appear in Spain, where service in a school classified as especially difficult adds points in later transfer contests; in New South Wales, where teachers accumulate school-specific transfer points and may qualify for priority transfers after service in rural or remote schools; and in Turkey, where school conditions determine service points used in later transfer rounds (Government of Spain, 2010; New South Wales Department of Education, 2026b,a; Republic of Türkiye Ministry of National Education, 2026b).

of institutions, which we call schools, and a continuum of agents. At every date, a new cohort enters, remains in the market for a finite number of periods, and is assigned once per period. An agent's state contains a vector of school-specific flow payoffs and a vector of school-specific scores; both components may evolve with age, the current state, and the current assignment. Thus, attending a school can affect both how attractive future schools are and the priority with which the agent competes for them.

We introduce two solution concepts admitting a dynamic supply-and-demand representation. At each date, school selectivity is summarized by a vector of cutoffs. Given an anticipated cutoff path, an agent chooses among the schools whose cutoffs they meet, taking into account the effect of the current assignment on future types and future feasible sets. Current demand therefore depends on the entire anticipated path of cutoffs. Moreover, past assignments affect the current distribution of scores and payoffs, so aggregate demand is endogenous to the history of the market.

We call a cutoff path *market clearing* when demand never exceeds capacity and every school with a positive cutoff fills its capacity. A stationary market-clearing cutoff is a constant cutoff path that generates a stationary cross-sectional distribution of types. These concepts are dynamic counterparts of the supply-and-demand representation of static matching developed by Azevedo and Leshno (2016). The crucial difference is that, in our model, a cutoff vector affects aggregate demand both directly, through current feasibility and continuation values, and indirectly, through the endogenous distribution of agents at future ages.

Positive analysis. We begin by providing a strategic foundation for our solution concepts. Suppose that, at every date, agents submit rank-order lists and current assignments are determined by agent-proposing Deferred Acceptance (DA), as in several of the recurring centralized systems discussed above.[3] In the continuum economy, a single agent takes current and future cutoffs as given. We show that cutoff paths induced by equilibria of the induced game are exactly the market-clearing cutoff paths, and that stationary equilibrium cutoff vectors are exactly the stationary market-clearing cutoffs. At any equilibrium cutoff path, each agent is therefore assigned to a dynamically optimal feasible school, where dynamic optimality incorporates both the current flow payoff and the continuation value induced by the assignment.

We also show that the same cutoff-based representation approximates equilibrium behavior in large finite economies, with an error that vanishes as the market grows. This provides a foundation for using the continuum model in applications with a large but finite number of participants. Further, this characterization has an empirical implication: rankings submitted in a dynamic assignment system reveal preferences over the lifetime consequences of current assignments, not merely over their contemporaneous payoffs. This points toward dynamic extensions of the revealed-preference methods used to estimate demand in centralized assignment systems (Fack et al., 2019; Artemov et al., 2023; Che et al., 2023).

DA also retains a useful truthfulness property, although truthfulness is dynamic and endogenously determined rather than based on a current exogenous flow payoff. A truthful report ranks schools according to their dynamic choice-specific values in equilibrium. We show that every market-clearing cutoff path, and hence every equilibrium allocation outcome, can be supported by such reports. Even though, contrary to the static environment, multiple truthful equilibria may coexist in our setting, restricting attention to truthful equilibria is without loss in terms of outcomes.

We then study the structural properties of stationary market-clearing cutoffs and hence of

[3]Examples include the French teacher-assignment system, São Paulo's annual teacher-transfer market, and the U.S. Army's recurring officer marketplace. Turkey's teacher-transfer system uses a common administrative point ranking across institutions and can therefore be represented as the common-priority special case of DA (Combe et al., 2022; Elacqua and Rosa, 2025; Davis et al., 2026; Republic of Türkiye Ministry of National Education, 2026a).

equilibrium outcomes. In the static continuum model, market-clearing cutoffs form a nonempty complete lattice and are generically unique (Azevedo and Leshno, 2016). The dynamic setting preserves existence but not the remaining structure. A change in cutoffs alters current feasibility, continuation values, and the endogenous stationary distribution of types. These channels can overturn the monotonicity and gross-substitutes properties that discipline static demand. We provide explicit examples in which stationary demand violates these properties and construct economies with robustly multiple stationary cutoffs. Finally, under aggregate scarcity, the set of stationary market-clearing cutoffs satisfies the opposite order property from a lattice: it is an *anti-lattice*. That is, no two distinct stationary market-clearing cutoff vectors are coordinatewise comparable. The coordinatewise cutoff order therefore no longer identifies a canonical agent-optimal or institution-optimal stationary outcome and, thus, there is no longer a clear welfare ranking of equilibrium allocations.

Normative analysis. The positive analysis treats the score process as a primitive. We next reverse the perspective and treat the distribution of initial scores and the evolution of scores as design instruments. This perspective is empirically important: public authorities in many countries — including France, Turkey, Spain, and New South Wales among others — explicitly determine how current service is converted into future assignment priority (Combe et al., 2022; Republic of Türkiye Ministry of National Education, 2026b; Government of Spain, 2010; New South Wales Department of Education, 2026b). Our benchmark in this part is a broad class of anonymous, aggregative, and time-invariant mechanisms satisfying a no-message-records restriction: holding the aggregate state fixed, an agent's own past messages do not directly affect her current assignment.

Our main result is an *implementation principle*. Fix any stationary equilibrium allocation generated by a feasible mechanism in this class. We construct an entry distribution of school-specific scores and score-transition kernels such that repeated agent-proposing DA has a truthful stationary equilibrium inducing an equivalent allocation. Equivalence means that, at every age, the general mechanism and the DA implementation generate the same joint distribution of payoff types and assignments. The appendix also shows that the no-message-records restriction can be dispensed with when score transitions are allowed to condition on the agent's history of past reports of rank-order lists.

Within this class and for stationary allocation outcomes, the implementation principle implies that operating DA is without loss once the planner can design how priorities are initially distributed and how they evolve. Given a designer's objective, one may therefore first solve a comparatively general dynamic allocation problem and then translate the desired allocation into a points system that can be operated through a familiar matching procedure. Motivated by applications in which policymakers seek to address distributional imbalances, we build on this insight below to develop a systematic two-step computational procedure for designing an "optimal" dynamic priority system relative to the policymaker's stated objectives and constraints.

Application to dynamic priority design. We apply the implementation principle to environments in which the planner's objectives are summarized by a finite collection of linear constraints on aggregate assignments. Besides school capacities, these constraints can impose caps on the number of novices assigned to particular positions, floors on experience, or more general bounds on the age and quality composition of institutions, as in the literature on matching with distributional constraints (Kamada and Kojima, 2015; Suzuki et al., 2018; Hafalir et al., 2025).

Within the constrained feasible set, we seek an allocation that is Pareto efficient for agents. We identify such an allocation through a dynamic pseudo-market. Each entrant receives a random lifetime budget, and each aggregate constraint is assigned a shadow price. Because an age-a

assignment to school s can contribute differently to different constraints, these shadow prices induce age- and school-specific prices. Agents choose their preferred affordable career path, while equilibrium prices make the constraints feasible. We establish existence and show that every such pseudo-market equilibrium is constrained efficient.

The implementation principle then translates the pseudo-market allocation into an entry-score distribution and assignment-dependent score transitions under DA, thereby implementing the constrained-efficient allocation. The resulting priority rule is generally richer than the uniform service bonuses commonly observed in practice: the adjustment to an agent's priority at a future school may depend on the agent's current assignment, age, and prospective destination. This contrast suggests that the simpler priority systems used in practice need not be optimally designed.

Hence, methodologically, we offer a constructive two-step method. First, compute a dynamic pseudo-market equilibrium which we show can be obtained by computing a fixed point of a Walrasian-tâtonnement-type map. Second, apply the implementation principle to recover an initial score distribution and score-transition kernels that implement the allocation under DA.

We illustrate the method in a stylized economy with 50 assignment destinations and 40 annual career periods. The goal is to offer a quantitative proof of concept calibrated in scale to French teacher assignment. The 40 periods approximate a full career of potential reassignments and the number of destinations exceeds that of the French national inter-academic teacher-mobility round, which is a particularly appropriate reference market for our question. That round reallocates teachers annually across 30 académies and displays significant imbalances in the distribution of experienced teachers. About half of the tenured teachers who take part in the procedure come from two particularly disadvantaged and unattractive regions in the suburbs of Paris (Combe et al., 2022). The resulting large exit flows are compensated by assigning most new teachers to these two regions, which reproduces the concentration of novices in the same académies over the years. This pattern is a serious source of concern for policy makers as it contributes to undermining the attractiveness of the profession. Our exercise imposes stylized versions of common policy constraints and shows that the resulting priority system can be computed at an empirically relevant scale.

Related literature. This paper contributes to several strands of the literature. Our positive analysis builds most directly on Azevedo and Leshno (2016), who represent stable matchings in a continuum economy with finitely many institutions through market-clearing cutoffs. We extend their supply-and-demand framework to an environment in which agents are repeatedly assigned, demand is forward looking, and past assignments endogenously shape the distribution of future types. The cutoff representation remains tractable, but its structural implications change substantially: stationary demand may violate monotonicity and gross substitutes, stationary market-clearing cutoffs need not form a lattice, and robust multiplicity may arise.

Our normative analysis is related to static welfare-theorem and representation results that provide foundations for restricting attention to tractable market-like mechanisms without transfers. Ashlagi and Shi (2016) and Miralles and Pycia (2014) connect broad classes of efficient assignments to pseudomarkets, while Shi (2022) characterizes the outcomes attainable through priority-based mechanisms and develops a methodology for optimizing priorities and quotas. Our implementation principle provides a dynamic counterpart: any stationary equilibrium allocation generated by a broad class of mechanisms can be translated into initial scores and assignment-dependent score transitions and implemented through repeated DA. This also provides the foundation for a dynamic extension of Shi (2022)'s methodology, in which the evolution of priorities, rather than only their static distribution, becomes a design instrument.

Our application connects to the static literature on assignment under distributional constraints. Suzuki et al. (2018) construct a TTC-based mechanism satisfying efficiency, individual rationality,

and strategy-proofness under discrete-convex constraints; Hafalir et al. (2025) study mechanisms that improve a distributional objective while retaining these properties; and Combe et al. (2022) design and empirically evaluate a mechanism aimed at improving the distribution of teachers across schools. Echenique et al. (2021) construct constrained pseudomarkets in which prices internalize the externalities generated by assignment constraints. We introduce a dynamic dimension absent from these papers: current assignments alter agents' future payoffs, priorities, and opportunities, and therefore the future composition of institutions. Our application turns the corresponding intertemporal shadow values into an implementable system of dynamic priorities.

Our work also relates to three further strands of the literature. First, it connects to the theoretical literature on dynamic matching.[4] While Doval (2022) develops a recursive notion of stability for two-sided markets with dynamic arrivals and irreversible matches, and Nicolò et al. (2024) introduces a dynamic core and an intertemporal TTC mechanism for one-sided markets, we study the noncooperative game generated by repeatedly operating DA in an overlapping-generations economy in which assignments affect agents' future payoffs and priorities. Second, our implementation principle contributes to the dynamic mechanism-design literature surveyed by Bergemann and Välimäki (2019). Unlike much of that literature, we exclude monetary transfers and ask when future priority and assignment opportunities can serve as design instruments and be implemented through DA. Finally, our project relates to recent empirical models of dynamic matching. Ederer (2026) studies repeated two-sided matching without transfers and applies his framework to teachers in Peru, whereas Corblet et al. (2025) develop a transferable-utility model and estimate job-specific human-capital accumulation among Swedish engineers. We instead study a many-to-one market without transfers in which priorities are administratively designed, with an emphasis on their structural and normative properties and, ultimately, on a practical methodology to optimally design priority systems.

Organization. Section 2 presents the baseline continuum model and introduces market-clearing cutoff paths, and their stationary counterparts. Section 3 provides a game-theoretic foundation for these solution concepts under repeated agent-proposing Deferred Acceptance, studies the structural properties of stationary market-clearing cutoffs, and connects the analysis to large finite markets. Section 4 develops the normative analysis and establishes the implementation principle, showing how stationary equilibrium allocations generated by a broad class of mechanisms can be implemented through appropriately designed dynamic priorities under repeated DA. Section 5 applies this principle to allocation problems with distributional constraints, develops a systematic method for designing constrained-efficient dynamic priority systems, and provides a quantitative proof of concept demonstrating that the method is computationally tractable at the scale of major real-world centralized assignment systems. The Appendices contain the proofs, extensions, additional examples, and computational details.

2 The Continuum Model

This section develops a continuum model of dynamic matching without transfers, designed to capture the main features of the applications described in the Introduction: agents are repeatedly assigned to institutions over the course of their careers, and monetary transfers are unavailable. Time is discrete and indexed by $t \in \mathbb{N}_0$. At every date, a new cohort of mass $1/A$ enters the market and remains in the market for A periods. Hence, the market contains one cohort of each age, and

[4] For a recent survey, see Doval (2025).

so the total mass of agents present at any date is one. Let

$$\mathcal{A} := \{1, \ldots, A\}$$

denote the set of ages.

On the other side of the market is a finite set of institutions, which we refer to as schools,

$$\mathcal{S} := \{1, \ldots, S\}.$$

Each school $s \in \mathcal{S}$ has capacity $q^s > 0$. An agent may be assigned to a school or remain unmatched. We represent the latter outcome by the outside option $\emptyset$ and write

$$\mathcal{S}_0 := \mathcal{S} \cup \{\emptyset\}.$$

2.1 Primitives and Matchings

This subsection introduces the primitive state variables and the assignment objects used throughout the paper. We first describe agents' school-specific priorities and flow payoffs, and then specify how these characteristics evolve with age and past assignments. We next define pre-matchings, which allow us to trace the resulting cross-sectional distribution of types, and matchings, which additionally satisfy the schools' capacity constraints.

Types and entry. For a compact metric space X, let $\Delta(X)$ denote the set of Borel probability measures on X. At every age, an agent is described by a type $\theta = (r, u)$. The vector

$$r = (r^s)_{s \in \mathcal{S}} \in \mathcal{R} := [-1, 1]^S$$

records the agent's scores at the schools. School s ranks agents according to r^s, with a higher score corresponding to higher priority at that school. When we introduce the solution concept below, cutoffs will remain restricted to $[0, 1]$, so that a negative score is used to capture that the agent is ineligible for the school, even when its endogenous cutoff is zero. The vector

$$u = (u^s)_{s \in \mathcal{S}} \in \mathcal{U} := [0, 1]^S$$

records the agent's one-period payoff from attending each school. We normalize the payoff from the outside option to $u^\emptyset = 0$. Thus, the type space is

$$\Theta := \mathcal{R} \times \mathcal{U} = [-1, 1]^S \times [0, 1]^S,$$

endowed with the product topology and its Borel σ-algebra. The type θ is the agent's payoff-relevant state: together with the current assignment, it determines both current utility and, as we will see, the distribution of the agent's future type.

At every date, the types of entering agents have common law $p \in \Delta(\Theta)$. Because an entering cohort is a continuum, we identify p with its realized cross-sectional type distribution and therefore abstract from aggregate uncertainty.[5]

[5]This deterministic formulation can be given a stochastic microfoundation using an exact law of large numbers. Formally, one may place the continuum of agents and the underlying randomization on a rich Fubini extension and take individual types to be essentially pairwise independent with common law p. The realized cross-sectional distribution then equals p almost surely. See Sun (2006) for an exact law of large numbers based on Fubini extensions.

Individual type transitions. The dynamic link in the model operates through the evolution of types. For every age $a < A$ and current assignment $s \in \mathcal{S}_0$, let

$$Q_a(\,\cdot \mid \theta, s) \in \Delta(\Theta)$$

be the Markov kernel governing the agent's type at age $a+1$ when their current type is θ and their age-a assignment is s. The kernels are time invariant, but they may depend on age. The Markov formulation means that the current type and assignment summarize the aspects of the agent's history that are relevant for future payoffs and priorities.

This formulation allows assignments to affect both components of the agent's state. For example, in the teacher- and civil-service assignment systems discussed in the Introduction, service in a disadvantaged or remote post may generate bonus points, seniority, or eligibility in future assignment rounds, while the location and experience associated with that post may alter the agent's subsequent valuation of alternative assignments. The framework also allows these priority advantages to depreciate over time. The model also nests special cases in which scores, payoffs, or both remain fixed throughout an agent's lifetime. We assume that, for every $a < A$ and $s \in \mathcal{S}_0$, the map

$$\theta \longmapsto Q_a(\,\cdot \mid \theta, s)$$

is continuous when $\Delta(\Theta)$ is endowed with the topology of weak convergence.[6] This regularity condition ensures that small changes in the current type produce small changes in the distribution of the next-period type.

For the normative analysis, we will sometimes impose an additional separation between the evolution of priorities and that of payoffs. Specifically, we assume in those parts that

$$\begin{aligned} p(dr, du) &= \widehat{p}(dr)\, p^\star(du), \\ Q_a(dr', du' \mid r, u, s) &= \widehat{Q}_a(dr' \mid r, s)\, Q_a^\star(du' \mid u, s), \qquad a < A, \end{aligned}$$

where $\widehat{p} \in \Delta(\mathcal{R})$, $\widehat{Q}_a(\,\cdot \mid r, s) \in \Delta(\mathcal{R})$, $p^\star \in \Delta(\mathcal{U})$, and $Q_a^\star(\,\cdot \mid u, s) \in \Delta(\mathcal{U})$. This product structure is not needed for the positive results in this section and is imposed only when explicitly stated. It becomes useful in Section 4, where the distribution of agents' priorities and its evolution are design instruments of the planner.

Economies and initial conditions. Collecting the baseline primitives, an economy is a tuple

$$\mathcal{E} = \big[\mathcal{S}, q, p, (Q_a)_{a=1}^{A-1}\big], \qquad q := (q^s)_{s \in \mathcal{S}}.$$

The lifespan A, the type space Θ, and the normalization of the outside option are fixed throughout and are therefore suppressed from this notation. The economy is supplemented by an initial condition

$$\eta_0 := (\eta_{0,a})_{a=2}^{A},$$

which specifies the type distributions of the cohorts already present at date zero. Throughout the analysis, we fix an initial condition and suppress its dependence in the notation.

We will make use of the following assumption.

[6] Equivalently, for every bounded continuous function $f : \Theta \to \mathbb{R}$, the map

$$\theta \longmapsto \int_\Theta f(\theta') Q_a(d\theta' \mid \theta, s)$$

is continuous.

Assumption 2.1 (Absolute continuity and full support)**.** *The entry distribution p and, for every $a < A$, $\theta \in \Theta$, and $s \in \mathcal{S}_0$, the transition law $Q_a(\cdot \mid \theta, s)$ are absolutely continuous with respect to Lebesgue measure and have full support on Θ. The same conditions apply to the exogenously specified initial distributions.*

Assumption 2.1 is deliberately stronger than necessary. It excludes distributions with restricted support and singular transition laws, such as fixed scores or deterministic assignment-specific priority bonuses. We impose it in the main text only to streamline the exposition by collecting the relevant regularity and richness requirements in a single transparent condition. Lemma A.1 in the Appendix derives the weaker reduced-form properties actually used in the proofs and clarifies that the results continue to hold whenever those properties are imposed directly, including in economically relevant cases with singular transition laws.[7]

Pre-matchings and induced type distributions. The transition kernels describe how an individual agent's state evolves once an assignment is specified. To obtain the corresponding aggregate dynamics, we need to specify an assignment rule for every date, age, and current type. A *pre-matching* is a sequence

$$\mu = (\mu_{t,a})_{t \in \mathbb{N}_0,\, a \in \mathcal{A}}$$

of measurable maps

$$\mu_{t,a} : \Theta \longrightarrow \mathcal{S}_0.$$

Thus, $\mu_{t,a}(\theta)$ is the assignment of an age-a agent of type θ at date t. The qualifier "pre" emphasizes that these assignment rules need not satisfy the schools' capacity constraints.

Given a pre-matching and an initial condition, let $\eta^{\mu}_{t,a} \in \Delta(\Theta)$ denote the type distribution conditional on being age a at date t. New entrants satisfy

$$\eta^{\mu}_{t,1} = p \qquad \text{for every } t.$$

For every measurable set $X \subseteq \Theta$, the distributions of older agents evolve according to

$$\eta^{\mu}_{t+1,a+1}(X) = \int_{\Theta} Q_a\big(X \mid \theta, \mu_{t,a}(\theta)\big)\, \eta^{\mu}_{t,a}(d\theta), \qquad a = 1, \ldots, A-1. \tag{1}$$

Equation (1) first assigns each current type according to $\mu_{t,a}$ and then averages the corresponding transition laws. It makes explicit the dynamic link at the center of the model: current assignments change the composition of future cohorts, which in turn shapes future demand and assignments.

Matchings. A pre-matching describes a complete assignment policy but may require more seats than the schools can supply. We call a pre-matching μ a *matching* if it is feasible at every date. That is, for every $t \in \mathbb{N}_0$ and school $s \in \mathcal{S}$,

$$\sum_{a=1}^{A} \frac{1}{A}\, \eta^{\mu}_{t,a}\big(\{\theta : \mu_{t,a}(\theta) = s\}\big) \leq q^s.$$

[7]The full support assumption is used mainly to simplify the statements and our appendix treats the general case where it is not imposed. It is mainly used to obtain uniqueness of static market-clearing cutoffs which we also note can be obtained under weaker conditions. For instance, when we allow for uni-dimensional applicant scores, as in Serial Dictatorship.

The factor $1/A$ converts the type distribution conditional on age into the corresponding cross-sectional mass. The outside option has no capacity constraint.[8]

This definition separates aggregate feasibility from the question of which assignment policies agents will select. The next subsection introduces cutoffs and dynamic demand and uses them to formulate the relevant market-clearing solution concepts.

2.2 Solution concepts

We next formulate the solution concepts through a dynamic supply-and-demand representation that underpins our analysis. At each date, a school's selectivity is summarized by a cutoff. Agents choose among the schools for which they meet the current cutoff, taking into account that today's assignment affects both their future type and their future assignment opportunities. Consequently, current demand generally depends on the entire anticipated path of future cutoffs.

The present section introduces these solution concepts at the level of the continuum economy. Section 3 subsequently provides their game-theoretic foundation by showing how they arise from equilibrium behavior when repeated stable mechanisms are used,[9] first in the continuum model and then, approximately, in large finite economies.

2.2.1 Cutoffs and Dynamic Demand

We now define the demand generated by an anticipated path of school cutoffs. We first characterize agents' feasible sets and forward-looking individual choices, and then aggregate these choices across ages while accounting for the endogenous evolution of the type distribution.

Cutoff paths and feasible schools. A cutoff path is a sequence $\mathbf{P} = (P_t)_{t\in\mathbb{N}_0}$, where $P_t = (P_t^s)_{s\in\mathcal{S}} \in [0,1]^S$. The number P_t^s is the minimum score required for admission to school s at date t. Given a type $\theta = (r,u)$, define the set of feasible assignments at time t by[10]

$$\mathcal{F}_t(r;\mathbf{P}) := \{\emptyset\} \cup \left\{s \in \mathcal{S} : r^s \geq P_t^s\right\}.$$

The cutoff path is taken to be publicly observed and correctly anticipated. Because there is no aggregate uncertainty, it is deterministic.

Individual demand. For each date t, age $a \in \mathcal{A}$, and assignment $s \in \mathcal{S}_0$, define the choice-specific value of type $\theta = (r,u)$ by

$$W_{t,a}^s(\theta \mid \mathbf{P}) := \begin{cases} u^s + \displaystyle\int_\Theta V_{t+1,a+1}(\theta' \mid \mathbf{P}) Q_a(d\theta' \mid \theta, s), & a < A, \\ u^s, & a = A. \end{cases}$$

The value of type $\theta = (r,u)$ at date t is

$$V_{t,a}(\theta \mid \mathbf{P}) = \max_{s\in\mathcal{F}_t(r;\mathbf{P})} W_{t,a}^s(\theta \mid \mathbf{P}). \tag{2}$$

[8] We do not impose the analogue of the "open on the right" convention of Azevedo and Leshno (2016). That convention is not economically substantive; it merely eliminates artificial multiplicity among matchings that differ only on measure-zero sets.

[9] Stability is understood here in the static sense, in a given period and fixing (reported) preferences.

[10] Note that the dynamic program determining demand does not explicitly feature a discount factor. This is without loss of generality. An agent with payoff transition kernel $Q_a^\star(\cdot \mid u,s)$ who discounts the future at rate δ behaves like an agent who does not discount the future but has payoff transition kernel $\delta Q_a^\star(\cdot \mid u,s)$. Since discounting can be absorbed in the transitions, we omit it throughout.

Dynamic acceptability. To ensure that agents do not remain unmatched when a school is feasible, we maintain the following reduced-form restriction. For every cutoff path $\mathbf{P}$, date t, age $a < A$, type θ, and school $s \in \mathcal{S}$,

$$\int_{\Theta} V_{t+1,a+1}(\theta' \mid \mathbf{P}) Q_a(d\theta' \mid \theta, s) \geq \int_{\Theta} V_{t+1,a+1}(\theta' \mid \mathbf{P}) Q_a(d\theta' \mid \theta, \emptyset).$$

Thus, assignment to any school yields a continuation value weakly greater than that generated by the outside option. Since $u^s \geq u^{\emptyset} = 0$, every feasible school is weakly preferred to remaining unmatched. The assumption below together with Assumption 2.1 then implies that an agent with at least one feasible school chooses a school almost surely.

Assumption 2.2 (Generic no-indifference condition)**.** *For every cutoff path $\mathbf{P}$, date t, and age $a \in \mathcal{A}$, the set of types θ for which* $\arg\max_{s \in \mathcal{F}_t(r;\mathbf{P})} W^s_{t,a}(\theta \mid \mathbf{P})$ *has strictly more than one element has zero Lebesgue-measure.*

The condition rules out positive-Lebesgue-measure sets on which current payoff differences and continuation-value differences exactly offset one another. The terminology reflects the idea that such exact cancellations are knife-edge. Indeed, one simple environment in which Assumption 2.2 is satisfied is one in which agents experience small, school-specific transitory payoff (absolutely continuous) shocks, superimposed on an otherwise unrestricted payoff process. If these shocks are redrawn at entry and after every transition, and affect only current flow payoffs, then they eliminate dynamic ties even if they have arbitrarily small support .

We fix a measurable selection from the corresponding argmax of $V_{t,a}(\theta \mid \mathbf{P})$ and denote it by $d_{t,a}(\theta \mid \mathbf{P})$.[11] Thus, $d_{t,a}(\theta \mid \mathbf{P})$ is the assignment demanded by an age-a agent of type θ at date t, given the cutoff path $\mathbf{P}$. The collection

$$\mathbf{d}(\mathbf{P}) := \big(d_{t,a}(\cdot \mid \mathbf{P})\big)_{t,a}$$

is a pre-matching and therefore induces age-specific type distributions

$$\eta^{\mathbf{P}}_{t,a} := \eta^{\mathbf{d}(\mathbf{P})}_{t,a}$$

through (1). Under Assumptions 2.1 and 2.2, individual demand is unique $\eta^{\mathbf{P}}_{t,a}$-almost everywhere (see Lemma A.1 in the appendix). We maintain these assumptions throughout the positive analysis.

Aggregate demand. The demand at date t from age-a agents for school s is

$$D^s_{t,a}(\mathbf{P}) := \frac{1}{A}\, \eta^{\mathbf{P}}_{t,a}\big(\{\theta \in \Theta : d_{t,a}(\theta \mid \mathbf{P}) = s\}\big). \tag{3}$$

Aggregate demand for school s is obtained by aggregating across all ages[12]:

$$D^s_t(\mathbf{P}) := \sum_{a=1}^{A} D^s_{t,a}(\mathbf{P}). \tag{4}$$

Aggregate demand is dynamic in two distinct senses. First, an agent's current choice depends on continuation values and hence on future cutoffs. Second, past assignments affect the current cross-sectional distribution of types through the transition kernels.

[11] A measurable selection can be obtained by fixing an arbitrary ordering of $\mathcal{S}_0$ and selecting the first maximizer under that ordering.

[12] By Lemma A.1, aggregate demand and the induced type distributions do not depend on the tie-breaking rule used to define the measurable selection.

2.2.2 Market-clearing Cutoff Path

Let us now introduce our first solution concept which is the dynamic analogue of market-clearing cutoffs introduced by Azevedo and Leshno (2016) for the static economy.

Definition 2.1 (Market-clearing cutoff path)**.** *A cutoff path* $\mathbf{P}$ *is market clearing if, for every date* t *and school* s*,*

$$D_t^s(\mathbf{P}) \leq q^s, \tag{5}$$

$$D_t^s(\mathbf{P}) = q^s \qquad \text{whenever } P_t^s > 0. \tag{6}$$

The market-clearing conditions have a simple and standard interpretation: Demand cannot exceed capacity. A school with a strictly positive cutoff must fill its capacity, while a school with excess capacity has a zero cutoff. Importantly, and as we already mentioned, the difference is that the demand entering these equations is forward looking and depends on the endogenous distribution of types. This solution concept will be central to characterizing the equilibrium outcomes in a game induced by the use of a stable mechanism, such as DA, at each point in time. This will be developed in Section 3. In the next section, we consider a stationary version of our solution concept. This gives rise to a tractable solution concept, and also enables comparison with the static version of this solution concept as introduced by Azevedo and Leshno (2016).

2.2.3 Stationary Market-clearing Cutoff

We now consider a stationary environment where cutoff paths are constant over time, taking value $P = (P^s)_{s\in\mathcal{S}}$ which we call a stationary cutoff vector. In this context let the time-independent feasible set be

$$\mathcal{F}(r; P) := \{\emptyset\} \cup \{s \in \mathcal{S} : r^s \geq P^s\}.$$

Because the environment is time homogeneous and agents have finite lives, the value functions induced by a constant cutoff vector depend on age but not on calendar time. They satisfy

$$V_A(\theta \mid P) = \max_{s\in\mathcal{F}(r;P)} u^s, \tag{7}$$

and, for every $a < A$,

$$V_a(\theta \mid P) = \max_{s\in\mathcal{F}(r;P)} \left\{ u^s + \int_\Theta V_{a+1}(\theta' \mid P) Q_a(d\theta' \mid \theta, s) \right\}. \tag{8}$$

Let $d_a(\theta \mid P)$ denote the associated demand. The stationary age-specific distributions are generated recursively from the entry distribution:

$$\eta_1^P = p, \tag{9}$$

$$\eta_{a+1}^P(X) = \int_\Theta Q_a\big(X \mid \theta, d_a(\theta \mid P)\big)\, \eta_a^P(d\theta), \qquad a = 1, \ldots, A-1. \tag{10}$$

The resulting age-specific and **stationary aggregate demands** are

$$D_a^s(P) := \frac{1}{A}\, \eta_a^P\big(\{\theta : d_a(\theta \mid P) = s\}\big), \tag{11}$$

$$D^s(P) := \sum_{a=1}^{A} D_a^s(P). \tag{12}$$

Definition 2.2 (Stationary market-clearing cutoffs)**.** *A vector $P \in [0,1]^S$ is a stationary market-clearing cutoff vector if, for every school s,*

$$D^s(P) \leq q^s, \qquad D^s(P) = q^s \quad \text{if } P^s > 0.$$

By Lemma A.1, a stationary cutoff vector P determines the demand policies up to sets of measure zero and, through (9)–(10), therefore uniquely determines the age-specific type distributions $(\eta_a^P)_{a \in \mathcal{A}}$.

The following result provides a useful consistency check: any convergent market-clearing cutoff path must converge to a stationary market-clearing cutoff vector.

Lemma 2.1 (Convergence to stationarity)**.** *Suppose Assumptions 2.1 and 2.2 hold. Suppose that $\mathbf{P} = (P_t)_{t \in \mathbb{N}_0}$ is a market-clearing cutoff path and that $P_t \to P$. Then P is a stationary market-clearing cutoff vector.*

The proof is given in Appendix B.

3 Positive Analysis

This section provides a game-theoretic foundation for the market-clearing cutoff paths introduced in Section 2. At every date, agents submit strict rank-order lists and the agent-proposing Deferred Acceptance mechanism selects the current matching. We show that the cutoff paths generated in equilibrium coincide with the market-clearing cutoff paths of Definition 2.1; the stationary version of the result provides an analogous foundation for the stationary market-clearing cutoffs of Definition 2.2.

The argument proceeds in three steps. Subsection 3.1 first describes the static matching problem generated at a fixed date by the current distribution of types and the agents' reports. Subsection 3.2 then links these one-period problems over time and defines equilibrium under repeated DA. Finally, Subsection 3.3 establishes the equivalence between equilibrium cutoff outcomes and market-clearing cutoff paths. For ease of exposition, the main text focuses on DA. Appendix C develops the more general formulation with arbitrary stable cutoff selectors and explicitly allows the static market to have multiple market-clearing cutoff vectors.

3.1 Reported Demand and Deferred Acceptance

We begin with the matching problem at a fixed date, holding the current cross-sectional distribution and the agents' reports fixed. By the cutoff characterization of Azevedo and Leshno (2016), the outcome of any stable mechanism in this static market can be represented by selecting one of the market-clearing cutoff vectors associated with "reported demand". As will become clear, we therefore view stable mechanisms as cutoff selectors, with agent-proposing DA corresponding to the selector choosing the (coordinatewise) smallest market-clearing cutoff vector.

Let $\mathcal{L}$ denote the finite set of strict rank-order lists over $\mathcal{S}_0$. For a cutoff vector $P \in [0,1]^S$, a score vector $r \in \mathcal{R}$, and a report $\ell \in \mathcal{L}$, let

$$\widehat{d}(P \mid r, \ell)$$

be the $\succ_\ell$-maximal element of

$$\mathcal{F}(r; P) = \{\emptyset\} \cup \{s \in \mathcal{S} : r^s \geq P^s\}.$$

Thus, $\widehat{d}(P \mid r, \ell)$ is the agent's most preferred assignment, according to the reported ranking, among those for which the agent meets the cutoff.

A symmetric reporting strategy is a collection $\sigma = (\sigma_{t,a})_{t,a}$ of measurable probability kernels

$$\sigma_{t,a}(\cdot \mid \theta) \in \Delta(\mathcal{L}),$$

where $\sigma_{t,a}(\cdot \mid \theta)$ is the distribution of reports used by an age-a agent of type θ at date t. Given a profile $\sigma_t = (\sigma_{t,a})_{a \in \mathcal{A}}$ and age-specific type distributions $\eta_t = (\eta_{t,a})_{a \in \mathcal{A}}$, the associated reported aggregate demand for school s at cutoff vector P is

$$\widehat{D}^s(P \mid \sigma_t, \eta_t) := \frac{1}{A} \sum_{a=1}^{A} \int_{\Theta} \int_{\mathcal{L}} \mathbf{1}\Big\{\widehat{d}(P \mid r, \ell) = s\Big\}\, \sigma_{t,a}(d\ell \mid \theta)\, \eta_{t,a}(d\theta). \tag{13}$$

For any reported-demand function $\widehat{D} = (\widehat{D}^s)_{s \in \mathcal{S}}$, let

$$\mathcal{C}(\widehat{D}) := \Big\{P \in [0,1]^S : \widehat{D}^s(P) \le q^s \text{ for every } s, \text{ and } \widehat{D}^s(P) = q^s \text{ whenever } P^s > 0\Big\}$$

be the set of static market-clearing cutoff vectors.

Once σ_t and η_t are fixed, this is a static continuum matching market. Lemma A.1 implies that its score distribution places no mass on cutoff boundaries. The cutoff characterization of Azevedo and Leshno (2016) therefore applies: $\mathcal{C}(\widehat{D})$ is nonempty and forms a complete lattice under the coordinatewise order. We denote its smallest element by

$$\mathrm{DA}(\widehat{D}) := \min \mathcal{C}(\widehat{D}).$$

This is the cutoff representation of the agent-optimal stable matching selected by agent-proposing Deferred Acceptance. In particular, DA remains well defined even when the static market admits several market-clearing cutoff vectors. Appendix C represents other stable mechanisms by alternative selections from $\mathcal{C}(\widehat{D})$ and proves the characterization results at that level of generality.

3.2 The Dynamic Reporting Game

We now link the one-period matching problems across dates. Fix the initial distributions $(\eta_{0,a})_{a=2}^{A}$ and suppose that DA is operated at every date. A strategy profile σ induces an outcome recursively. New entrants satisfy $\eta^{\sigma}_{t,1} = p$ for every t, while the distributions of the cohorts already present at date zero are given by the initial condition. At each date t, define

$$\widehat{D}^{\sigma}_t(\cdot) := \widehat{D}(\cdot \mid \sigma_t, \eta^{\sigma}_t), \tag{14}$$

$$P^{\sigma}_t := \mathrm{DA}\big(\widehat{D}^{\sigma}_t\big). \tag{15}$$

The age-specific type distributions evolve according to

$$\eta^{\sigma}_{t+1,a+1}(X) = \int_{\Theta} \int_{\mathcal{L}} Q_a\Big(X \mid \theta, \widehat{d}(P^{\sigma}_t \mid r, \ell)\Big)\, \sigma_{t,a}(d\ell \mid \theta)\, \eta^{\sigma}_{t,a}(d\theta), \qquad a = 1, \dots, A-1. \tag{16}$$

Thus, current reports determine the current DA cutoff and assignments, while those assignments affect the distribution of types—and hence reported demand and cutoffs—at subsequent dates. Although individual agents may randomize over reports, exact aggregation in the continuum implies that the realized distribution of reports coincides with its expected distribution, hence, mixed

reports create no aggregate uncertainty. The induced cutoff path $\mathbf{P}^\sigma = (P_t^\sigma)_{t\in\mathbb{N}_0}$ and the associated cross-sectional distributions are therefore deterministic.

The equilibrium notion is particularly simple in the continuum. A single agent cannot affect aggregate reported demand and hence cannot alter either the current cutoff or the future cutoff path. At the same time, conditional on these cutoffs, an agent can obtain any currently feasible assignment by placing it at the top of their report. Sequential optimality can therefore be expressed directly through the dynamic choice problem of Section 2.

Definition 3.1 (Equilibrium under Deferred Acceptance)**.** *A reporting strategy profile σ is an equilibrium of the game induced by DA if, for every date t, age a, and type $\theta = (r, u)$,*

$$\widehat{d}(P_t^\sigma \mid r, \ell) \in \underset{s\in\mathcal{F}_t(r;\mathbf{P}^\sigma)}{\arg\max} W_{t,a}^s(\theta \mid \mathbf{P}^\sigma) \quad \textit{for } \sigma_{t,a}(\cdot \mid \theta)\textit{-almost every } \ell. \tag{17}$$

Condition (17) requires every report used by an agent to induce a dynamically optimal feasible assignment. It is equivalent to the absence of a profitable unilateral deviation: a deviation does not change the cutoff path, while the choice-specific values already incorporate the effect of the current assignment on future types and future assignment opportunities. Because agents have finite remaining lifetimes, backward induction—or, equivalently, the one-shot deviation principle—then rules out profitable multi-period deviations. Notice that the condition disciplines the assignment induced by a report, rather than the report itself. Several reports may therefore be optimal when they induce the same assignment.

We call an equilibrium σ with induced cutoff path $\mathbf{P}^\sigma$ *truthful* if, for every date t, age $a \in \mathcal{A}$, and $\eta_{t,a}^\sigma$-almost every type θ, the strategy $\sigma_{t,a}(\cdot \mid \theta)$ is concentrated on the strict rank-order list obtained by ordering all assignments $s \in \mathcal{S}_0$ according to their dynamic choice-specific values $W_{t,a}^s(\theta \mid \mathbf{P}^\sigma)$ (with ties resolved by a fixed ordering). Thus, truthfulness is dynamic: agents rank assignments according to their total values, which incorporate both current flow payoffs and continuation values, rather than according to current payoffs alone.

For the stationary analysis, we restrict both behavior and the induced aggregate state to be time invariant.

Definition 3.2 (Stationary equilibrium under Deferred Acceptance)**.** *A time-invariant reporting strategy $\sigma = (\sigma_a)_{a\in\mathcal{A}}$ is a stationary equilibrium of the game induced by DA if there exist initial age-specific type distributions $(\eta_{0,a})_{a=2}^A$ such that:*

1. *when σ_a is used at every date, the induced cutoff path and age-specific type distributions are time invariant; and*
2. *the resulting strategy profile is an equilibrium under these initial distributions in the sense of Definition 3.1.*

3.3 Equilibrium Characterization

We can now connect the strategic formulation to the supply-and-demand solution concepts of Section 2. The first part of the next result concerns arbitrary initial conditions and nonstationary cutoff paths; the second specializes the equivalence to stationary outcomes.

Proposition 3.1 (Deferred Acceptance and market-clearing cutoffs)**.** *Suppose Assumptions 2.1 and 2.2 hold.*

1. *For a fixed initial condition, DA-equilibrium cutoff paths are precisely the market-clearing cutoff paths. Moreover, every market-clearing cutoff path can be induced by a truthful equilibrium.*

2. *Stationary DA-equilibrium cutoff vectors are precisely the stationary market-clearing cutoff vectors. Moreover, every stationary market-clearing cutoff vector can be induced by a truthful stationary equilibrium.*

The equivalence has a simple interpretation. In any equilibrium, the assignment generated by an agent's report coincides almost everywhere with the agent's true dynamic demand at the equilibrium cutoffs. Reported and dynamic aggregate demand therefore agree at the selected cutoff, so DA's one-period market-clearing conditions imply dynamic market clearing. Conversely, given a market-clearing cutoff path, agents can report the rankings induced by their dynamic choice-specific values (and so they report truthfully). At the target cutoff, the resulting reported demand coincides with dynamic demand; the full-support conclusion of Lemma A.1 makes this cutoff the unique static market-clearing cutoff for the constructed reports, so it is selected by DA.

The formal proof is given in Appendix C, and separates the two ingredients of this argument. It first proves the equilibrium–market-clearing equivalence for arbitrary stable cutoff selectors, allowing the current static market to have several market-clearing cutoffs. It then shows that, under the maintained richness conditions, every stable selector implements every market-clearing path. Proposition 3.1 follows as a corollary of those more general results.

3.4 Existence and Structural Properties

The equilibrium characterization in Subsection 3.3 identifies stationary market-clearing cutoffs with the stationary cutoff outcomes of equilibria under repeated DA. We now study this equilibrium set, asking first whether it is nonempty under our maintained assumptions and, more importantly, whether it retains the structural properties familiar from static continuum matching. To isolate the effects of intertemporal linkages from transitional dynamics, we focus throughout on stationary cutoffs.

Hence, the analysis addresses three questions. We first establish that a stationary market-clearing cutoff vector always exists. We then study the order structure of the solution set and identify a weak condition under which distinct stationary cutoffs cannot be coordinatewise comparable. Finally, we explain why the demand properties supporting the static lattice structure need not survive in the dynamic environment.

Total assignment and composition effects. For every stationary cutoff vector P, let ρ_a^P denote the score marginal of η_a^P and define the stationary cross-sectional score distribution by

$$\bar{\rho}^P := \frac{1}{A}\sum_{a=1}^{A} \rho_a^P.$$

Let

$$E(P) := \big\{r \in \mathcal{R} : r^s \geq P^s \text{ for some } s \in \mathcal{S}\big\}$$

be the set of score vectors that meet at least one cutoff. Dynamic acceptability and our assumptions (see Lemma A.1 in the appendix) imply that, except on a null set, an agent is assigned to a school if and only if their score vector belongs to $E(P)$. Hence the total mass assigned under P is

$$M(P) := \sum_{s\in\mathcal{S}} D^s(P) = \bar{\rho}^P\big(E(P)\big). \tag{18}$$

Consider two distinct cutoff vectors $P^- \leq P^+$. Since $E(P^+) \subseteq E(P^-)$, total assignment satisfies

$$M(P^-) - M(P^+) = \bar{\rho}^{P^-}\big(E(P^-) \setminus E(P^+)\big) - \Big[\bar{\rho}^{P^+}\big(E(P^+)\big) - \bar{\rho}^{P^-}\big(E(P^+)\big)\Big]. \tag{19}$$

The first term is the *direct eligibility effect*: lowering cutoffs expands the set of agents who meet at least one cutoff. The term in brackets is the *composition effect*: changing cutoffs also changes past assignments and therefore the stationary distribution of scores. Our order result only requires that these two effects not cancel exactly.

Assumption 3.1 (No exact compensation)**.** *For every pair of distinct cutoff vectors $P^- \leq P^+$ with $P^+ \in (0,1]^S$,*

$$M(P^-) \neq M(P^+).$$

Equivalently, the two terms on the right-hand side of (19) *are not equal.*

The condition is a natural non-compensation requirement which rules out the knife-edge case in which the endogenous composition effect exactly offsets the direct eligibility effect. It also follows from transparent primitive restrictions. For example, suppose that next-period scores depend on the current state only through current scores and are independent of the current assignment. The stationary score distribution is then independent of P. With full support of entry scores, lowering a positive cutoff vector strictly increases the mass of agents who meet at least one cutoff, and therefore implies $M(P^-) > M(P^+)$.

Order structure. We say that a set $X \subseteq [0,1]^S$ is an **anti-lattice** under the coordinatewise order if no two distinct elements are comparable: for every $P, P' \in X$,

$$P \leq P' \quad \Longrightarrow \quad P = P'.$$

We obtain the following result.

Theorem 3.1 (Existence, order structure, and robust multiplicity)**.** *Suppose Assumptions 2.1 and 2.2 hold. Consider an economy $\mathcal{E} = [\mathcal{S}, q, p, (Q_a)_{a=1}^{A-1}]$.*

(i) There exists a stationary market-clearing cutoff vector.

(ii) If $\mathcal{R} = [0,1]^S$, $\sum_{s\in\mathcal{S}} q^s < 1$, and Assumption 3.1 holds, then the stationary market-clearing cutoffs form an anti-lattice under the coordinatewise order.[13]

(iii) There exists an economy $\mathcal{E} = [\mathcal{S}, q, p, (Q_a)_{a=1}^{A-1}]$ with multiple stationary market-clearing cutoff vectors. Moreover, this multiplicity is robust: every sufficiently small perturbation[14] *of the primitives $(q, p, (Q_a)_{a=1}^{A-1})$ also admits multiple stationary market-clearing cutoff vectors.*

[13] A set that is both a lattice and an anti-lattice is necessarily a singleton. This observation provides one way to interpret the uniqueness result of Azevedo and Leshno: generically, the static set of market-clearing cutoffs has both properties and hence contains a single element.

[14] Appendix C.4 makes this notion precise. Perturbations of capacities are measured in the Euclidean norm, those of the entry distribution in total variation, and those of the transition kernels in total variation uniformly over ages, current types, and assignments.

Part (i) is the dynamic counterpart of a standard competitive-equilibrium existence argument. Once continuity of stationary demand is established (which follows from the weak continuity of the transition kernels and the null-boundary properties established in Lemma A.1), a standard Walrasian-tâtonnement-type map and Brouwer's fixed-point theorem deliver a stationary market-clearing cutoff.

The role of the assumptions in part (ii) is transparent. Under nonnegative scores, a school with a zero cutoff is feasible for every agent. When total capacity is strictly below one, dynamic acceptability therefore implies that every stationary market-clearing cutoff has strictly positive components. Every school then fills and total assignment equals $\sum_s q^s$. The no-exact-compensation condition rules out two distinct comparable cutoffs, since they would have to generate both equal and unequal total assignment.

This conclusion contrasts sharply with the static model. In a static continuum matching economy, market-clearing cutoffs form a nonempty complete lattice and therefore have extremal elements, including an agent-optimal cutoff. Under part (ii), by contrast, multiple stationary cutoffs are necessarily incomparable, so the coordinatewise cutoff order provides no analogue of the static extremal-cutoff ranking. Part (iii) shows that multiplicity can occur even under Assumptions 2.1 and 2.2 and it can persist under perturbations which is in contrast with the generic uniqueness result obtained in Azevedo and Leshno (2016). We provide an example below which illustrates this phenomenon.

Example on multiplicity. Let $\mathcal{S} = \{1, 2\}$ and $A = 2$, so each age group has mass $1/2$. Age-one scores are independent and uniformly distributed:

$$r_1^1, r_1^2 \overset{\text{i.i.d.}}{\sim} \text{Unif}[0, 1].$$

The flow-payoff vectors when young and old are, respectively,

$$u^Y = \left(\frac{1}{4}, \frac{1}{2}\right), \qquad u^O = \left(1, \frac{1}{2}\right).$$

Thus, agents prefer school 2 when young and school 1 when old. If an agent attends school s when young, their age-two score at s becomes equal to 1, while their score at the other school is independently redrawn from $\text{Unif}[0, 1]$. Following the outside option, both scores are independently redrawn. Attending a school when young therefore guarantees access to it when old.

Fix school 2's cutoff at zero and write $x := P^1$. Conditional on school 1 being feasible when young, choosing it yields lifetime payoff

$$\frac{1}{4} + 1 = \frac{5}{4},$$

whereas choosing school 2 yields

$$\frac{1}{2} + (1 - x) \cdot 1 + x \cdot \frac{1}{2} = \frac{3}{2} - \frac{x}{2}.$$

The lifetime-payoff advantage of choosing school 1 is therefore

$$-\frac{1}{4} + \frac{x}{2}. \tag{20}$$

As x rises, waiting to enter school 1 becomes riskier. Once $x > 1/2$, the value of insuring future access exceeds the current payoff loss from forgoing school 2. Hence, when $x < 1/2$, young agents

choose school 2 even when school 1 is feasible. When $x > 1/2$, they instead choose school 1 whenever feasible, thereby securing access to it when old.

Let D^1 denote stationary demand in this benchmark economy. The preceding choices imply

$$D^1(x,0) = \begin{cases} \dfrac{1-x}{2}, & x < \dfrac{1}{2}, \\ \dfrac{2-x-x^2}{2}, & x > \dfrac{1}{2}. \end{cases} \tag{21}$$

Indeed, when $x < 1/2$, no young agent chooses school 1, while a fraction $1-x$ of old agents meets its cutoff. When $x > 1/2$, a fraction $1-x$ chooses school 1 when young and retains access when old; among the remaining fraction x, a fraction $1-x$ meets the newly drawn age-two cutoff. The value of demand at $x = 1/2$, where agents are indifferent, is immaterial for the argument.

Set capacities equal to

$$q^1 = \frac{11}{32}, \qquad q^2 = \frac{3}{4}. \tag{22}$$

Equation (21) gives

$$D^1\left(\frac{5}{16}, 0\right) = D^1\left(\frac{3}{4}, 0\right) = \frac{11}{32}.$$

Consequently,

$$P^- = \left(\frac{5}{16}, 0\right), \qquad P^+ = \left(\frac{3}{4}, 0\right) \tag{23}$$

both clear the market for school 1. Moreover, because school 2 has a zero cutoff and all benchmark scores at that school are nonnegative, every agent has at least one feasible school and total assignment is one. Hence, at either cutoff vector $P \in \{P^-, P^+\}$,

$$D^2(P) = 1 - D^1(P) = \frac{21}{32} < \frac{3}{4} = q^2.$$

Since school 2 has a zero cutoff, this strict inequality is consistent with market clearing. Thus, both P^- and P^+ are stationary market-clearing cutoff vectors.

The economic mechanism is simple. At the lower cutoff, agents attend school 2 when young and try to enter school 1 only when old. At the higher cutoff, future access to school 1 is sufficiently difficult that agents attend it early to secure priority. This insurance response creates the upward jump in stationary demand and hence the second market-clearing cutoff.

The benchmark is deliberately singular. Appendix C.4 replaces its concentrated payoff and transition laws with arbitrarily close atomless, absolutely continuous, full-support laws satisfying Assumptions 2.1 and 2.2 (hence, the jump in Figure 1 becomes a steep continuous increase). Because the two intersections above lie away from the indifference cutoff and demand crosses capacity strictly on each branch, one intersection survives near each of P^- and P^+. The strict slack at school 2 is also preserved. The appendix consequently shows that multiplicity persists under all sufficiently small perturbations of the capacity vector, the entry distribution, and the transition kernels.

Why static demand properties may fail. As we just saw, structural properties familiar from static matching markets need not carry over to the stationary dynamic environment. The underlying reason is that stationary demand need not satisfy the comparative-statics properties that discipline demand in the static model. There, raising a school's own cutoff weakly lowers demand for that school (monotonicity), whereas raising the cutoffs of other schools weakly increases it (gross substitute). In the dynamic model, however, a cutoff change also alters continuation values and

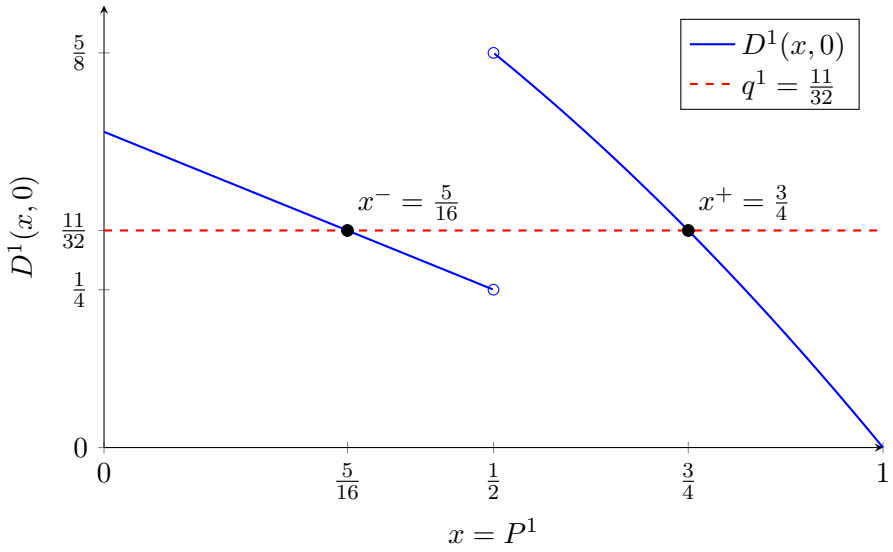


Figure 1: Stationary demand for school 1 in the benchmark multiplicity example. The dashed line is school 1's capacity. Its two intersections with stationary demand generate the cutoff vectors P^- and P^+.

the endogenous stationary distribution of types. These additional channels may reinforce, offset, or even reverse the direct substitution effect. This is particularly clear in the above example where we see that stationary aggregate demand is non-monotonic which critically drives the multiplicity (See Figure 1). Below we further illustrate that stationary demand can violate gross substitute.[15,16]

Example (Failure of gross substitutes). Let $\mathcal{S} = \{1, 2\}$ and $A = 2$. All entrants have

$$u^1 = 1, \qquad u^2 = \frac{1}{2}, \qquad r = (1, 1),$$

with payoffs fixed over time. Age-two scores are

$$r' = \begin{cases} (2/3, 0), & \text{after attending school 2,} \\ (0, 0), & \text{otherwise.} \end{cases}$$

Thus, school 2 provides future priority at school 1.

Compare

$$\underline{P} = \left(\frac{1}{2}, \frac{1}{2}\right) \quad \text{and} \quad \overline{P} = \left(\frac{3}{4}, \frac{1}{2}\right).$$

[15]Recall that gross substitutes requires that, for every school s and cutoff vectors $P \leq P'$ satisfying $P^s = P'^s$,

$$D^s(P') \geq D^s(P).$$

[16]Cox et al. (2025) study school choice when preferences depend on the composition of previous cohorts. They show that the resulting intercohort process may cycle because a gross-substitutes-like condition fails. Their peer-composition channel is different from our assignment-dependent evolution of individual payoffs and priorities, but both papers emphasize that objects endogenously generated by past assignments can reshape current demand and invalidate familiar static structure.

Under $\underline{P}$,

$$W_1^2(\theta \mid \underline{P}) = \frac{1}{2} + 1 > 1 = W_1^1(\theta \mid \underline{P}),$$

so age-one agents choose school 2 and $D^2(\underline{P}) = 1/2$. Under $\overline{P}$, the priority bonus no longer gives access to school 1, so

$$W_1^2(\theta \mid \overline{P}) = \frac{1}{2} < 1 = W_1^1(\theta \mid \overline{P}),$$

and $D^2(\overline{P}) = 0$. Hence,

$$\underline{P}^1 < \overline{P}^1, \qquad \underline{P}^2 = \overline{P}^2, \qquad D^2(\overline{P}) < D^2(\underline{P}),$$

violating gross substitutes.[17] The intuition behind the example is simple: Raising school 1's cutoff destroys the future priority reward attached to attending school 2, and therefore lowers current demand for school 2.

3.5 Finite Economies

The continuum model is intended as a tractable approximation of a large but finite market, in which an individual's market power is essentially negligible. In general, a version of the model with a finite number of agents in each generation introduces significant complications, both from the perspective of the analyst and the agents. In a finite economy, agents' reports directly influence the allocation of other agents, which indirectly determines tomorrow's scores, hence reported preferences not only impact future prospects through individual scores but also alter the competition that will be faced. This creates a complex inference problem, as agents need to form expectations about the types of others over time, the induced sequence of random cutoffs, and strategize accordingly. However, in large enough markets, it is natural to conjecture that a given agent's influence over the outcome becomes negligible. If so, at the individual level it should be approximately optimal to behave in a large enough market as in the continuum model – i.e. best respond to the predictable equilibrium cutoffs determined by aggregate behavior. The aim of this section is to establish such a result formally.

The finite economy. Assumptions and notations on time, schools, types, their distribution and transition probabilities are unchanged. For any continuum economy $\mathcal{E} = [\mathcal{S}, q, p, (Q_a)_{a=1}^{A-1}]$, denote $\mathcal{E}^n = [\mathcal{S}, q^n, p, (Q_a)_{a=1}^{A-1}]$ the corresponding finite economy with n agents per generation (i.e. nA agents present in every period), where $q_s^n = \frac{\lceil nAq_s \rceil}{nA}$ is the capacity of school s in the finite economy, expressed in normalized mass so that $q_s^n \to q_s$ as $n \to \infty$.

In each period, agents report a rank-ordered list of schools from $\mathcal{L}$; a fixed mechanism, which we typically take to be DA, maps the collection of reports and individual score vectors (determining priorities) to a stable allocation. A new cohort of n agents is born every period, drawing their type $\theta = (r, u)$ independently from p; an agent with type θ assigned to school $s \in \mathcal{S}_0$ at age a draws their next period type from $Q_a(\cdot|\theta, s)$. All draws are conditionally independent given the history of past realized types of all agents and allocations. This is exactly the finite economy analogue of the reporting game introduced in Sections 3.1 and 3.2, and it admits the same representation in terms of demands and cutoffs. The main difference is that aggregate demand is now discrete, and both realized demands and cutoffs depend on the (finite sample of) random type draws.

[17] The construction can easily be extended to continuous distributions concentrated near the specified types and transitions.

Index individuals who entered in t by $i(t) \in \{1(t), ..., n(t)\} =: C(t)$, so that $C(t)$ is the cohort of entrants at t. Let $I(t) = \cup_{a=1}^{A} C(t-a+1)$ the set of agents present at t, and $I_a(t) = C(t-a+1)$ the set of agents that are of age a at time t.

It is well known even in static environments that, in general, market power need not vanish as the number of agents grow large. Indeed, individual deviations may, in general, trigger arbitrary long chains of reallocations (see for instance Artemov et al., 2023). This may lead to failures of lower-hemicontinuity of equilibria when passing from the large-but-finite market to the continuum: some limit continuum equilibria are not the limit of any sequence of finite equilibria. In the dynamic environments, both the incentives for such strategic manipulations and their consequences compound, and the entire sequence of future type distributions and allocations may be altered as a result of a single deviation in a finite, yet arbitrarily large market. For this reason, which is in particular driven by the endogeneity of the evolution of the type distribution, we cannot apply existing results either from the matching literature (e.g. Artemov et al., 2023) or the literature on large games (e.g. Carmona and Podczeck, 2022). Furthermore, new issues arise in choosing an appropriate equilibrium concept for the finite economy, since we need to account for dynamic inference, which also requires taking a stance on what the players observe. To solve the issues above, we define an ex post notion of approximate equilibrium in the finite economy which enables direct comparison of strategic behaviour.

Approximate equilibrium in the finite economy. We need a notion of equilibrium in the finite economy which captures that there are no profitable deviations when the market is large enough, regardless of agents' private information and the random realization of types. Intuitively, say that strategy profile σ is an *ex-post asymptotic equilibrium* of the finite economy $\mathcal{E}^n$ if at any date t, the expected mass of agents with a profitable deviation from $\sigma_{t,a}$, using any reporting strategy conditioning on any subset of the full realized public history and computing continuation values accordingly, is vanishing as n goes to infinity. The following definition formalizes this idea.

Definition 3.3. *Fix a strategy profile $\sigma_{t,a} : \Theta \to \mathcal{L}$. For any t and $i \in I(t)$ let $G_i^n(\pi|\sigma)$ the expected continuation value, conditional on date t history (of types and assignments), from following reporting strategy $\pi \in \Pi_{t,i}$ when all other agents report according to σ, where $\Pi_{t,i}$ is the class of all reporting strategies over agent i's remaining lifetime conditioning on the date t history. Similarly denote $G_i^n(\sigma|\sigma)$ the expected payoff from following σ. The strategy profile σ is an* ex-post asymptotic equilibrium *if for every t and every $\varepsilon > 0$:*

$$\mathbb{E}\left[\frac{1}{nA}\left|\left\{i \in I(t) : \sup_{\pi} G_i^n(\pi|\sigma) - G_i^n(\sigma|\sigma) > \varepsilon\right\}\right|\right] \xrightarrow[n\to\infty]{} 0.$$

This notion is meant to capture a deliberately strong notion of robustness against finite market manipulations. First, it allows for deviations that exploit any information about the history, including the case where the deviator knows every other contemporaneous agent's type, so that the possible deviations of any less informed agents are subsumed. Second, it controls for the mass of agents with a strictly profitable deviation of arbitrary size.[18]

Main result: the continuum approximates large economies. Our central objective is to state that (i) it is an asymptotic equilibrium for agents to behave in a large economy as if it was

[18]Our definition controls the mass of agents with a profitable deviation at every threshold ε, and so implies the more common formulation in which the size of the deviation and the mass of agents are quantified by the same vanishing sequence. The proof in fact establishes the stronger statement that the expected average gain across agents vanishes, from which the definition follows by Markov's inequality.

the continuum economy and (ii) the outcome (types, cutoffs, assignments) in such a finite economy indeed converge to the corresponding continuum outcomes. To do so, we need to formalize one last relatively mild regularity condition on continuum strategies.

Definition 3.4. *A strategy profile $\sigma_{t,a} : \Theta \to \mathcal{L}$ is called* regular *if for all t, a, the set $\bigcup_{\ell \in \mathcal{L}} \partial\big(\sigma_{t,a}^{-1}(\{\ell\})\big) \subset \Theta$ has Lebesgue measure zero.*

Notice in particular that truth-telling strategies given equilibrium market-clearing cutoffs (which are effectively without loss within continuum equilibria in terms of induced outcomes, see Proposition 3.1) are always regular under Assumption 2.2. We are now ready to state the main result on the finite economy.[19]

Theorem 3.2. *Suppose Assumptions 2.1 and 2.2 hold. Let $\sigma = (\sigma_{t,a})$ be an equilibrium in the game induced by DA for the continuum economy $\mathcal{E}$, with regular strategies. Then σ is an ex-post asymptotic equilibrium in the finite economy $\mathcal{E}^n$ and for every date t, in the finite economy when agents report according to σ, we have that almost surely:*

(i) empirical type distributions converge weakly to $\eta_{t,a}^{\sigma}$;

(ii) agreggate demand functions converge uniformly to $\hat{D}_t^{\sigma}$;

(iii) realized market-clearing cutoffs converge to P_t.

There are two main reasons why this result is not a direct generalization of existing results for convergence of *static* economies to the continuum – both of them economically substantial. First, the market-manipulating deviations which must be ruled out have more breadth in the dynamic environment. In a static finite economy, a manipulation typically displaces a chain of other agents while leaving the deviator reassigned to the same school; as a result, the equilibria such manipulations support are payoff-equivalent to truthful ones. This is no longer true in the dynamic world: displacing others alters their assignments, hence their scores tomorrow, hence the cutoffs that the deviator will face. A manipulation that is payoff-neutral in the current period can therefore be strictly profitable through its effect on future competition. The relevant class of deviations is correspondingly larger, and ruling them out requires showing not only that a single agent moves current cutoffs by a vanishing amount, but that this displacement, compounded through the endogenous evolution of the type distribution over their lifetime, remains negligible.

Second, existing convergence arguments in static markets treat the finite economy as an independent sample from a fixed distribution of types, and obtain convergence of empirical demands from a law of large numbers applied once. In the dynamic environment, the date-t type distribution is itself generated by the date-$(t-1)$ allocation: it is random, endogenous, and correlated across agents. Hence, the proof requires transferring convergence across periods via a law of large numbers for types that are only conditionally independent and not identically distributed, applied at each date to a distribution that is itself part of the induction.

The proof is given in Appendix C.5. We first show that if the empirical type distributions at date t converge to their continuum counterparts, then so do the induced demands, uniformly in cutoffs. This implies that market-clearing cutoffs in the finite economy converge to the unique (datewise, given strategies) continuum market clearing cutoff, and that such aggregate convergence is robust to single-agent deviations. This is preserved by the transition to $t+1$ through the

[19] Although the theorem is stated for the game induced by DA, the same conclusion applies to any stable cutoff selector. Under our maintained assumptions, the reported aggregate demand at each date admits a unique static market-clearing cutoff vector, so all stable selectors generate the same cutoff.

continuity of the type kernels, once the dependence coming from realized cutoffs is removed by comparing each agent's finite market assignment with the one they would have received at the deterministic continuum cutoffs. Lastly, approximate optimality follows by comparing an agent's realized trajectory with the trajectory they would have had against continuum cutoffs: the two diverge only if one of the agent's scores falls close to a cutoff, which has vanishing probability. Whether the approximation can be made uniform in t is a distinct question, requiring conditions that limit the compounding of errors across dates – which is possible under additional regularity conditions.

4 Normative Analysis

The positive analysis treated the distribution of scores and its evolution as primitives and characterized the outcomes generated by repeated Deferred Acceptance. We now reverse the perspective. We fix a "payoff process" and ask whether restricting attention to score-based mechanisms entails a loss when the planner can design the entry distribution of scores and their transition kernels. This perspective is directly relevant in applications such as the assignment of teachers and other civil servants discussed in the Introduction, where the public authority explicitly determines both how priority points are initially assigned and how they evolve with agents' past assignments. Our *implementation principle* shows that every stationary equilibrium allocation generated by a broad class of mechanisms can be reproduced by repeated DA under an appropriately designed signed-score system.

The result provides the bridge between an abstract design problem and the applications studied below. One may first identify a desirable allocation using a comparatively rich mechanism and then use the implementation principle to translate that allocation into initial priorities and dynamic-priority rules that can be operated through a familiar matching procedure such as DA.

4.1 General Mechanisms

Throughout this section, we take an agent's primitive type to consist solely of their payoff vector. Payoffs upon entry have distribution $p^\star \in \Delta(\mathcal{U})$ and, for every $a < A$, evolve according to

$$Q_a^\star(\cdot \mid u, s) \in \Delta(\mathcal{U}), \qquad (u, s) \in \mathcal{U} \times \mathcal{S}_0.$$

The objects to be designed are the score distribution $\widehat{p}$ and the score-transition kernels $(\widehat{Q}_a)_{a<A}$ introduced in Section 2.

Let $\mathcal{Y} := [0, 1]$, endowed with the uniform distribution $\lambda_{\mathcal{Y}}$. A draw $y \in \mathcal{Y}$ is observed by the agent and the mechanism designer before the agent sends a message. The draw is a universal randomization device: any desired distribution on a standard Borel space can be generated from it through a measurable map. Let $\mathcal{X}$ be a compact standard Borel message space.

At age a, we distinguish two parts of an agent's observable history. The *physical history*

$$\mathcal{H}_a := \big(\mathcal{S}_0 \times \mathcal{Y}\big)^{a-1}, \qquad \mathcal{H}_1 := \{\varnothing\},$$

records past assignments and past realizations of the public randomization device. The *message record*

$$\mathcal{R}_a^m := \mathcal{X}^{a-1}, \qquad \mathcal{R}_1^m := \{\varnothing\},$$

records the agent's past messages. We write $(h, \rho) \in \mathcal{H}_a \times \mathcal{R}_a^m$ for the full observable history.

The mechanism may also condition on aggregate information. Let

$$\Xi := \prod_{a \in \mathcal{A}} \Delta\big(\mathcal{X} \times \mathcal{H}_a \times \mathcal{R}_a^m \times \mathcal{Y}\big).$$

An element $\xi = (\xi_a)_{a \in \mathcal{A}}$ records, for each age, the cross-sectional distribution of current messages, physical histories, message records, and current public draws. A mechanism is a collection of measurable allocation rules

$$\varphi_a : \mathcal{X} \times \mathcal{H}_a \times \mathcal{R}_a^m \times \mathcal{Y} \times \Xi \longrightarrow \mathcal{S}_0, \qquad a \in \mathcal{A}. \tag{24}$$

Thus, $\varphi_a(x, h, \rho, y \mid \xi)$ is the school assigned to an age-a agent who sends message x, has observable history (h, ρ) and current draw y, when the aggregate state is ξ. We restrict attention to feasible mechanisms, meaning that every allocation generated by the mechanism satisfies the capacity constraints.[20] We also maintain the outside-option convention of Section 2: at every observable state, an agent can choose a message that leads to $\emptyset$.

The formulation in (24) already incorporates three substantive restrictions on the allocation rule. It is anonymous because the allocation rule has no agent index; it is time invariant because it has no calendar-time index; and it is aggregative because the behavior of the rest of the market enters only through the cross-sectional distribution ξ, rather than through the identities or full profile of other agents. We therefore do not state these three properties as separate axioms.

We note that repeated agent-proposing DA is a particular case of the mechanisms considered here. Given an entry-score distribution and score-transition kernels, the score process can be represented by measurable mappings $R_a : \mathcal{H}_a \times \mathcal{Y} \to \mathcal{R}$, where $R_a(h, y)$ is the current score vector generated by the agent's physical history and public draw. Taking rank-order lists as messages, the aggregate state determines reported demand and the DA cutoff vector, and each agent is assigned to their highest-ranked school whose cutoff is met by $R_a(h, y)$. The resulting rule is feasible, anonymous, time invariant, and aggregative; the formal construction is provided in Appendix E.1.

Remark 4.1. *As will become clear, we do not allow for random draws realized after messages are submitted. Hence, conditional on the current message, history, public draw, and aggregate state, the assignment is deterministic. If an ex-post randomization device were introduced, the menus below would instead consist of lotteries over schools. The implementation principle could then be extended using an appropriate ex-post equilibrium notion, under which optimality is required conditional on every realization of the ex-post randomization device.*

4.2 The Implementation Principle

This subsection develops the implementation principle in three steps. We first introduce our main axiom to represent equilibrium behavior under a general mechanism as dynamic choice from menus, and introduce the stationary-equilibrium and outcome-equivalence notions needed for the comparison. We then show that every such stationary equilibrium allocation can be reproduced by repeated agent-proposing DA under an appropriately designed score process, before discussing extensions and outlining the construction.

[20] Formally, for every aggregate state $\xi = (\xi_a)_{a \in \mathcal{A}}$ and every school $s \in \mathcal{S}$,

$$\frac{1}{A} \sum_{a=1}^{A} \int \mathbf{1}\{\varphi_a(x, h, \rho, y \mid \xi) = s\}\, \xi_a(dx, dh, d\rho, dy) \le q^s.$$

Axiom and menus. The main implementation result requires the following restriction on the use of an agent's past messages.

Axiom 1 (No message records)**.** *For every age a, message $x \in \mathcal{X}$, physical history $h \in \mathcal{H}_a$, current draw $y \in \mathcal{Y}$, aggregate state $\xi \in \Xi$, and message records $\rho, \rho' \in \mathcal{R}_a^m$,*

$$\varphi_a(x, h, \rho, y \mid \xi) = \varphi_a(x, h, \rho', y \mid \xi).$$

Without Axiom 1, a past message can serve as a commitment device: agents with the same ordinal ranking but different cardinal intensities may self-select between a risky future claim to a preferred school and a safe assignment to a less preferred school. Example E.2 in the Appendix constructs an economy in which a mechanism that records messages implements the utilitarian first best, whereas every stationary equilibrium of every mechanism satisfying Axiom 1 yields strictly lower welfare.[21] We return to the role and substantive implications of this axiom after stating our implementation principle.

By Axiom 1, conditional on the aggregate state, the mechanism does not depend on the agent's own past messages. We henceforth suppress the message-record argument and write $\varphi_a(x, h, y \mid \xi)$. For $h \in \mathcal{H}_a$, let $h \oplus (s, y) \in \mathcal{H}_{a+1}$ denote the physical history obtained by appending the current assignment s and draw y.

Given an aggregate state ξ, the *menu* offered to an age-a agent with physical history h and current draw y is

$$\mathcal{M}_a(h, y \mid \xi) := \{\emptyset\} \cup \{\varphi_a(x, h, y \mid \xi) : x \in \mathcal{X}\}. \tag{25}$$

This is the set of schools the agent can obtain by varying their current message, holding their observable state and the aggregate state fixed. We assume throughout that the menu correspondence is measurable: for every a, s, and ξ, the set

$$\{(h, y) \in \mathcal{H}_a \times \mathcal{Y} : s \in \mathcal{M}_a(h, y \mid \xi)\}$$

is Borel.

Stationary equilibria for general mechanisms. We now define the stationary outcomes to which the implementation principle will apply. A time-invariant strategy is a collection of measurable maps

$$\sigma_a : \mathcal{U} \times \mathcal{H}_a \times \mathcal{R}_a^m \times \mathcal{Y} \longrightarrow \mathcal{X}, \qquad a \in \mathcal{A}.$$

Suppose that the age-specific distribution of (u, h, ρ, y) is η_a. The strategy induces the aggregate distribution

$$\xi_a^\sigma(B) := \int \mathbf{1}\{(\sigma_a(u, h, \rho, y), h, \rho, y) \in B\}\, \eta_a(du, dh, d\rho, dy) \tag{26}$$

for every measurable $B \subseteq \mathcal{X} \times \mathcal{H}_a \times \mathcal{R}_a^m \times \mathcal{Y}$. In words, ξ_a^σ is the cross-sectional distribution of the current message and observable state among age-a agents. Let $\xi^\sigma = (\xi_a^\sigma)_{a \in \mathcal{A}}$. The corresponding assignment is

$$\pi_a^\sigma(u, h, \rho, y) := \varphi_a(\sigma_a(u, h, \rho, y), h, y \mid \xi^\sigma). \tag{27}$$

[21] The underlying screening logic is closely related to Abdulkadiroğlu et al. (2011): in their static setting, the Boston mechanism can outperform DA in ex-ante welfare because strategic reports convey cardinal preference intensities that truthful ordinal reports under DA do not.

For compactness, define the next state generated by current state (u, h, ρ, y), next payoff u', and next public draw y' as

$$T_a^\sigma(u, h, \rho, y; u', y') := \big(u', h \oplus (\pi_a^\sigma(u, h, \rho, y), y), \\ \rho \oplus \sigma_a(u, h, \rho, y), y'\big).$$

The age-specific distributions are stationary when entrants satisfy

$$\eta_1 = p^\star \otimes \delta_\varnothing \otimes \delta_\varnothing \otimes \lambda_\mathcal{Y}, \tag{28}$$

and, for every measurable $B \subseteq \mathcal{U} \times \mathcal{H}_{a+1} \times \mathcal{R}_{a+1}^m \times \mathcal{Y}$,

$$\eta_{a+1}(B) = \int\int \mathbf{1}\{T_a^\sigma(u, h, \rho, y; u', y') \in B\}\lambda_\mathcal{Y}(dy') \\ \times Q_a^\star(du' \mid u, \pi_a^\sigma(u, h, \rho, y))\eta_a(du, dh, d\rho, dy). \tag{29}$$

For a fixed stationary aggregate state ξ, define values directly from the menus. At the terminal age,

$$U_A^\varphi(u, h, y \mid \xi) = \max_{s \in \mathcal{M}_A(h,y|\xi)} u^s. \tag{30}$$

For every $a < A$,

$$U_a^\varphi(u, h, y \mid \xi) = \max_{s \in \mathcal{M}_a(h,y|\xi)} \left\{ u^s + \int_\mathcal{U}\int_\mathcal{Y} U_{a+1}^\varphi(u', h \oplus (s, y), y' \mid \xi)\lambda_\mathcal{Y}(dy')Q_a^\star(du' \mid u, s) \right\}. \tag{31}$$

Let Λ_a be the product of Lebesgue measure on payoff types and public draws and counting measure on past assignments. We next impose two assumptions that adapt Assumptions 2.1 and 2.2 to the broader mechanism-design environment considered in this section.

Assumption 4.1 (Regular payoff dynamics)**.** *The entry distribution $p^\star$ and, for every $a < A$, $u \in \mathcal{U}$, and $s \in \mathcal{S}_0$, the transition law $Q_a^\star(\cdot \mid u, s)$ are absolutely continuous with respect to Lebesgue measure on $\mathcal{U}$.*

Assumption 4.2 (Generic no-indifference condition)**.** *For every aggregate state $\xi \in \Xi$ and age $a \in \mathcal{A}$, the set*

$$\mathcal{I}_a^\varphi(\xi) := \left\{ (u, h, y) : \begin{array}{l} \text{the maximization problem in (30) or (31)} \\ \text{has more than one solution} \end{array} \right\}$$

has Λ_a-measure zero.

Informally, a stationary equilibrium consists of a time-invariant strategy and age-specific state distributions that reproduce the same aggregate state at every date, while assigning almost every agent a dynamically optimal school from their current menu; the definition below parallels Definition 3.2, which introduced the same notion for the specific game induced by repeated DA.

Definition 4.1 (Stationary equilibrium under a general mechanism)**.** *A time-invariant strategy $\sigma = (\sigma_a)_{a \in \mathcal{A}}$ is a stationary equilibrium of the game induced by φ if there exist age-specific distributions $(\eta_a)_{a \in \mathcal{A}}$ such that:*

1. *equations* (26)– (29) *induce time-invariant aggregate distributions and assignments; and*

2. *for every age a, the assignment $\pi_a^\sigma(u,h,\rho,y)$ is dynamically optimal among the schools in $\mathcal{M}_a(h,y \mid \xi^\sigma)$ for η_a-almost every (u,h,ρ,y).*

Assumptions 4.2 and 4.1 ensure that, in every stationary equilibrium, the dynamically optimal school is unique for almost every state at every age. Under Axiom 1, this also means that the equilibrium assignment is independent of the agent's message record almost everywhere. Lemma D.1 in Appendix D states and proves these claims formally.

4.2.1 Result

We first make precise the outcome comparison used below in our implementation principle. Consider a stationary allocation π generated by a general mechanism, with age-specific state distributions $(\eta_a)_a$, and a stationary allocation $\widehat{\pi}$ generated by a score economy, with age-specific type distributions $(\widehat{\eta}_a)_a$. We say that π and $\widehat{\pi}$ are *equivalent* if, for every age a, Borel set $B \subseteq \mathcal{U}$, and assignment $s \in \mathcal{S}_0$,

$$\widehat{\eta}_a\big(\{(r,u) : u \in B,\ \widehat{\pi}_a(r,u) = s\}\big) = \eta_a\big(\{(u,h,\rho,y) : u \in B,\ \pi_a(u,h,\rho,y) = s\}\big). \tag{32}$$

Thus, at every age, the two systems induce the same joint distribution of payoff types and assignments.

Definition 4.2 (Implementation by Deferred Acceptance)**.** *We say that repeated agent-proposing DA* implements *a stationary allocation π if there exist an entry-score distribution $\widehat{p}$, score-transition kernels $(\widehat{Q}_a)_{a<A}$ and a stationary equilibrium of the dynamic game induced by repeated DA whose allocation $\widehat{\pi}$ is equivalent to π.*

Our result below states that, once the payoff process is fixed, any stationary equilibrium allocation generated by a feasible mechanism satisfying Axiom 1 can be reproduced under repeated DA by appropriately designing agents' initial scores and their assignment-dependent evolution. The implementation preserves, at every age, the joint distribution of payoff types and assignments, even though the underlying scores, messages, and histories may differ across the two systems.

Theorem 4.1 (Implementation principle)**.** *Suppose that Assumptions 4.1 and 4.2 hold. Let φ be a feasible mechanism satisfying Axiom 1, and let σ be a stationary equilibrium. Then π^σ is implementable by repeated agent-proposing DA.*

The result admits two complementary interpretations. First, it provides a foundation for the score- and priority-based procedures used in many applications, including the assignment of teachers to schools in France discussed in the Introduction. It shows that these existing practices cover the whole set of mechanisms satisfying Axiom 1. Although these mechanisms form a broad class, the common restriction imposed by this axiom may be seen substantive: allowing assignments to depend on agents' past messages may improve welfare, as illustrated by Example E.2.

Second, conditional on Axiom 1, the theorem shows that restricting attention to repeated DA entails no loss in terms of stationary-equilibrium allocations. Although the axiom is substantive, the implementation argument remains surprisingly robust. In particular, one can prove that the same implementation conclusion continues to hold without Axiom 1 if repeated DA retains the rank-order lists submitted in previous rounds and the designed score-transition kernels are allowed to condition on the accumulated ROL record. The current matching is still computed by ordinary DA. If the original message space $\mathcal{X}$ contains at most $S!$ messages, then, at each state, its messages can be encoded by distinct rank-order lists that induce the corresponding current assignments. The ROL record can therefore reproduce the message-dependent continuation of the original

mechanism. Likewise, stationarity is used only to obtain a time-invariant score system. Nonstationary equilibrium allocations can be accommodated by allowing the entry-score distribution and the score-transition kernels to depend on calendar time. Finally, there are settings in which one may wish to retain scores in $[0,1]^S$. In particular, legal or administrative rules sometimes prevent the designer from excluding an agent from a school that has available seats. Appendix D shows formally that, under the corresponding non-exclusion condition, the target allocation can also be implemented by stable mechanisms.

4.2.2 Proof Sketch

Fix the stationary equilibrium σ and its aggregate state ξ^σ. The proof proceeds in four steps.

First, Axiom 1 makes the agent's current message relevant only through the school it produces. Conditional on the physical state (h, y), the agent can therefore be viewed as choosing their dynamically preferred school from the menu $\mathcal{M}_a(h, y \mid \xi^\sigma)$.

Second, encode both this menu and the physical state in a score vector. In order to do this, choose a Borel injection $\kappa_a : \mathcal{H}_a \times \mathcal{Y} \to (-1/4, 1/4)$, with a Borel inverse on its image, and define

$$g_a^s(h, y) := \kappa_a(h, y) + \begin{cases} \frac{1}{2}, & s \in \mathcal{M}_a(h, y \mid \xi^\sigma), \\ -\frac{1}{2}, & s \notin \mathcal{M}_a(h, y \mid \xi^\sigma). \end{cases} \tag{33}$$

At the zero cutoff, a school is feasible exactly when its score is positive, so the sign of each coordinate reproduces the current menu. At the same time, adding or subtracting $1/2$ from any coordinate recovers the common term $\kappa_a(h, y)$. Since κ_a is injective, the score vector then recovers the precise state (h, y), not merely the menu.

Third, use this decoding to construct the score process. The entry-score distribution is the law of $g_1(\varnothing, y)$. After a current assignment s, the score-transition kernel decodes (h, y) from the current score, appends (s, y) to the physical history, draws the next public randomization, and applies g_{a+1}. Hence, after every possible current school, the score economy generates exactly the same conditional distribution of next-period physical states and menus as the original mechanism. Starting from the terminal age and working backward, each current school therefore has the same expected continuation value in the two systems. Together with the same current menu and flow payoff, this gives the same dynamic choice problem and, by the almost-everywhere uniqueness result proved in the Appendix, the same assignment almost everywhere.

Finally, set the cutoff vector equal to $P = (0, \ldots, 0)$. Because the constructed feasible sets reproduce the mechanism's menus, feasibility of φ implies that demand at P does not exceed capacity at any school. Since no component of P is positive, these are sufficient for market clearing. Moreover, the zero vector is the smallest admissible cutoff vector and is therefore selected by agent-proposing DA. Applying arguments similar to those for Proposition 3.1, we can then construct a truthful stationary equilibrium under repeated DA that induces this cutoff vector. The allocation induced by this equilibrium is therefore equivalent to the stationary equilibrium allocation induced by σ under the original mechanism φ.

Thus, subject to Axiom 1, the planner may search for desirable stationary equilibrium allocations in a rich mechanism class and translate the selected allocation into a score system operated through repeated DA. The next section uses this principle to design dynamic priorities in the presence of distributional objectives.

5 Applications

5.1 Implementation of a Constrained Efficient Allocation via Deferred Acceptance

In this section, we return to the paper's general design question. Suppose that the designer's objectives can be expressed as distributional constraints on the allocation. How can one construct a priority system that selects, among the allocations satisfying these constraints, an allocation that is Pareto efficient for the agents? Drawing on our implementation principle, we provide a two-step computational procedure for identifying and implementing such an allocation under DA.

We describe here our method in a simple environment in which each agent knows at entry the entire path of flow payoffs over schools. The problem can then be understood as a lifetime assignment problem. Our method however extends to more general environments where agents learn their future payoffs over time according to transition kernels, as detailed in Appendix F.1.2.

We focus on a complete information environment. Agents are born with a payoff path $\bar{u} := (\bar{u}_1, \cdots, \bar{u}_A) \in ([0,1]^S)^A$, drawn from a probability measure p on $([0,1]^S)^A$. Future payoffs are discounted at a common factor $\delta \in (0,1)$. The value of the outside option is normalized to 0, so that all schools are acceptable to the agents. We assume that indifferences have mass zero.

Assumption 5.1. *For every pair of distinct lifetime paths $(s_1, \cdots, s_A)$, $(s'_1, \cdots, s'_A) \in \mathcal{S}_0^A$,*

$$p\Big(\Big\{\bar{u} : \sum_{a=1}^{A} \delta^{a-1} \bar{u}_a(s_a) = \sum_{a=1}^{A} \delta^{a-1} \bar{u}_a(s'_a)\Big\}\Big) = 0. \tag{34}$$

An allocation specifies, for each agent, the school she attends at every age of her career. Formally, an allocation π assigns to each age a and each payoff type a school $\pi_a \in \mathcal{S}_0$, and induces age-specific distributions $(\eta_a)_{a \in \mathcal{A}}$ over states, as in Section 4. Such an allocation can be generated by a general mechanism from the class described in Section 4, but for the moment we treat π as a primitive object. The designer's objectives are defined on the aggregate assignment induced by the allocation: for every age a and school s, let

$$M_{a,s}(\pi) := \eta_a\big(\{\cdot \ : \ \pi_a(\cdot) = s\}\big)$$

denote the share of the age-a generation assigned to school s, so that the mass of agents attending s at age a is $M_{a,s}(\pi)/A$.

The designer's objective takes the form of a finite collection Ω of linear inequalities on aggregate assignments. Each $c \in \Omega$ is described by coefficients $(\alpha^c_{a,s})_{a,s}$, with $\alpha^c_{a,\emptyset} = 0$, and a bound b^c, and requires

$$\Phi_c(\pi) := \sum_{a=1}^{A} \sum_{s \in \mathcal{S}_0} \alpha^c_{a,s} M_{a,s}(\pi) \ \leq \ b^c. \tag{35}$$

Coefficients are unrestricted in sign, so that (35) accommodates floors as well as ceilings. We write

$$\mathcal{C} := \{\pi : \ \Phi_c(\pi) \leq b^c \ \text{ for every } c \in \Omega\}$$

for the set of feasible allocations, which we assume nonempty.

Such constraints capture a wide range of desiderata. First, capacity constraints are naturally expressed as linear inequalities: $\frac{1}{A}\sum_a M_{a,s}(\pi) \leq q^s$. Richer instruments are also available, such as caps on the mass of young agents assigned to given positions, experience floors for some positions,

or bounds on the average seniority of an institution. Examples of such constraints are found in many applications. The Japanese medical residency match imposes regional caps on the number of doctors (Kamada and Kojima, 2015). In the Italian judiciary, magistrates at their first appointment cannot be assigned to certain positions before obtaining their first professional evaluation — an experience floor attached to specific positions — while tenure in the same investigative positions is capped at ten years.[22]

Within $\mathcal{C}$, we seek an allocation that is Pareto efficient for the agents. Write $U_a(\pi; \bar{u})$ for the payoff from age a onwards of an agent with payoff path $\bar{u}$ and assignment path $(s_1, \cdots, s_A)$ under π,

$$U_a(\pi; \bar{u}) := \sum_{t=a}^{A} \delta^{t-a} \bar{u}_t(s_t).$$

Definition 5.1. *An allocation $\pi^\star \in \mathcal{C}$ is dynamically $\mathcal{C}$-constrained efficient if there is no $\pi \in \mathcal{C}$ such that, for every age $a = 1, \cdots, A$,*

$$U_a(\pi; \bar{u}) \;\geq\; U_a(\pi^\star; \bar{u})$$

for p-almost every payoff path $\bar{u}$, with strict inequality for some age a on a set of paths of positive p-measure.

We show that such an allocation can be implemented under DA by designing appropriately entry scores and their evolution.

Proposition 5.1. *Fix a finite collection Ω of linear constraints and suppose Assumption 5.1 holds. Then there exists a dynamically $\mathcal{C}$-constrained efficient allocation that is implementable by repeated agent-proposing DA in the sense of Definition 4.2.*

The proof is constructive and follows a two-step procedure: (i) we first obtain a constrained-efficient allocation as the equilibrium outcome of a pseudo-market in which each entrant receives a random lifetime budget and constraints receive shadow prices; (ii) the implementation principle then translates the pseudo-market allocation into an entry-score distribution and assignment-dependent score transitions implementing it under DA.

The dynamic pseudo-market we build draws on the methodology of Echenique et al. (2021). Assigning a shadow price to each constraint in Ω directly is not possible as inequalities imposing floors would carry negative prices. We therefore work with a finite collection $\Omega^\star$ of linear inequalities with nonnegative coefficients that characterize the lower contour set of $\mathcal{C}$. Each constraint in $\Omega^\star$ receives a shadow price, which internalizes the externalities that an agent attending school s at age a exerts on others by potentially tightening the constraints involving s. A first-welfare-theorem argument guarantees that any equilibrium allocation is efficient within the original constrained set $\mathcal{C}$.

Formally, an entrant's budget B is drawn from an atomless distribution ν on $[0, 1]$, independently of her payoff path. An allocation is described by an array $x := (x_{a,B,\bar{u},s})_{a,B,\bar{u},s}$ such that $x_{a,B,\bar{u},s}$ takes value 1 if and only if an agent with budget B and payoff path $\bar{u}$ is assigned school s at age a. Such an array induces an allocation in the sense above, the entrant's budget being a component of the type; the definitions of the aggregate assignment, of feasibility and of constrained efficiency therefore apply unchanged, with $\nu \otimes p$ in place of p. In particular, the share of age-a agents assigned to school s is $M_{a,s}(x) = \int x_{a,B,\bar{u},s}\, \nu(dB)\, p(d\bar{u})$.

[22] Decreto legislativo 5 aprile 2006, n. 160, artt. 10–13 and 19, as amended by legge 30 luglio 2007, n. 111.

Every constraint $c \in \Omega^\star$ is described by nonnegative coefficients $(\alpha^{\star c}_{a,s})_{a,s}$, with $\alpha^{\star c}_{a,\emptyset} = 0$, where $\alpha^{\star c}_{a,s}$ measures the contribution of an age-a agent assigned to school s to the use of c, and by a bound $b^{\star c}$. The use of c under x is

$$\Phi^\star_c(x) := \sum_{a=1}^{A} \sum_{s \in \mathcal{S}_0} \alpha^{\star c}_{a,s} M_{a,s}(x),$$

and the set of allocations satisfying the collection of constraints $\Omega^\star$ is

$$\mathcal{C}^\star := \{x : \ \Phi^\star_c(x) \leq b^{\star c} \ \text{ for every } c \in \Omega^\star\}.$$

By construction of $\Omega^\star$, any allocation in $\mathcal{C}^\star$ satisfies every constraint in Ω.

Shadow prices on constraints translate into prices on schools, and agents buy their favorite sequence of schools among the affordable ones. Since an agent assigned to school s contributes to a given constraint c differently depending on her age, school prices vary across generations. Given a vector of constraint prices $P = (P_c)_{c \in \Omega^\star} \in \mathbb{R}^{\Omega^\star}_+$, define generation-school prices by

$$P_{a,s} := \sum_{c \in \Omega^\star} \alpha^{\star c}_{a,s} P_c, \qquad \forall s \in \mathcal{S}_0,$$

so that the outside option is free, $P_{a,\emptyset} = 0$. These prices allow us to select a constrained-efficient allocation, through the following equilibrium concept.

Definition 5.2. *A pair $(x^\star, P^\star)$ is a stationary pseudo-market equilibrium if:*

1. *(feasibility) $x^\star \in \mathcal{C}^\star$;*
2. *(optimality) for $(\nu \otimes p)$-almost every type $(B, \bar{u})$, the school path $(s^\star_1, \cdots, s^\star_A)$ that $x^\star$ assigns to type $(B, \bar{u})$ solves*

$$\max_{(s_1, \cdots, s_A) \in \mathcal{S}_0^A} \ \sum_{a=1}^{A} \delta^{a-1} \bar{u}_a(s_a) \qquad \textit{subject to} \qquad \sum_{a=1}^{A} P^\star_{a,s_a} \leq B;$$

3. *(positive price for binding constraints) for every $c \in \Omega^\star$,*

$$P^\star_c > 0 \quad \Longrightarrow \quad \Phi^\star_c(x^\star) = b^{\star c}.$$

In words, an equilibrium requires that the allocation satisfy all the constraints; that almost every agent be assigned a lifetime sequence of schools maximizing her discounted lifetime payoff among the sequences whose cumulated price does not exceed her budget; and that any constraint left slack be free.

We show in Appendix F that a stationary pseudo-market equilibrium exists and that every stationary pseudo-market equilibrium allocation is constrained efficient. This dynamic pseudo-market mechanism hence identifies a constrained-efficient allocation.

We now show that our implementation principle translates the pseudo-market allocation into an entry-score distribution and score transitions under DA that implement the constrained-efficient allocation.

First, note that in the dynamic pseudo-market defined above, agents choose a full path of assignments given their budget, which differs from the sequential mechanisms considered in Section 4. However, because the entire payoff path is known at age 1, the lifetime problem above can be

implemented sequentially using residual budgets. In that case, using backward induction, after any history of past choices $(s_1^*, \ldots, s_{a-1}^*)$, agents would pick their most favorite school given their current residual budget. The resulting budget mechanism is anonymous, time-invariant and aggregative. The message sent by an agent in each period is his dynamically preferred school among feasible ones. In equilibrium, this message coincides with the assignment, and the mechanism thus conditions only on the history of past assignments, through gradual depletion of the budget. The mechanism therefore satisfies Axiom 1, and the allocation identified through the dynamic pseudo-market is a stationary equilibrium allocation of a mechanism in the class described in Section 4.

Remark 5.1. *The complete-information environment does not satisfy Assumption 4.1: the entire payoff path is drawn at entry and realized deterministically thereafter, so the transition laws $Q_a^\star(\cdot \mid u, s)$ are not absolutely continuous with respect to Lebesgue measure on $\mathcal{U}$. This is immaterial here. Absolute continuity is used in the proof of Theorem 4.1 only to establish that the dynamic optimization problem admits a unique solution at almost every state and every age. Assumption 5.1 delivers this directly. Theorem 4.1 thus applies.*

While the proof of Theorem 4.1 establishes the existence of an entry-score distribution and of score-transition kernels implementing a given allocation, it does not deliver them in closed form because the encoding of menus into scores relies on an abstract Borel isomorphism. For the dynamic pseudo-market, we show in Appendix F.1.3 that the structure of the budget mechanism yields an explicit construction, in which entry scores are an affine function of the entrant's budget and scores evolve only through additive school- and age-dependent bonuses in each period. Note, however, that the priority rule delivered by this construction is generally richer than the uniform bonuses commonly observed in practice. The additive bonuses may indeed depend on the current assignment, the age, and the prospective destination as we now document in simulations.

5.2 Quantitative Exploration: The Welfare-Inequality Frontier

We now turn to a quantitative analysis illustrating how the method exposed above can be used in practice.

We simulate a stationary economy with a set of 50 assignment destinations, $S = \{1, \cdots, 50\}$ and an outside option. This quantitative exercise offers a proof of concept that priority systems can be computed at an empirically relevant scale. In particular, 50 assignment destinations is larger than the number of destinations in the French national inter-academic teacher-mobility round, which reallocates teachers across 30 "académies" each year. Capacities satisfy $q^s = 1/50$ for all s and $\sum_s q^s = 1$. Agents live for $A = 40$ periods, which approximates a full career of assignments.

Relying on the extension of our dynamic pseudo-market analysis to the case where preferences are not known ex ante, we adopt a logit utility specification for flow payoffs: the utility of attending school s in period a is the sum of a possibly age-dependent intrinsic value term $\bar{u}_a^s$ and an idiosyncratic shock. We define

$$u_a^s = \bar{u}_a^s + \varepsilon_{a,s},$$

where $\varepsilon_{a,s}$ is an i.i.d. Gumbel$(0, 1)$ shock. Intrinsic values are age-specific:

$$\bar{u}_a^s \;=\; \lambda \, \frac{s}{N} \, \varphi(a), \qquad \varphi(a) \;=\; \varphi_0 + (1 - \varphi_0) \, \frac{a-1}{A-1}, \qquad a = 1, \ldots, A,$$

with $\lambda = 3$, $\varphi_0 = 0.10$ and $N = 50$. The ladder s/N is evenly spaced on $(0, 1]$; the profile rises linearly from $\varphi(1) = 0.10$ to $\varphi(A) = 1$, so a new entrant perceives one tenth of the quality differences that a teacher at the end of her career perceives. This represents the fact that agents tend to settle

as they grow older, making them more sensitive to where they are matched. Agents discount the future at rate $\delta = 0.95$.

We follow the application section by assuming that all schools are acceptable, and the outside option's value is lower than the flow payoff of being matched at any school. Since the support for the realized flow payoffs is $\mathbb{R}$, we set the value of the external option to $u_a(\emptyset) = -\infty$ for all a.

We reproduce the dynamic pseudo-market described above and endow each agent with an initial budget $B \sim U[0, 1]$. School prices are computed from the prices on constraints, as described in Section 5.1. Each equilibrium is computed using a deterministic simulation of the continuum economy. Demands are computed using exact propagation of the joint budget-discount factor distribution on a fixed grid, and a complementarity Newton solver is used to compute equilibrium prices. Prices below the grid resolution are approximated to 0 and allocations to free schools are determined by lottery rationing. Appendix F.6 describes the method in detail and reports a certification of every described point.

The capacity-only economy In the absence of specific distributional constraints of the designer, the allocation delivered through the dynamic pseudo-market features a sorting of experience with school quality. Because quality is worth less to young agents, the optimal policy for an agent is to save budget early to buy quality later.

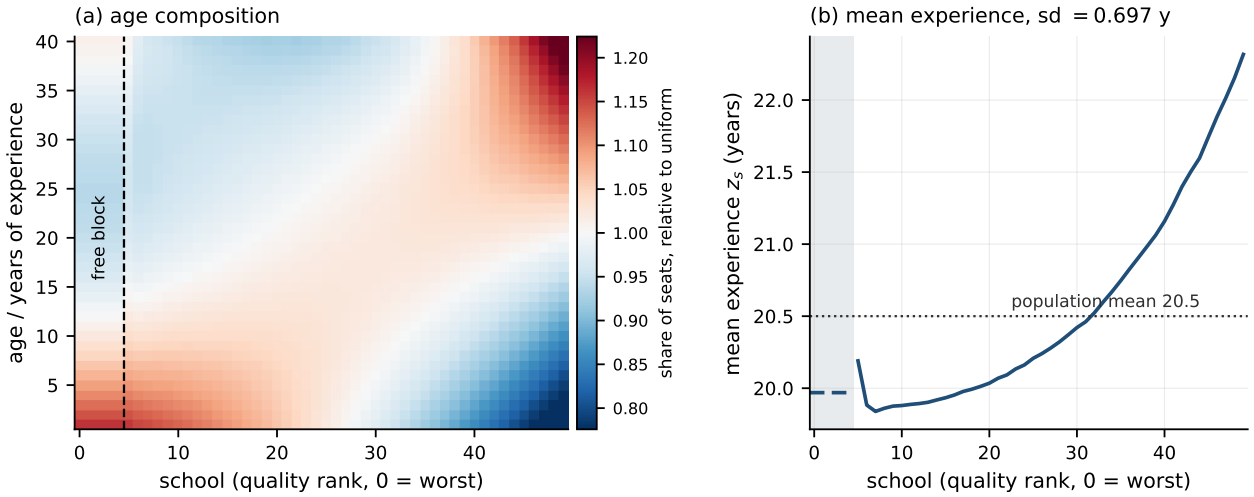


Figure 2: Age composition of each school, relative to the uniform composition $1/A$. Red: over-represented. The free block appears as a single aggregate column: every school in it has a zero price, so its seats are rationed by lottery rather than by price and the split across them carries no economic content (see Appendix F.6.2 for more details).

Mean experience is spread from 19.8 years at the least experienced school to 22.3 at the most experienced, and the novice share runs from 0.199 at the best school to 0.278 at the worst, against a uniform value of 0.250 (see Figure 2). We summarize dispersion in experience by the capacity-weighted standard deviation of mean experience:0.697 years.

Prices are the highest at the best schools and constant across ages since capacity constraints weight all ages equally. A year at the best school costs 9.0% of the mean entrant's lifetime budget, while a year at the worst school is free. Figure 3 converts the whole schedule into the number of years each budget can sustain at each school: the richest possible entrant can buy at most 22 years at the top school, and the median entrant can spend her whole career at school 33 or worse.

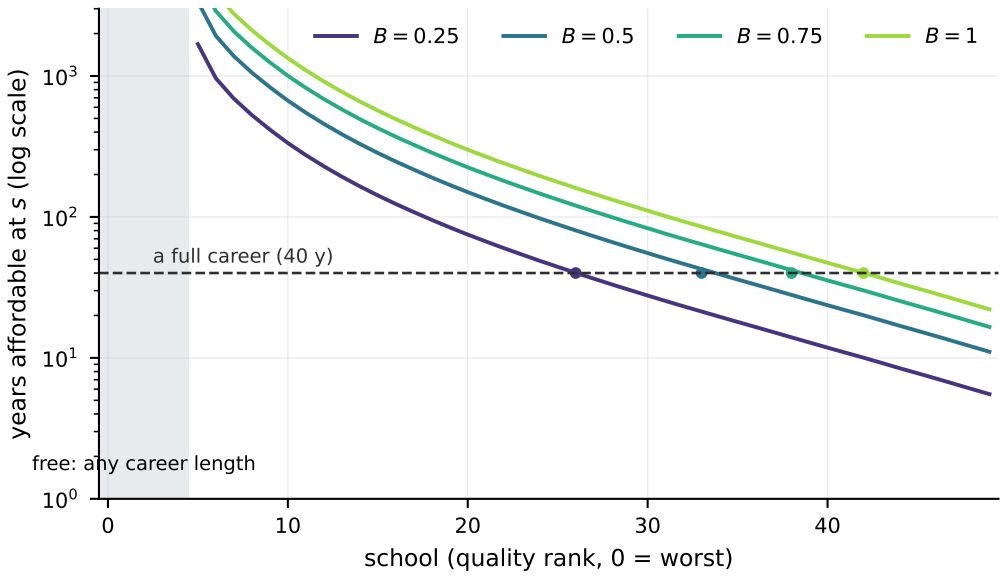


Figure 3: Affordability of the capacity-only price schedule. For four entry budgets B, the curve gives B/P_s, the number of years an agent could finance at school s if she spent her whole budget there; the horizontal line marks a full career of 40 years and the dots mark, for each budget, the best school that can be sustained for a full career. Schools in the free block carry a zero price and are omitted from the log scale.

Policy instrument We consider the case of a policy maker who is interested in compressing the dispersion of agents' experience across schools, moving toward equal compositions.

An instrument often used in practice is a novice cap[23]: no school is allowed to fill more than a fraction κ of its capacity with agents in their first a_0 years,

$$\frac{1}{Aq^s}\sum_{a\le a_0} D_{a,s} \le \kappa, \quad \forall s$$

In the following we fix $a_0 = 10$. Note that imposing the constraint on all schools makes sure that the constraint is always satisfied at the most novice-heavy school, whose identity is not exogenously fixed but varies with the equilibrium. The population share of novices is $a_0/A = 0.25$ so $\kappa = 0.25$ is targeting equal composition. The capacity-only allocation attains $\frac{1}{Aq^s}\sum_{a\le a_0} D_{a,s} = 0.278$ at its most novice-heavy school so no κ above that value will bind.

This instrument formalizes policies that are used or discussed in practice: the French ministry recommends that trainee teachers not be assigned to the most difficult schools ("éducation prioritaire") and the Cour des Comptes recommended in a 2018 report to restrict positions in these schools to teachers with at least two years of seniority[24]. More generally, until 2016, Judges in Italy

[23]Another possible instrument would be to constrain the average age at all schools to be above a threshold, thus ensuring that the average experience at the worst-off school is above this threshold. Such constraints are naturally defined as floor constraints and require some pre-processing before computing the equilibrium. The $|S|$ floor constraints on the average age at each school however translate to $(2^S - 1)(A + 1)$ pre-processed constraints to be used in the equilibrium computation as shown in Appendix F.5. This imposes a burden on equilibrium computation which proves particularly expensive for a frontier computation exercise where the equilibrium is to be recomputed several times.

[24]Cour des comptes, L'éducation prioritaire, rapport public thématique, octobre 2018, p. 113 (Recommandation n° 8), https://www.ccomptes.fr/fr/publications/leducation-prioritaire

could not be assigned to certain positions before their first professional evaluation which takes place in general around four years after the entry in service [25].

Constrained equilibrium under the policy instrument We now compute a constrained equilibrium of this economy for $\kappa = 0.274$, showing the tractability of the equilibrium computation at a relevant scale. Each equilibrium involves $2,000$ generation-school cells and 100 constraint prices and is solved in minutes, certified against an independent finer grid. We then derive bonuses from the price schedule, allowing us to implement the equilibrium allocation under DA.

At $\kappa = 0.274$, the cap binds at the two most novice-heavy schools. The price system reacts to the distributional constraints in the following way: the schools that were free at the capacity-only equilibrium remain free for old agents while two of them acquire a positive price for novices, representing 0.010% and 0.007% of the mean lifetime budget. This price schedule allows dispersion to fall from 0.697 to 0.684.

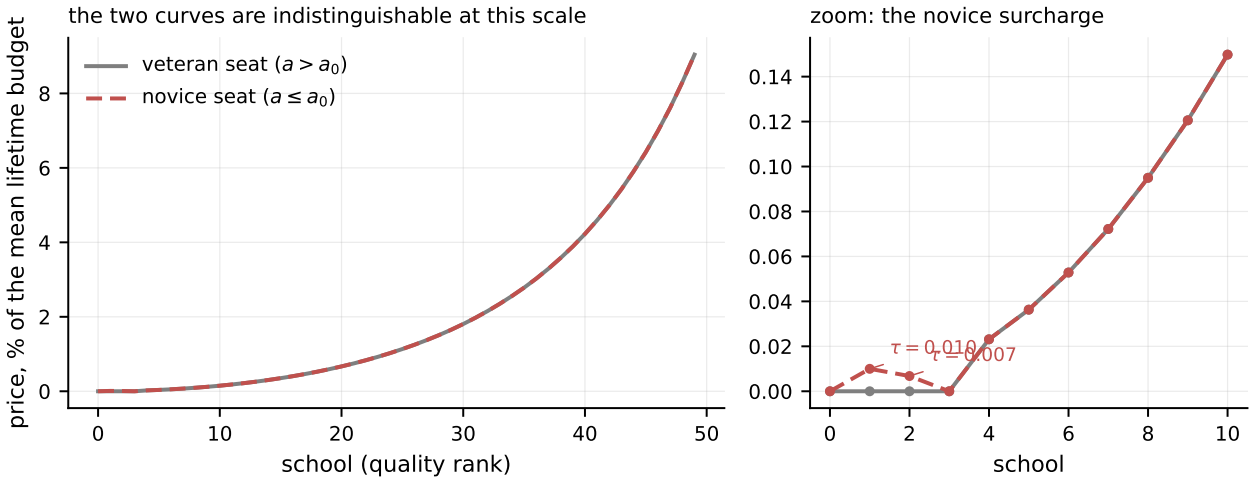


Figure 4: Equilibrium prices at $\kappa = 0.274$, in percent of the mean lifetime budget. Left: the price of a novice seat and of a veteran seat at each school; the two schedules are indistinguishable at the scale of the quality ladder. Right: the same two schedules over the cheapest ten schools, where the novice surcharge sits.

A relevant policy question is whether more equality is attained at a high welfare cost and the above exercise can be repeated along several bounds to trace the frontier of attainable equality-welfare pairs.

Table 1 reports 3 points from that frontier. We measure the welfare cost of a constraint ordinally, by the average number of rungs of the fifty-school ladder that agents move: at every point, we compare the constrained and capacity-only assignments and report the mean displacement for novices and for agents in their last decade. The table shows that the novice cap is essentially free in these terms.[26] However, below $\kappa = 0.274$, more equality is obtained at the cost of excluding some of the novices from the market: as the bound tightens more initially free schools start to bear a positive price at young age until no schools are free at young ages. From that point, a novice must pay a positive price in each of her a_0 first periods and the excluded mass is equal to a_0 times the cheapest novice price. Only novices are excluded as the novice cap constraints only weigh young

[25] Decreto legislativo 5 aprile 2006, n. 160, art. 11, as amended by legge 30 luglio 2007, n. 111.

[26] In general, note that the magnitudes we work with are small throughout because the capacity-only equilibrium is already fairly balanced.

generations. Old agents thus always keep access to free schools. As a consequence, exclusion is temporary. Poor novices drop out during their first years before entering free schools again when old. This phenomenon undermines the effectiveness of the instrument for two reasons: (i) an agent who spends her first years outside of the market is unlikely to have accumulated experience; (ii) in practice vacant seats are generally not left empty but filled with substitutes whose experience is arguably lower than that of novices. Appendix F.6.5 shows that finer instruments, that treat each generation differently rather than through two coarse age blocks can reduce distributional imbalances without creating exclusion. In particular, putting a ceiling on the average age at all schools allows the designer to compress dispersion by 31% while matching every agent.

κ	schools surcharged	sd(z_s)	rank shift, novices	rank shift, veterans	excluded mass
0.288 (slack)	0	0.697	—	—	0
0.274	2	0.684	−0.01	−0.01	0
0.266	20	0.700	+0.12	−0.02	0.0052

Table 1: The novice cap instrument. Rank shifts are the mean change in assigned school, in rungs of the fifty-school ladder, for agents with $a \leq 10$ and $a > 30$, obtained by coupling the constrained and capacity-only assignments monotonically at each age. Every point is certified in the sense of Appendix F.6.

Implementation under DA Whichever instrument the designer settles on, the allocation is implementable under agent-proposing DA. We compute below the bonuses for the $\kappa = 0.274$ constrained economy described above. The same computation applies to the equilibria obtained for the ceiling on average age. The market-clearing vector of cutoffs is fixed to the zero vector and agents are endowed with an initial score vector $(r_s^{(1),b})_s$ that depends on their budget under the pseudo-market mechanism and evolves via additive bonuses

$$r_{s_a}^{(a),b} = r_{s_a}^{(a-1),b} + \beta_{s_a}(s_{a-1}, a-1), \quad \forall s$$

where s_{a-1} is the school the agent was assigned to at $a-1$ and s_a the school whose score is being updated. The method used to derive the bonuses from prices is detailed in Appendix F.1.3.

Two properties of the bonuses are worth emphasizing. First, they are generation specific. A seat at a novice-heavy school costs priority points to a novice and nothing to an older agent. Moreover, at $a_0 + 1$, the scores for the novice-heavy schools are credited with the surcharge the agent will no longer incur at these schools. Second, bonuses depend on the school attended in the previous period. A novice who attends a novice-heavy school incurs a penalty of 0.01 priority points while attending another school leaves the score unaffected. Table 5 shows the initial scores and relevant bonuses.

Two lessons for design follow from our analysis. First, coarse constraints targeting uniformly a block of agents can backfire by creating exclusion in the market. Second, when priority scores are built directly to address policy concerns, they often feature a suboptimal uniform structure. Our construction and quantitative exercise show that bonuses implementing a constrained efficient allocation are school and generation specific and depend on the shape of the constraints.

Initial score $r_1^s = B + b_1(s)$						
school s	0	1	2	4	25	49
$b_1(s)$	0	−0.0100	−0.0068	−0.0231	−1.135	−9.042

update $b_a(s_{a-1}, s)$	$2 \le a \le a_0$			$a = a_0 + 1$			$a > a_0 + 1$		
attended s_{a-1} \ scored s	1	2	else	1	2	else	1	2	else
0	0	0	0	+0.0100	+0.0068	0	0	0	0
1	−0.0100	−0.0100	−0.0100	0	−0.0032	−0.0100	0	0	0
2	−0.0068	−0.0068	−0.0068	+0.0032	0	−0.0068	0	0	0
4	−0.0231	−0.0231	−0.0231	−0.0131	−0.0163	−0.0231	−0.0231	−0.0231	−0.0231
25	−1.135	−1.135	−1.135	−1.125	−1.128	−1.135	−1.135	−1.135	−1.135
49	−9.042	−9.042	−9.042	−9.032	−9.035	−9.042	−9.042	−9.042	−9.042

Figure 5: Scores and bonuses at $\kappa = 0.274$. Cutoffs are zero at every school. An agent enters with score $r_1^s = B + b_1(s)$ and updates it by $b_a(s_{a-1}, s)$, where s_{a-1} is the school she attended at $a - 1$; all entries are in percent of the mean lifetime budget. Schools 1 and 2 are capped, school 0 is free and uncapped, schools 4 and above carry a capacity price; rows 25 and 49 are shown for scale. Cells are shaded on the palette of Figure 2: red for a positive update, blue for a negative one, graded by size. Updates shown in bold include the graduation credit; note that a credited update can still be a net debit, when the price of the seat just occupied exceeds the credit. Bottom: the score carried for the capped school over the career, for two agents entering with the same budget, one spending her novice years at the capped school and one at the free school beside it. Scores are in the same units as budgets.

References

Abdulkadiroğlu, Atila, Yeon-Koo Che, and Yosuke Yasuda, "Resolving Conflicting Preferences in School Choice: The Boston Mechanism Reconsidered," *American Economic Review*, February 2011, *101* (1), 399–410.

Artemov, Georgy, Yeon-Koo Che, and YingHua He, "Stable matching with mistaken agents," *Journal of Political Economy Microeconomics*, 2023, *1* (2), 270–320.

Ashlagi, Itai and Peng Shi, "Optimal Allocation Without Money: An Engineering Approach," *Management Science*, 2016, *62* (4), 1078–1097.

Azevedo, Eduardo M. and Jacob D Leshno, "A supply and demand framework for two-sided matching markets," *Journal of Political Economy*, 2016, *124* (5), 1235–1268.

Bellman, Richard, *Dynamic Programming*, Princeton University Press, 1957.

Bergemann, Dirk and Juuso Välimäki, "Dynamic Mechanism Design: An Introduction," *Journal of Economic Literature*, 2019, *57* (2), 235–274.

Bobba, Matteo, Tim Ederer, Gianmarco Leon-Ciliotta, Christopher Neilson, and Marco G. Nieddu, "Teacher Compensation and Structural Inequality: Evidence from Centralized Teacher School Choice in Peru," 2021. Working Paper.

Carmona, Guilherme and Konrad Podczeck, "Approximation and characterization of Nash equilibria of large games," *Economic Theory*, 2022, *73* (2), 679–694.

Che, Yeon-Koo, Dong Woo Hahm, and YingHua He, "Leveraging Uncertainties to Infer Preferences: Robust Analysis of School Choice," Technical Report, SSRN Electronic Journal September 2023. Available at SSRN: `https://ssrn.com/abstract=4580093`.

Combe, Julien, Olivier Tercieux, and Camille Terrier, "The design of teacher assignment: Theory and evidence," *The Review of Economic Studies*, 2022, *89* (6), 3154–3222.

Corblet, Pauline, Jeremy T. Fox, and Alfred Galichon, "Repeated Matching Games: An Empirical Framework," NBER Working Paper 34352, National Bureau of Economic Research 2025.

Cox, Natalie, Ricardo Fonseca, Bobak Pakzad-Hurson, and Matthew Pecenco, "Peer Preferences in Centralized School Choice Markets: Theory and Evidence," 2025. Working paper.

Davis, Jonathan M. V., Kyle Greenberg, and Damon Jones, "An Experimental Evaluation of Deferred Acceptance: Evidence From Over 100 Army Officer Labor Markets," *Econometrica*, 2026, *94* (2), 641–662.

Doval, Laura, "Dynamically Stable Matching," *Theoretical Economics*, 2022, *17*, 687–724.

_ , "Dynamic Matching," in Yeon-Koo Che, Pierre-André Chiappori, and Bernard Salanié, eds., *Handbook of the Economics of Matching*, Vol. 2, Elsevier, 2025, chapter 6, pp. 489–577.

Echenique, Federico, Antonio Miralles, and Jun Zhang, "Constrained Pseudo-Market Equilibrium," *American Economic Review*, November 2021, *111* (11), 3699–3732.

Ederer, Tim, "An Empirical Model of Dynamic Two-Sided Matching: Application to the Teacher Labor Market," 2026. Working paper.

Eisenhauer, Philipp, "The Approximate Solution of Finite-Horizon Discrete-Choice Dynamic Programming Models," *Journal of Applied Econometrics*, 2019, *34* (1), 149–154.

Elacqua, Gregory and Leonardo Rosa, "Teacher Transfers and the Disruption of Teacher Staffing in the City of São Paulo," *International Journal of Educational Development*, 2025, *117*, 103353.

Fack, Gabrielle, Julien Grenet, and Yinghua He, "Beyond Truth-Telling: Preference Estimation with Centralized School Choice and College Admissions," *American Economic Review*, 2019, *109* (4), 1486–1529.

Government of Spain, "Royal Decree 1364/2010 Regulating Nationwide Transfer Competitions for Public-School Teachers," Boletín Oficial del Estado 2010.

Guillen, Pablo and Onur Kesten, "Matching Markets with Mixed Ownership: The Case for a Real-life Assignment Mechanism," *International Economic Review*, 2012, *53* (3), 1027–1046.

Hafalir, Isa E., Fuhito Kojima, and M. Bumin Yenmez, "Market Design with Distributional Objectives: Efficiency, Incentives, and Property Rights," *Journal of Economic Theory*, December 2025, *230*, 106099.

Hahm, Dong Woo and Minseon Park, "A Dynamic Framework of School Choice: Effects of Middle Schools on High School Choice," 2024. Working paper, November 2024; available at SSRN 3996418.

Hanushek, Eric A., John F. Kain, and Steven G. Rivkin, "Why Public Schools Lose Teachers," *The Journal of Human Resources*, 2004, *39* (2), 326–354.

Jackson, Kirabo C., "Student Demographics, Teacher Sorting, and Teacher Quality: Evidence from the End of School Desegregation," *Journal of Labor Economics*, 2009, *27* (2), 213–256.

Kamada, Yuichiro and Fuhito Kojima, "Efficient Matching under Distributional Constraints: Theory and Applications," *American Economic Review*, January 2015, *105* (1), 67–99.

Keane, Michael P. and Kenneth I. Wolpin, "The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence," *The Review of Economics and Statistics*, 1994, *76* (4), 648–672.

Miralles, Antonio and Marek Pycia, "Prices and Efficient Assignments Without Transfers," 2014. Working Paper.

New South Wales Department of Education, "Benefits and Incentives for Teaching in Rural and Remote Schools," Teach.NSW webpage 2026.

_ , "Transfer System: Appendix 2 of the Staffing Procedure for the Teaching Service in NSW Public Schools," Departmental policy webpage 2026.

Nicolò, Antonio, Pietro Salmaso, and Riccardo D. Saulle, "Dynamic One-Sided Matching," 2024. Working paper.

Republic of Türkiye Ministry of National Education, "Announcement of the 2026 Interprovincial Voluntary Teacher-Transfer Process," Ministry announcement, May 11, 2026 2026. Teachers could rank up to 40 schools; assignments were based on preferences, vacancies, and service points.

_ , "Regulation on the Appointment and Transfer of Teachers Affiliated with the Ministry of National Education," Official regulation announced January 9, 2026 2026.

Rust, John, "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 1987, *55* (5), 999–1033.

_ , "Using Randomization to Break the Curse of Dimensionality," *Econometrica*, 1997, *65* (3), 487–516.

Shi, Peng, "Optimal Priority-Based Allocation Mechanisms," *Management Science*, January 2022, *68* (1), 171–188.

Sidibe, Modibo, Bocar Ba, Patrick Bayer, Nayoung Rim, and Roman Rivera, "Police Officer Assignment and Neighborhood Crime," 2021. Working Paper.

Sun, Yeneng, "The Exact Law of Large Numbers via Fubini Extension and Characterization of Insurable Risks," *Journal of Economic Theory*, January 2006, *126* (1), 31–69.

Suzuki, Takamasa, Akihisa Tamura, and Makoto Yokoo, "Efficient allocation mechanism with endowments and distributional constraints," *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems*, 2018, *July 10-15, Stockholm, Sweden.*

Thakur, Ashutosh, "Matching in the Civil Service: A Market Design Approach to Public Administration and Development," ECONtribute Discussion Paper 087, University of Bonn and University of Cologne 2021.

Young, Eric R., "Solving the Incomplete Markets Model with Aggregate Uncertainty Using the Krusell-Smith Algorithm and Non-Stochastic Simulations," *Journal of Economic Dynamics and Control*, 2010, *34* (1), 36–41.

Appendix

A Distributional Implications of the Maintained Assumptions

Let λ_Θ denote Lebesgue measure on Θ. For any cutoff path $\mathbf{P}$, define the dynamic indifference set

$$\mathcal{I}_{t,a}(\mathbf{P}) := \left\{ \theta = (r,u) \in \Theta : \left| \arg\max_{s \in \mathcal{F}_t(r;\mathbf{P})} W^s_{t,a}(\theta \mid \mathbf{P}) \right| > 1 \right\},$$

and the corresponding exceptional set

$$\mathcal{B}_{t,a}(\mathbf{P}) := \bigcup_{j \in \mathcal{S}} \{\theta = (r,u) : r^j = P^j_t\} \ \cup \ \mathcal{I}_{t,a}(\mathbf{P}).$$

Lemma A.1 (Distributional richness and null boundaries)**.** *Suppose that Assumptions 2.1 and 2.2 hold. Let $(\eta_{t,a})_{t,a}$ be any family of age-specific distributions generated from p and the initial condition by measurable probability kernel*[27]

$$\pi_{t,a}(\cdot \mid \theta) \in \Delta(\mathcal{S}_0) \text{ for each } t \text{ and } a.$$

Then:

(i) Every $\eta_{t,a}$ is absolutely continuous with respect to λ_Θ and has full support on Θ.

(ii) For every school $j \in \mathcal{S}$ and $x \in [-1,1]$,

$$\eta_{t,a}\big(\{\theta = (r,u) : r^j = x\}\big) = 0.$$

Moreover, for every cutoff path $\mathbf{P}$,

$$\eta_{t,a}\big(\mathcal{I}_{t,a}(\mathbf{P})\big) = 0, \qquad \eta_{t,a}\big(\mathcal{B}_{t,a}(\mathbf{P})\big) = 0.$$

(iii) For every $a < A$, current type $\theta \in \Theta$, assignment $s \in \mathcal{S}_0$, cutoff path $\mathbf{P}$, and date t,

$$Q_a\big(\mathcal{B}_{t+1,a+1}(\mathbf{P}) \mid \theta, s\big) = 0.$$

Proof. The age-one distribution is p, and the distributions of cohorts already present at date zero satisfy part (i) by Assumption 2.1. Suppose that a current distribution η is followed by a measurable assignment kernel $\pi(\cdot \mid \theta) \in \Delta(\mathcal{S}_0)$. The next-age distribution is

$$\eta'(X) = \int_\Theta \sum_{s \in \mathcal{S}_0} Q_a(X \mid \theta, s)\, \pi(s \mid \theta)\, \eta(d\theta).$$

If N is λ_Θ-null, absolute continuity of every conditional transition law gives $Q_a(N \mid \theta, s) = 0$ for every θ and s, and hence $\eta'(N) = 0$. If $O \subseteq \Theta$ is nonempty and open, full support gives

[27] A pre-matching $\mu = (\mu_{t,a})_{t,a}$—and hence, in particular, any matching—is identified with the degenerate kernels

$$\pi_{t,a}(\cdot \mid \theta) = \delta_{\mu_{t,a}(\theta)}.$$

We state the lemma for general probability kernels because mixed reporting strategies may induce randomized assignments conditional on type. The formulation therefore covers both the deterministic assignments generated by cutoff paths and the randomized assignments generated by mixed strategies.

$Q_a(O \mid \theta, s) > 0$ for every θ and s. The integrand above is therefore strictly positive for every θ, so $\eta'(O) > 0$. Part (i) follows recursively.

For part (ii), every score boundary $\{\theta : r^j = x\}$ is a λ_Θ-null hyperplane, while $\mathcal{I}_{t,a}(\mathbf{P})$ is λ_Θ-null by Assumption 2.2. Part (i) therefore implies the stated null-set properties, and the same is true of their finite union $\mathcal{B}_{t,a}(\mathbf{P})$. Part (iii) follows because every conditional transition law is absolutely continuous with respect to λ_Θ. □

Remark A.1. *The proofs below use only the null-set conclusions in parts (ii)–(iii), except for the selector-independent characterization, which additionally uses the full-support conclusion in part (i). Thus, these reduced-form properties may be imposed directly in place of the stronger Assumptions 2.1 and 2.2; in particular, singular transition laws are admissible for virtually all our results.*

B Proof of Lemma 2.1

We separate the argument into three steps. Since $\mathcal{U} = [0,1]^S$ and $u^\emptyset = 0$, every flow payoff is bounded in absolute value by $\bar{u} := 1$. Since an age-a agent receives at most $A - a + 1$ additional flow payoffs, backward induction gives

$$\sup_{t \in \mathbb{N}_0} \sup_{\mathbf{P}} \sup_{\theta \in \Theta} |V_{t,a}(\theta \mid \mathbf{P})| \leq (A - a + 1)\bar{u}, \qquad \sup_{P \in [0,1]^S} \sup_{\theta \in \Theta} |V_a(\theta \mid P)| \leq (A - a + 1)\bar{u}. \tag{36}$$

Thus, the value functions are bounded uniformly over calendar dates, types, and cutoff paths and, in the stationary problem, uniformly over types and stationary cutoff vectors.

Step 1: convergence of individual values and demand. For each age a, let $\mathcal{I}_a(P)$ be the set of types with more than one maximizer in the stationary problem (8), and define

$$\mathcal{B}_a(P) := \bigcup_{s \in \mathcal{S}} \{\theta = (r, u) : r^s = P^s\} \ \cup \ \mathcal{I}_a(P).$$

Consider the constant cutoff path whose value is P at every date. For any cohort that enters after date zero, its age-a distribution under this path is η_a^P. Part (ii) of Lemma A.1, applied to this constant path, implies

$$\eta_a^P\big(\mathcal{B}_a(P)\big) = 0 \qquad \text{for every } a. \tag{37}$$

Note that, by definition of $\mathcal{B}_a(P)$: if $\theta \notin \mathcal{B}_a(P)$, then no score of θ is equal to the cutoff and the stationary maximizer is unique.

We prove by backward induction on age the following statement. If $t_n \to \infty$ and $\theta_n \to \theta \notin \mathcal{B}_a(P)$, then

$$V_{t_n,a}(\theta_n \mid \mathbf{P}) \longrightarrow V_a(\theta \mid P), \qquad d_{t_n,a}(\theta_n \mid \mathbf{P}) = d_a(\theta \mid P) \quad \text{for all sufficiently large } n. \tag{38}$$

At age A, the definition of $\mathcal{B}_A(P)$ implies that every score is strictly on one side of the corresponding cutoff. Since $P_{t_n} \to P$ and $\theta_n \to \theta$, the feasible set $\mathcal{F}_{t_n}(r_n; \mathbf{P})$ is eventually equal to $\mathcal{F}(r; P)$ where r_n (r) is the score vector associated with θ_n (θ). Further, the maximizer over this finite set is unique. Hence the payoff gap between the best and the second-best feasible assignments is strictly positive. Because the flow payoffs in θ_n converge to those in θ, the maximizing assignment is unchanged for all sufficiently large n, and its value converges. This proves (38) at age A.

Suppose the statement holds at age $a + 1$. We first hold the current type θ and assignment s fixed. Part (iii) of Lemma A.1, applied to the constant cutoff path P, gives

$$Q_a\big(\mathcal{B}_{a+1}(P) \mid \theta, s\big) = 0. \tag{39}$$

For every $\theta' \notin \mathcal{B}_{a+1}(P)$, the induction hypothesis, applied to the constant sequence $\theta'_n = \theta'$, yields

$$V_{t+1,a+1}(\theta' \mid \mathbf{P}) \longrightarrow V_{a+1}(\theta' \mid P).$$

Equation (39) gives pointwise convergence $Q_a(\cdot \mid \theta, s)$-almost everywhere, and the dominated convergence theorem yields

$$\int_\Theta V_{t_n+1,a+1}(\theta' \mid \mathbf{P}) Q_a(d\theta' \mid \theta, s) \longrightarrow \int_\Theta V_{a+1}(\theta' \mid P) Q_a(d\theta' \mid \theta, s) \tag{40}$$

given that the functions inside the integrals are bounded in absolute value by the common constant $(A-a)\bar{u}$ in (36).

We also need convergence when the current type varies. Let $\theta_n \to \theta$. By the continuity of the transition kernel in the weak topology,

$$Q_a(\cdot \mid \theta_n, s) \Rightarrow Q_a(\cdot \mid \theta, s)$$

where we write $\Rightarrow$ for weak convergence. We use the following standard extension of the Portmanteau theorem: if $\nu_n \Rightarrow \nu$, the functions g_n are uniformly bounded, and $g_n(x_n) \to g(x)$ whenever $x_n \to x$ outside a ν-null set, then $\int g_n \, d\nu_n \to \int g \, d\nu$. Here, the induction hypothesis supplies this continuous convergence outside $\mathcal{B}_{a+1}(P)$, the uniform bound comes from (36), and (39) makes the exceptional set null under the limiting transition law. Consequently,

$$\int_\Theta V_{t_n+1,a+1}(\theta' \mid \mathbf{P}) Q_a(d\theta' \mid \theta_n, s) \longrightarrow \int_\Theta V_{a+1}(\theta' \mid P) Q_a(d\theta' \mid \theta, s). \tag{41}$$

Adding the convergent current flow payoff shows that the objective associated with every assignment s converges to its stationary counterpart. If $\theta \notin \mathcal{B}_a(P)$, its feasible set is eventually unchanged and its stationary maximizer is unique. Since $\mathcal{S}_0$ is finite, the strict gap between the best and second-best stationary objectives is preserved for all sufficiently large n. This proves (38) at age a and completes the backward induction.

Step 2: convergence of age-specific type distributions and aggregate demand. For any fixed age a and all sufficiently large dates t, the age-a cohort entered after date zero. Its age-one distribution is therefore p; the arbitrary initial distributions of cohorts already present at date zero no longer matter. We show by induction on age that

$$\eta^{\mathbf{P}}_{t,a} \Rightarrow \eta^P_a. \tag{42}$$

The claim is immediate at age one because $\eta^{\mathbf{P}}_{t,1} = p = \eta^P_1$ for every t.

Suppose (42) holds at age a, and let $f : \Theta \to \mathbb{R}$ be bounded and continuous. For each assignment s, define

$$K^s_a f(\theta) := \int_\Theta f(\theta') Q_a(d\theta' \mid \theta, s).$$

The weak continuity of $\theta \mapsto Q_a(\cdot \mid \theta, s)$ implies that $K^s_a f$ is continuous. In addition, by definition of the law of motion, we must have

$$\int_\Theta f(\theta') \eta^{\mathbf{P}}_{t+1,a+1}(d\theta') = \sum_{s \in \mathcal{S}_0} \int_\Theta \mathbf{1}\{d_{t,a}(\theta \mid \mathbf{P}) = s\} K^s_a f(\theta) \eta^{\mathbf{P}}_{t,a}(d\theta),$$

$$\int_\Theta f(\theta') \eta^P_{a+1}(d\theta') = \sum_{s \in \mathcal{S}_0} \int_\Theta \mathbf{1}\{d_a(\theta \mid P) = s\} K^s_a f(\theta) \eta^P_a(d\theta).$$

By (38), for every sequence $t_n \to \infty$ and $\theta_n \to \theta \notin \mathcal{B}_a(P)$, the first integrand evaluated at θ_n converges to the second integrand evaluated at θ. The functions are uniformly bounded, and (37) makes the exceptional set null under the limiting measure η_a^P. Applying the same extension of the Portmanteau theorem, together with the induction hypothesis $\eta_{t,a}^{\mathbf{P}} \Rightarrow \eta_a^P$, gives convergence of the two integrals. Since this holds for every bounded continuous f, (42) follows at age $a+1$.

The same argument, now applied to the demand indicators, yields

$$D_{t,a}^s(\mathbf{P}) = \frac{1}{A}\int_\Theta \mathbf{1}\{d_{t,a}(\theta \mid \mathbf{P}) = s\}\eta_{t,a}^{\mathbf{P}}(d\theta) \longrightarrow \frac{1}{A}\int_\Theta \mathbf{1}\{d_a(\theta \mid P) = s\}\eta_a^P(d\theta) = D_a^s(P). \tag{43}$$

Summing (43) over the finite set of ages gives

$$D_t^s(\mathbf{P}) \longrightarrow D^s(P) \qquad \text{for every school } s. \tag{44}$$

Step 3: passage to the market-clearing conditions. Since $\mathbf{P}$ is market clearing, $D_t^s(\mathbf{P}) \le q^s$ for every t and s. Taking limits in (44) gives $D^s(P) \le q^s$. If $P^s > 0$, then $P_t^s > 0$ for all sufficiently large t. Market clearing along the path then gives $D_t^s(\mathbf{P}) = q^s$ eventually, and (44) implies $D^s(P) = q^s$. Hence P satisfies the stationary market-clearing conditions.

C Proofs for Section 3

C.1 Stable Cutoff Selectors and the General Reporting Game

The main text focuses on the agent-proposing DA mechanism. We now formulate the reporting game for an arbitrary stable mechanism, allowing the static market at a given date to have several market-clearing cutoff vectors. Recall that $\mathcal{C}(\widehat{D})$ denotes the set of static market-clearing cutoffs associated with a reported-demand function $\widehat{D}$. A *stable cutoff selector* is a map Φ defined on the relevant class of reported-demand functions such that

$$\Phi(\widehat{D}) \in \mathcal{C}(\widehat{D}) \qquad \text{for every } \widehat{D} \text{ in its domain.}$$

The agent-proposing DA rule used in the main text is the smallest stable selector,

$$\Phi^-(\widehat{D}) := \min \mathcal{C}(\widehat{D}) = \mathrm{DA}(\widehat{D}).$$

Other stable mechanisms correspond to other selections from the same complete lattice.

Fix an initial condition and a sequence of stable cutoff selectors $\mathbf{\Phi} = (\Phi_t)_{t\in\mathbb{N}_0}$. Given a reporting strategy profile σ, set $\eta_{t,1}^\sigma = p$ for every t and define recursively

$$\widehat{D}_t^\sigma(\cdot) := \widehat{D}(\cdot \mid \sigma_t, \eta_t^\sigma),$$
$$P_t^\sigma := \Phi_t\big(\widehat{D}_t^\sigma\big),$$

with the remaining age-specific distributions generated by

$$\eta_{t+1,a+1}^\sigma(X) = \int_\Theta \int_{\mathcal{L}} Q_a\Big(X \mid \theta, \widehat{d}(P_t^\sigma \mid r, \ell)\Big)\, \sigma_{t,a}(d\ell \mid \theta)\, \eta_{t,a}^\sigma(d\theta).$$

The notions of equilibrium and stationary equilibrium introduced here are the natural counterparts of Definitions 3.1 and 3.2 in the main text. We restate them here to make the appendix self-contained.

Definition C.1 (Equilibrium under a stable cutoff selector)**.** *A reporting strategy profile σ is an equilibrium of the game induced by $\mathbf{\Phi}$ if, for every date t, age a, and type $\theta = (r, u)$,*

$$\widehat{d}(P_t^\sigma \mid r, \ell) \in \underset{s \in \mathcal{F}_t(r; \mathbf{P}^\sigma)}{\arg\max} W_{t,a}^s(\theta \mid \mathbf{P}^\sigma) \quad \textit{for } \sigma_{t,a}(\cdot \mid \theta)\textit{-almost every } \ell.$$

Definition C.2 (Stationary equilibrium under a stable cutoff selector)**.** *Fix a time-invariant stable cutoff selector Φ. A time-invariant reporting strategy $\sigma = (\sigma_a)_{a \in \mathcal{A}}$ is a stationary equilibrium of the game induced by Φ if there exist initial age-specific type distributions $(\eta_{0,a})_{a=2}^A$ such that:*

1. *when σ_a is used at every date, the induced cutoff path and age-specific type distributions are time invariant; and*
2. *the resulting strategy profile is an equilibrium under these initial distributions in the sense of Definition C.1.*

C.2 Equilibrium Characterization

For the next proposition and its corollary, and to make clear that we do not need the whole strength of Assumptions 2.1 and 2.2, we simply assume that part (ii) of Lemma A.1 holds.

Proposition C.1 (Equilibrium and market-clearing cutoff paths under stable selectors)**.** *Fix an initial condition.*

(i) For any sequence of stable cutoff selectors $\mathbf{\Phi}$, every equilibrium-induced cutoff path is market clearing.

(ii) Every market-clearing cutoff path can be induced by an equilibrium under some sequence of stable cutoff selectors.

Proof. For part (i), let σ be an equilibrium under $\mathbf{\Phi}$. By Definition C.1, every report used by an agent induces a dynamically optimal feasible assignment under the cutoff path $\mathbf{P}^\sigma$. By part (ii) of Lemma A.1, this assignment is unique for $\eta_{t,a}^\sigma$-almost every type. Consequently,

$$\widehat{d}(P_t^\sigma \mid r, \ell) = d_{t,a}(\theta \mid \mathbf{P}^\sigma)$$

for $\eta_{t,a}^\sigma$-almost every θ and every report ℓ used with positive probability.

A simple induction argument on dates shows that

$$\eta_{t,a}^\sigma = \eta_{t,a}^{\mathbf{P}^\sigma} \qquad \text{for every } t \text{ and } a.$$

Indeed, at date zero, the two systems have the same initial distributions, and the distribution of age-one agents is p at every date. Suppose the distributions coincide at date t. Under this common distribution, the preceding equality shows that the equilibrium assignment coincides almost everywhere with $d_{t,a}(\cdot \mid \mathbf{P}^\sigma)$. Substituting the same assignment rule into the two laws of motion yields

$$\eta_{t+1,a+1}^\sigma = \eta_{t+1,a+1}^{\mathbf{P}^\sigma} \qquad \text{for every } a < A,$$

which completes the induction.

Reported and dynamic aggregate demand therefore coincide at the selected cutoff:

$$\left(\widehat{D}_t^\sigma\right)^s(P_t^\sigma) = D_t^s(\mathbf{P}^\sigma) \qquad \text{for every school } s.$$

Because $P_t^\sigma = \Phi_t(\widehat{D}_t^\sigma)$ belongs to $\mathcal{C}(\widehat{D}_t^\sigma)$, it clears reported demand. Hence

$$D_t^s(\mathbf{P}^\sigma) \le q^s, \qquad D_t^s(\mathbf{P}^\sigma) = q^s \quad \text{whenever } (P_t^\sigma)^s > 0.$$

Thus $\mathbf{P}^\sigma$ is a market-clearing cutoff path.

For part (ii), let $\mathbf{P}$ be a market-clearing cutoff path. For each t, a, and θ, choose a report $\ell_{t,a}^{\mathbf{P}}(\theta)$ that ranks assignments according to the choice-specific values $W_{t,a}^s(\theta \mid \mathbf{P})$, resolving ties with a fixed ordering of $\mathcal{S}_0$. By part (ii) of Lemma A.1,

$$\widehat{d}\big(P_t \mid r, \ell_{t,a}^{\mathbf{P}}(\theta)\big) = d_{t,a}(\theta \mid \mathbf{P})$$

for $\eta_{t,a}^{\mathbf{P}}$-almost every type. Let $\sigma_{t,a}^{\mathbf{P}}(\cdot \mid \theta)$ be degenerate at $\ell_{t,a}^{\mathbf{P}}(\theta)$ and define

$$\widehat{D}_t^{\mathbf{P}}(\cdot) := \widehat{D}\big(\cdot \mid \sigma_t^{\mathbf{P}}, \eta_t^{\mathbf{P}}\big).$$

At P_t, reported demand coincides with dynamic demand:

$$\big(\widehat{D}_t^{\mathbf{P}}\big)^s(P_t) = D_t^s(\mathbf{P}) \qquad \text{for every } s.$$

Since $\mathbf{P}$ is market clearing, $P_t \in \mathcal{C}(\widehat{D}_t^{\mathbf{P}})$. For each date t, choose a stable cutoff selector Φ_t satisfying

$$\Phi_t\big(\widehat{D}_t^{\mathbf{P}}\big) = P_t.$$

An induction on dates then shows that $\sigma^{\mathbf{P}}$ generates the distributions $\eta_t^{\mathbf{P}}$, the cutoff path $\mathbf{P}$, and the assignments $d_{t,a}(\cdot \mid \mathbf{P})$. Because its reports rank current assignments by their dynamic choice-specific values, $\sigma^{\mathbf{P}}$ is an equilibrium under $\mathbf{\Phi}$ and induces $\mathbf{P}$. □

Corollary C.1 (Stationary equilibrium characterization under stable selectors)**.** *A vector P is a stationary market-clearing cutoff vector if and only if it is induced by a stationary equilibrium under some time-invariant stable cutoff selector.*

Proof. Apply Proposition C.1 to the constant cutoff path. If P is stationary market clearing, the dynamic rankings, reported-demand function, and age-specific distributions are the same at every date, so the same selector can be chosen to select P in every period. The reports constructed in part (ii) rank assignments by their stationary dynamic choice-specific values, so the resulting stationary equilibrium is truthful. The converse follows from part (i) of the proposition. □

C.3 Selector-independent Characterization

The preceding result allows the selector to depend on the market-clearing path to be implemented. The next result shows that, under the maintained richness conditions, this dependence disappears.

Proposition C.2 (Selector independence)**.** *Suppose that Assumptions 2.1 and 2.2 hold. Fix any sequence of stable cutoff selectors*

$$\mathbf{\Phi} = (\Phi_t)_{t \in \mathbb{N}_0}.$$

The cutoff paths induced by equilibria under $\mathbf{\Phi}$ are exactly the market-clearing cutoff paths.

Proof. By part (i) of Proposition C.1, every cutoff path induced by an equilibrium under $\mathbf{\Phi}$ is market clearing.

Conversely, let $\mathbf{P}$ be a market-clearing cutoff path and construct the pure reporting strategy $\sigma^{\mathbf{P}}$ as in the proof of Proposition C.1. Let

$$\widehat{D}_t^{\mathbf{P}}(\cdot) := \widehat{D}\big(\cdot \mid \sigma_t^{\mathbf{P}}, \eta_t^{\mathbf{P}}\big)$$

be the corresponding reported-demand function. We show that $\widehat{D}_t^{\mathbf{P}}$ has a unique static market-clearing cutoff.

Part (i) of Lemma A.1 implies that $\eta_{t,a}^{\mathbf{P}}$ has full support on Θ for every t and a. In particular,

$$\operatorname{supp} \eta_{t,A}^{\mathbf{P}} = \Theta.$$

At age A, choice-specific values coincide with flow payoffs:

$$W_{t,A}^{s}(\theta \mid \mathbf{P}) = u^{s}.$$

For every school $s \in \mathcal{S}$, let

$$\mathcal{U}_s^{\text{top}} := \left\{ u \in [0,1]^S : u^s > u^j \text{ for every } j \neq s, \text{ and } u^s > 0 \right\}.$$

This is a nonempty open subset of $\mathcal{U}$. Hence, for every school s and every nonempty open score set $O \subseteq \mathcal{R}$,

$$\eta_{t,A}^{\mathbf{P}}\big(O \times \mathcal{U}_s^{\text{top}}\big) > 0.$$

Thus, for each school and each nonempty open score set, a positive mass of age-A agents reports that school as its first choice and has a score vector in that set.

This top-choice richness implies strict gross substitutes for $\widehat{D}_t^{\mathbf{P}}$.[28] If $R \leq R'$, $R \neq R'$, and

$$J := \{s \in \mathcal{S} : R^s < R'^s\},$$

then

$$\sum_{s \in J} \big(\widehat{D}_t^{\mathbf{P}}\big)^s(R) > \sum_{s \in J} \big(\widehat{D}_t^{\mathbf{P}}\big)^s(R').$$

Indeed, choose $s \in J$. Full support gives a positive mass of agents who rank s first and whose scores satisfy

$$R^s < r^s < R'^s, \qquad r^j < R'^j \quad \text{for every } j \in J \setminus \{s\}.$$

These agents demand a school in J, namely s, at R, but no school in J at R'. Moreover, raising only the cutoffs of schools in J cannot cause any other agent to switch into J.

Because $\mathcal{C}(\widehat{D}_t^{\mathbf{P}})$ is a complete lattice, strict gross substitutes implies that it is a singleton. To see this, suppose its smallest and largest elements R^- and R^+ were distinct and let

$$J := \{s \in \mathcal{S} : (R^-)^s < (R^+)^s\}.$$

For every $s \in J$, $(R^+)^s > 0$, and market clearing gives

$$\big(\widehat{D}_t^{\mathbf{P}}\big)^s(R^+) = q^s, \qquad \big(\widehat{D}_t^{\mathbf{P}}\big)^s(R^-) \leq q^s.$$

Summing over J contradicts strict gross substitutes.

At P_t, reported and dynamic demand coincide, and $\mathbf{P}$ is market clearing. Therefore

$$P_t \in \mathcal{C}(\widehat{D}_t^{\mathbf{P}}).$$

Uniqueness implies

$$\mathcal{C}(\widehat{D}_t^{\mathbf{P}}) = \{P_t\},$$

so every stable cutoff selector satisfies

$$\Phi_t(\widehat{D}_t^{\mathbf{P}}) = P_t.$$

An induction on dates now shows that $\sigma^{\mathbf{P}}$ generates $\mathbf{P}$, the distributions $(\eta_{t,a}^{\mathbf{P}})_{t,a}$, and the assignments $d_{t,a}(\cdot \mid \mathbf{P})$. The strategy is therefore an equilibrium under the prescribed sequence $\mathbf{\Phi}$. □

[28] This condition was introduced by Che, Hahm, and He (2023), who show that it guarantees uniqueness of the market-clearing cutoff in a static environment.

Corollary C.2 (Selector-independent stationary characterization)**.** *Suppose that Assumptions 2.1 and 2.2 hold. Fix any time-invariant stable cutoff selector Φ. A vector P is a stationary market-clearing cutoff vector if and only if it is induced by a stationary equilibrium under Φ.*

Proof. Suppose first that a stationary equilibrium under Φ induces the constant cutoff vector P. Corollary C.1 implies that P is stationary market clearing.

Conversely, let P be stationary market clearing and initialize the economy at

$$\eta_{0,a} = \eta_a^P, \qquad a = 2, \dots, A.$$

Part (i) of Lemma A.1 implies that every η_a^P has full support. Apply Proposition C.2 to the constant cutoff path $(P, P, \dots)$. The cutoff path, dynamic rankings, and age-specific distributions are time invariant, so the constructed reporting strategy is time invariant. Since it ranks assignments by their stationary dynamic choice-specific values, it is truthful. It is therefore a truthful stationary equilibrium under Φ and induces P. □

Completion of the proof of Proposition 3.1. The map $\mathrm{DA} = \Phi^-$ is a stable cutoff selector. Part (1) of the proposition follows from Proposition C.2 applied to the constant sequence $\Phi_t = \Phi^-$. Part (2) follows from Corollary C.2 applied to Φ^-. □

C.4 Proof of Theorem 3.1

We prove parts (i), (ii) and (iii).

Part (i). We first verify that stationary aggregate demand is continuous. Let $P_n \to P$. For every age a, let $\mathcal{B}_a(P)$ be the union of the stationary score-boundary sets $\{\theta : r^s = P^s\}$ and the set of types with more than one stationary maximizer. By Lemma A.1, $\eta_a^P(\mathcal{B}_a(P)) = 0$, and every conditional transition law assigns zero mass to $\mathcal{B}_{a+1}(P)$. Backward induction on age, using boundedness of the value functions and weak continuity of the transition kernels, therefore implies that whenever $\theta_n \to \theta \notin \mathcal{B}_a(P)$,

$$V_a(\theta_n \mid P_n) \longrightarrow V_a(\theta \mid P), \qquad d_a(\theta_n \mid P_n) = d_a(\theta \mid P) \quad \text{eventually.}$$

An induction through the stationary law of motion then gives $\eta_a^{P_n} \Rightarrow \eta_a^P$ for every age. Applying the same continuous-convergence argument to the demand indicators yields $D_a^s(P_n) \to D_a^s(P)$ for every a and s. Hence $D : [0,1]^S \to [0,1]^S$ is continuous.

For each school s, let $z^s(P) := D^s(P) - q^s$ and define $F : [0,1]^S \to [0,1]^S$ by

$$F^s(P) := \min\bigl\{\max\{P^s + z^s(P), 0\}, 1\bigr\}.$$

The map F is continuous, so Brouwer's fixed-point theorem gives $P^* \in [0,1]^S$ such that $F(P^*) = P^*$.

Fix a school s. If $(P^*)^s = 0$, the fixed-point condition implies $z^s(P^*) \le 0$, and hence $D^s(P^*) \le q^s$. If $0 < (P^*)^s < 1$, then $z^s(P^*) = 0$. Finally, $(P^*)^s = 1$ is impossible: a fixed point at the upper boundary would require $z^s(P^*) \ge 0$, whereas $D^s(1, P^*_{-s}) = 0 < q^s$. Indeed, when $(P^*)^s = 1$, only agents with $r^s = 1$ can be assigned to school s, and part (ii) of Lemma A.1 assigns zero mass to this boundary. Thus,

$$D^s(P^*) \le q^s, \qquad D^s(P^*) = q^s \quad \text{whenever } (P^*)^s > 0.$$

Since this holds for every school, P^* is a stationary market-clearing cutoff vector.

Part (ii). Recall from (18) that

$$M(P) = \sum_{s \in \mathcal{S}} D^s(P)$$

is the total mass assigned under the stationary cutoff vector P. Suppose, toward a contradiction, that two distinct stationary market-clearing cutoff vectors P^- and P^+ are comparable. Without loss of generality,

$$P^- \leq P^+, \qquad P^- \neq P^+.$$

Under the nonnegative-score specification of part (ii), every component of both cutoff vectors is strictly positive. Indeed, if $P^s = 0$ for some school s, every agent meets that school's cutoff. Dynamic acceptability and the null-indifference conclusion of Lemma A.1 would then imply that a mass one is assigned to schools, contradicting $\sum_s q^s < 1$. Hence every school fills under both vectors. Therefore,

$$M(P^-) = \sum_{s \in \mathcal{S}} q^s = M(P^+).$$

But Assumption 3.1, applied to the distinct comparable pair $P^- \leq P^+$, implies

$$M(P^-) \neq M(P^+),$$

a contradiction. Hence no two distinct stationary market-clearing cutoff vectors are coordinatewise comparable, and the solution set is an anti-lattice.

Part (iii). We complete the proof of Theorem 3.1(iii). The benchmark economy presented in the main text is deliberately singular: payoff vectors are deterministic, and attending a school generates a deterministic priority bonus. It also has a positive mass of indifferent agents when $P^1 = 1/2$. We show that the benchmark can be approximated by an economy satisfying Assumptions 2.1 and 2.2, while preserving the two market-clearing cutoffs. We then show that this multiplicity persists under small perturbations of all the primitives.

Proof of Theorem 3.1(iii). Recall the benchmark payoff vectors

$$u^Y = \left(\frac{1}{4}, \frac{1}{2}\right), \qquad u^O = \left(1, \frac{1}{2}\right),$$

and capacities

$$q^1 = \frac{11}{32}, \qquad q^2 = \frac{3}{4}.$$

When the cutoff of school 2 is zero and $x := P^1$, benchmark stationary demand for school 1 is

$$D^1(x, 0) = \begin{cases} \dfrac{1-x}{2}, & x < \dfrac{1}{2}, \\ \dfrac{2-x-x^2}{2}, & x > \dfrac{1}{2}. \end{cases} \tag{45}$$

As shown in the main text,

$$D^1\left(\frac{5}{16}, 0\right) = D^1\left(\frac{3}{4}, 0\right) = \frac{11}{32}.$$

A regularized economy. Let $\bar{\lambda}_\Theta$ denote normalized Lebesgue measure on

$$\Theta = [-1,1]^2 \times [0,1]^2.$$

For every Borel set C of positive Lebesgue measure, let $\mathsf{U}(C)$ denote the uniform probability measure on C. For $v \in [0,1]^2$ and $\varepsilon > 0$, define

$$C_\varepsilon(v) := \left\{u \in [0,1]^2 : \|u - v\|_\infty \le \varepsilon\right\}.$$

The entry law in the regularized economy is

$$p_\varepsilon := (1-\varepsilon)\left[\mathsf{U}([0,1]^2) \otimes \mathsf{U}\big(C_\varepsilon(u^Y)\big)\right] + \varepsilon\bar{\lambda}_\Theta. \tag{46}$$

Thus, with probability $1-\varepsilon$, entry scores have the same law as in the benchmark and entry payoffs lie in an ε-neighborhood of u^Y. With probability ε, the entire type is drawn from a full-support distribution.

We next define the age-one transition kernel. Draw independently

$$X = (X^1, X^2) \sim \mathsf{U}([0,1]^2)$$

and

$$U_\varepsilon^O \sim \mathsf{U}\big(C_\varepsilon(u^O)\big).$$

Following the outside option, let

$$R^{(\emptyset)} := X.$$

Following assignment to school $s \in \{1,2\}$, define $R^{(s)}$ by

$$R^{(s),s} := 1 - \varepsilon + \varepsilon X^s, \qquad R^{(s),j} := X^j, \quad j \neq s. \tag{47}$$

Let $\overline{Q}_{\varepsilon,s}$ be the law of

$$\big(R^{(s)}, U_\varepsilon^O\big), \qquad s \in \mathcal{S}_0,$$

and set

$$Q_{1,\varepsilon}\big(\cdot \mid \theta, s\big) := (1-\varepsilon)\overline{Q}_{\varepsilon,s} + \varepsilon\bar{\lambda}_\Theta, \qquad (\theta, s) \in \Theta \times \mathcal{S}_0. \tag{48}$$

The transition law does not depend on the current type θ.

Both p_ε and every conditional transition law in (48) are absolutely continuous with respect to Lebesgue measure. Because each contains $\bar{\lambda}_\Theta$ with strictly positive weight, they also have full support on Θ. Moreover, the transition kernel is constant in θ and is therefore weakly continuous. Hence Assumption 2.1 holds.

The regularization also preserves dynamic acceptability. Couple the transition following school s and that following the outside option by using the same draws of the mixture component, X, and U_ε^O. On the concentrated component,

$$R^{(s),s} - R^{(\emptyset),s} = (1-\varepsilon)(1 - X^s) \ge 0,$$

while every other score coordinate and the payoff vector are unchanged. On the full-support component, use the same next type following all current assignments. Assignment to school s therefore weakly expands the next-period feasible set relative to the outside option, without changing next-period payoffs. Its continuation value is consequently weakly higher.

Generic no indifference. We next verify Assumption 2.2. Fix an arbitrary cutoff path $\mathbf{P}$. At age two, indifference between two feasible schools requires

$$u^1 = u^2,$$

and indifference between school s and the outside option requires

$$u^s = 0.$$

These conditions define Lebesgue-null hyperplanes in Θ.

At age one, the transition kernels do not depend on the current type. Consequently, for every current assignment $s \in \mathcal{S}_0$, its expected continuation value is a scalar $C_s(\mathbf{P})$ that depends on the cutoff path but not on the current payoff vector. Indifference between schools s and j therefore requires

$$u^s - u^j = C_j(\mathbf{P}) - C_s(\mathbf{P}),$$

while indifference between school s and the outside option requires

$$u^s = C_\emptyset(\mathbf{P}) - C_s(\mathbf{P}).$$

Again, these are Lebesgue-null hyperplanes. Since there are finitely many pairs of assignments, the dynamic indifference set is contained in a finite union of null sets. Assumption 2.2 follows.

Convergence to benchmark demand. Let $D^1_\varepsilon(x, 0)$ denote stationary demand for school 1 in the regularized economy. Consider the two disjoint intervals

$$I^- := \left[\frac{1}{4}, \frac{3}{8}\right], \qquad I^+ := \left[\frac{11}{16}, \frac{13}{16}\right]. \tag{49}$$

These intervals lie strictly on opposite sides of the insurance threshold $x = 1/2$.

We show that, for sufficiently small ε, agents drawn from the concentrated components of both (46) and (48) behave exactly as in the benchmark throughout $I^- \cup I^+$.

First, if

$$\varepsilon < \frac{3}{16},$$

then the priority generated by (47) satisfies

$$R^{(s),s} \geq 1 - \varepsilon > \frac{13}{16}.$$

Thus, throughout $I^- \cup I^+$, attending a school when young continues to guarantee access to it when old.

Second, if $\varepsilon < 1/4$, every payoff vector in $C_\varepsilon(u^O)$ ranks school 1 strictly above school 2. Define

$$m_\varepsilon := \int_{C_\varepsilon(u^O)} (u'^1 - u'^2) \, \mathsf{U}\big(C_\varepsilon(u^O)\big)(du').$$

Then

$$\left| m_\varepsilon - \frac{1}{2} \right| \leq 2\varepsilon.$$

For an entrant whose payoff vector belongs to $C_\varepsilon(u^Y)$, the payoff advantage of choosing school 1 rather than school 2, conditional on both being feasible, is

$$u^1 - u^2 + (1-\varepsilon)x\, m_\varepsilon. \tag{50}$$

The full-support component of the transition does not appear in the difference because its law is the same after both assignments. Moreover, uniformly over $x \in [0,1]$ and $u \in C_\varepsilon(u^Y)$,

$$\left| u^1 - u^2 + (1-\varepsilon)x\, m_\varepsilon - \left(-\frac{1}{4} + \frac{x}{2}\right) \right| \le 5\varepsilon. \tag{51}$$

On I^-, the benchmark advantage satisfies

$$-\frac{1}{4} + \frac{x}{2} \le -\frac{1}{16},$$

whereas on I^+ it satisfies

$$-\frac{1}{4} + \frac{x}{2} \ge \frac{3}{32}.$$

It follows from (51) that, for sufficiently small ε, concentrated-component entrants choose school 2 whenever both schools are feasible throughout I^- and choose school 1 whenever it is feasible throughout I^+, exactly as in the benchmark.

For sufficiently small ε, agents drawn from the concentrated components of the entry and transition laws make the same choices as in the benchmark throughout $I^- \cup I^+$. At age one, behavior can therefore differ from the benchmark only for the ε fraction of entrants drawn from the full-support component. At age two, behavior can differ only if either the entry draw or the transition draw comes from its full-support component, which occurs with probability at most 2ε. Since each age group has mass $1/2$, it follows that

$$\sup_{x \in I^- \cup I^+} \left| D^1_\varepsilon(x,0) - D^1(x,0) \right| \le \frac{\varepsilon}{2} + \frac{2\varepsilon}{2} = \frac{3}{2}\varepsilon \longrightarrow 0. \tag{52}$$

Two market-clearing cutoffs in the regularized economy. At the endpoints of the first interval,

$$D^1\left(\frac{1}{4}, 0\right) = \frac{3}{8} > \frac{11}{32} > \frac{5}{16} = D^1\left(\frac{3}{8}, 0\right).$$

At the endpoints of the second interval,

$$D^1\left(\frac{11}{16}, 0\right) = \frac{215}{512} > \frac{176}{512} = \frac{11}{32} > \frac{135}{512} = D^1\left(\frac{13}{16}, 0\right).$$

All four inequalities are strict. By (52), for sufficiently small ε they remain true after replacing D^1 by D^1_ε:

$$\begin{aligned} D^1_\varepsilon\left(\frac{1}{4}, 0\right) > q^1 > D^1_\varepsilon\left(\frac{3}{8}, 0\right), \\ D^1_\varepsilon\left(\frac{11}{16}, 0\right) > q^1 > D^1_\varepsilon\left(\frac{13}{16}, 0\right). \end{aligned} \tag{53}$$

Under Assumptions 2.1 and 2.2, stationary demand is continuous in the cutoff vector. The intermediate-value theorem applied to (53) therefore gives

$$x^-_\varepsilon \in \left(\frac{1}{4}, \frac{3}{8}\right)$$

and

$$x_\varepsilon^+ \in \left(\frac{11}{16}, \frac{13}{16}\right)$$

such that

$$D_\varepsilon^1(x_\varepsilon^-, 0) = D_\varepsilon^1(x_\varepsilon^+, 0) = q^1.$$

Hence the two cutoff vectors

$$P_\varepsilon^- := (x_\varepsilon^-, 0), \qquad P_\varepsilon^+ := (x_\varepsilon^+, 0)$$

are distinct and clear the market for school 1.

At either vector, total assignment cannot exceed one. Consequently,

$$D_\varepsilon^2(P_\varepsilon^\pm) \leq 1 - D_\varepsilon^1(P_\varepsilon^\pm) = 1 - \frac{11}{32} = \frac{21}{32} < \frac{3}{4} = q^2.$$

Because the cutoff of school 2 is zero, this strict inequality is consistent with market clearing. Thus, P_ε^- and P_ε^+ are both stationary market-clearing cutoff vectors in an economy satisfying Assumptions 2.1 and 2.2.

Robustness. We finally show that the multiplicity just obtained persists under small perturbations of the primitives. To make the statement precise, consider the metric

$$\begin{aligned} d(\mathcal{E}, \widetilde{\mathcal{E}}) := \|q - \widetilde{q}\| + \|p - \widetilde{p}\|_{\mathrm{TV}} \\ + \sup_{\theta \in \Theta,\, s \in \mathcal{S}_0} \left\| Q_1(\cdot \mid \theta, s) - \widetilde{Q}_1(\cdot \mid \theta, s) \right\|_{\mathrm{TV}}, \end{aligned} \tag{54}$$

on the class of economies satisfying the maintained regularity assumptions.

In the present two-age environment, stationary demand at any fixed cutoff depends continuously on the primitives under this metric. To see this, terminal-age values are bounded and independent of the transition kernel. Uniform total-variation convergence of the transition kernels therefore implies uniform convergence of age-one continuation values. Since the limiting economy places zero mass on score boundaries and dynamic indifference sets, the induced age-one assignments converge in probability. The entry distributions then give convergence of age-one demand, while the law of motion and uniform convergence of the transition kernels give convergence of the age-two distributions and demand.

Fix a sufficiently small ε for which (53) holds. The four endpoint excess-demand values in (53) are all nonzero. Hence their signs are preserved under sufficiently small perturbations of the primitives. Moreover,

$$q^1 + q^2 - 1 = \frac{3}{32} > 0.$$

Continuity with respect to the primitives therefore implies that there is a neighborhood $\mathcal{N}$ of

$$\mathcal{E}_\varepsilon = \left[\{1, 2\}, q, p_\varepsilon, Q_{1,\varepsilon}\right]$$

such that every

$$\widetilde{\mathcal{E}} = \left[\{1, 2\}, \widetilde{q}, \widetilde{p}, \widetilde{Q}_1\right] \in \mathcal{N}$$

satisfies

$$\begin{aligned} \widetilde{D}^1\left(\frac{1}{4}, 0\right) > \widetilde{q}^1 > \widetilde{D}^1\left(\frac{3}{8}, 0\right), \\ \widetilde{D}^1\left(\frac{11}{16}, 0\right) > \widetilde{q}^1 > \widetilde{D}^1\left(\frac{13}{16}, 0\right), \end{aligned}$$

as well as

$$\widetilde{q}^1 + \widetilde{q}^2 > 1.$$

Continuity of stationary demand in the cutoff again yields one solution $\widetilde{x}^-$ in $(1/4, 3/8)$ and another solution $\widetilde{x}^+$ in $(11/16, 13/16)$ to

$$\widetilde{D}^1(x, 0) = \widetilde{q}^1.$$

At either solution,

$$\widetilde{D}^2(x, 0) \leq 1 - \widetilde{q}^1 < \widetilde{q}^2.$$

Thus,

$$(\widetilde{x}^-, 0) \qquad \text{and} \qquad (\widetilde{x}^+, 0)$$

are two distinct stationary market-clearing cutoff vectors.

Hence multiplicity persists under every sufficiently small admissible perturbation of the capacity vector, entry distribution, and transition kernel. □

C.5 Finite economies: proof of Theorem 3.2

Throughout, $\mathrm{d_{BL}}$ denotes the bounded-Lipschitz metric, which metrizes weak convergence; product spaces with a finite set are endowed with the metric in which distinct elements of the finite set are at distance one. For a distribution ν on $\mathcal{R} \times \mathcal{L}$ we write D^ν for the induced aggregate demand, and $\|D^{\nu'} - D^\nu\|_\infty := \max_s \sup_P |D_s^{\nu'}(P) - D_s^\nu(P)|$.

We write $\hat{\eta}_{t,a}^n$ for the empirical distribution of the date-t types of agents of age a in $\mathcal{E}^n$, $\hat{\nu}_t^n$ for the induced empirical distribution of reports and scores, and P_t^n for the realised market-clearing cutoffs; $\eta_{t,a}$, ν_t, P_t denote their continuum counterparts under σ.

Preliminary step: a law of large numbers for conditionally independent types.

Lemma C.1. *Let $\mathsf{X} = [0,1]^k \times \mathcal{K}$ with $\mathcal{K}$ finite, let $(\mathcal{G}_n)$ be σ-fields, and for each n let $X_1^n, \ldots, X_{N_n}^n$ be conditionally independent given $\mathcal{G}_n$, with $\mathcal{G}_n$-measurable conditional laws $\mu_1^n, \ldots, \mu_{N_n}^n$, not necessarily identical. Suppose $N_n/n \to c > 0$. Then*

$$\mathrm{d_{BL}}\left(\frac{1}{N_n}\sum_i \delta_{X_i^n} \ , \ \frac{1}{N_n}\sum_i \mu_i^n\right) \xrightarrow[n\to\infty]{} 0 \qquad \textit{almost surely.}$$

Proof. Partitioning X into $m^k|\mathcal{K}|$ cells of diameter $1/m$ and bounding each cell's fluctuation by its conditional standard deviation gives $\mathbb{E}[\mathrm{d_{BL}} \mid \mathcal{G}_n] \leq (m^k|\mathcal{K}|/N_n)^{1/2} + 2/m$ for every m. Since replacing a single X_i^n changes $\mathrm{d_{BL}}$ by at most $2/N_n$, McDiarmid's inequality applied conditionally on $\mathcal{G}_n$ gives Gaussian concentration at scale $N_n^{-1/2}$; the resulting tail bounds are summable in n, and Borel–Cantelli concludes. □

Note that neither identical distribution nor unconditional independence is used. This is what allows the lemma to be applied at each date, to type distributions that are themselves random and endogenous.

Convergence of demands.

Lemma C.2. *Let $\nu, \nu^n \in \Delta(\mathcal{R} \times \mathcal{L})$ with $\mathrm{d_{BL}}(\nu^n, \nu) \to 0$, and suppose the score marginal of ν is atomless in each coordinate, in the sense that $\omega^\nu(\varepsilon) := \max_s \sup_x \nu(\{|r_s - x| \le \varepsilon\}) \to 0$ as $\varepsilon \to 0$. Then $\|D^{\nu^n} - D^\nu\|_\infty \to 0$.*

Proof. Bracket $\mathbb{1}_{\{d(P|r,\ell)=s\}}$ between products of one-dimensional Lipschitz functions $g^\pm$ agreeing with it outside ε-tubes around the cutoffs, with Lipschitz constant at most S/ε uniformly in (P, s, ℓ). This yields, for every $\varepsilon > 0$, $\|D^{\nu^n} - D^\nu\|_\infty \le (3 + S/\varepsilon)\,\mathrm{d_{BL}}(\nu^n, \nu) + S\omega^\nu(\varepsilon)$. Letting $n \to \infty$ and then $\varepsilon \to 0$ gives the claim. □

Note that the bound is uniform in P, which is later an important requirement (pointwise convergence of demands would not suffice).

From demands to cutoffs; robustness to deviations.

Lemma C.3. *Let D be continuous with a unique market-clearing cutoff vector $P^\star$ for capacities q. Let (D^n, q^n) satisfy $\|D^n - D\|_\infty \to 0$ and $q^n \to q$, and let P^n be any market-clearing cutoff vector for (D^n, q^n). Then $P^n \to P^\star$.*

Proof. Cutoffs lie in a compact set, so it suffices to show every convergent subsequence has limit $P^\star$. Along such a subsequence with limit P, $D^n(P^n) \to D(P)$ by uniform convergence and continuity of D; passing to the limit in the market-clearing conditions shows P is market clearing for (D, q), hence $P = P^\star$. □

Lemma C.4. *Fix $m \in \mathbb{N}$ and let $\hat\nu_t^n$ and $\hat\nu_t^{n,-J}$ denote the empirical report distributions of, respectively, all agents present at t and all agents outside a set J with $|J| \le m$. Then $\|D^{\hat\nu_t^n} - D^{\hat\nu_t^{n,-J}}\|_\infty \le m/(nA)$, whatever the agents in J report. Consequently the conclusions of Lemmas C.2 and C.3 are unaffected by the behaviour of any fixed finite set of agents.*

Proof. The two are empirical measures whose atom sets differ in at most m points, and each demand region is a measurable set. □

Inductive step: propagation across dates. For an agent of age a with date-t type θ, write $s^n(\theta) := d(P_t^n \mid r, \sigma_{t,a}(\theta))$ for their realized assignment and $s^\infty(\theta) := d(P_t \mid r, \sigma_{t,a}(\theta))$ for the assignment they would receive at the continuum cutoffs, and let $\Xi_{t,a}(\theta) := (\theta, s^\infty(\theta))$. Write $\hat\lambda_{t,a}^n$ for the empirical distribution of the pairs $(\theta_i, s^n(\theta_i))$ and $\lambda_{t,a} := \Xi_{t,a}\#\eta_{t,a}$.

Lemma C.5. *Suppose that, along a realisation, $\hat\eta_{t,a}^n \Rightarrow \eta_{t,a}$ for every a and $P_t^n \to P_t$. Then $\hat\lambda_{t,a}^n \Rightarrow \lambda_{t,a}$ for every a, and consequently, almost surely, $\hat\eta_{t+1,a+1}^n \Rightarrow \eta_{t+1,a+1}$ for every $a \le A-1$.*

Proof. Three steps. *(i)* The discontinuity set of $\Xi_{t,a}$ is contained in $\bigcup_\ell \partial\sigma_{t,a}^{-1}(\{\ell\}) \cup \bigcup_j \{r_j = P_{t,j}\}$, null by regularity of σ and by absolute continuity of $\eta_{t,a}$ respectively; the continuous mapping theorem gives $\Xi_{t,a}\#\hat\eta_{t,a}^n \Rightarrow \lambda_{t,a}$. *(ii)* The realized and frozen-cutoff assignments differ only for agents whose score lies within $\|P_t^n - P_t\|_\infty$ of a coordinate of P_t; for fixed $\varsigma > 0$ this collar is contained in a closed set of $\eta_{t,a}$-measure at most $2S\varsigma\bar f$ once $\|P_t^n - P_t\|_\infty \le \varsigma$, so the portmanteau theorem and then $\varsigma \to 0$ give that the fraction of such agents vanishes; hence $\hat\lambda_{t,a}^n \Rightarrow \lambda_{t,a}$. *(iii)* Conditionally on the date-t history the date-$(t+1)$ types are independent with laws $Q_a(\cdot \mid \theta_i, s^n(\theta_i))$, so Lemma C.1 reduces the claim to $\hat\lambda_{t,a}^n Q_a \Rightarrow \lambda_{t,a} Q_a$, which follows from (ii) and weak continuity of Q_a. □

Approximate optimality.

Lemma C.6. *Fix t and let $g_i^n := \sup_{\pi \in \Pi_{t,i}} G_i^n(\pi|\sigma) - G_i^n(\sigma|\sigma)$. Then*

$$\mathbb{E}\left[\frac{1}{nA}\sum_{i \in I(t)} g_i^n\right] \xrightarrow[n\to\infty]{} 0.$$

Proof. By Lemma C.4 and the induction, $\|P_k^n - P_k\|_\infty \to 0$ for every k in agent i's remaining lifetime, uniformly over deviations. Couple the realized trajectory with the trajectory she would follow against the deterministic path (P_k); the two agree until the first date at which an assignment differs, which requires one of her scores to lie in a neighborhood around a cutoff, an event of vanishing probability by the argument of Lemma C.5(ii). Since payoffs are bounded and lifetimes finite, her value in $\mathcal{E}^n$ converges to her value against (P_k), uniformly over plans; continuum interim optimality of σ then bounds the limiting gain by zero. □

Proof of the Theorem. Parts (i)–(iii) follow by induction on t. At $t = 0$ types are i.i.d. draws, so Lemma C.1 gives (i); Lemmas C.2 and C.3 then give (ii) and (iii). The inductive step is Lemma C.5 followed again by Lemmas C.2 and C.3. The equilibrium statement is Lemma C.6 together with Markov's inequality, which converts the vanishing average gain into a vanishing mass of agents with gain above any fixed ε.

D Proofs for the Implementation Results

D.1 Proof of Theorem 4.1

We first record the regularity implication used in the proof.

Lemma D.1 (Almost-everywhere uniqueness)**.** *Suppose that Assumptions 4.2 and 4.1 hold. Then, at every stationary equilibrium, the dynamically optimal school is unique for η_a-almost every state at every age. Under Axiom 1, the equilibrium assignment is consequently independent of the message record almost everywhere.*

Proof. Let $\bar{\eta}_a$ denote the marginal of η_a on $\mathcal{U} \times \mathcal{H}_a \times \mathcal{Y}$. We show by induction that

$$\bar{\eta}_a \ll \Lambda_a \qquad \text{for every } a \in \mathcal{A}. \tag{55}$$

At age one, this follows from the absolute continuity of $p^\star$, the uniform distribution of the current draw, and the fact that $\mathcal{H}_1 = \{\varnothing\}$.

Suppose that (55) holds at age a. For every assignment $s \in \mathcal{S}_0$, define a measure on $\mathcal{U} \times \mathcal{H}_a \times \mathcal{Y}$ by

$$\nu_{a,s}(C) := \int \mathbf{1}\left\{(u,h,y) \in C,\ \pi_a^\sigma(u,h,\rho,y) = s\right\} \eta_a(du, dh, d\rho, dy).$$

Thus, $\nu_{a,s}$ is the distribution of the current payoff type, physical history, and public draw among agents assigned to s, without normalizing its total mass. For every measurable set C,

$$\nu_{a,s}(C) \le \bar{\eta}_a(C).$$

Hence $\nu_{a,s} \ll \Lambda_a$. Let $g_{a,s}$ denote its Radon–Nikodym density:

$$\nu_{a,s}(du, dh, dy) = g_{a,s}(u,h,y)\Lambda_a(du, dh, dy).$$

By Assumption 4.1, the kernel $Q_a^\star(\cdot \mid u,s)$ is absolutely continuous with respect to Lebesgue measure for every (u,s). Since it is a measurable probability kernel, it admits a jointly measurable density $q_{a,s}$ such that

$$Q_a^\star(du' \mid u,s) = q_{a,s}(u' \mid u)\lambda_{\mathcal{U}}(du').$$

Every history in $\mathcal{H}_{a+1}$ can be written uniquely as $h \oplus (s,y)$ for some $h \in \mathcal{H}_a$, $s \in \mathcal{S}_0$, and $y \in \mathcal{Y}$. Define

$$g_{a+1}\big(u', h \oplus (s,y), y'\big) := \int_{\mathcal{U}} q_{a,s}(u' \mid u) g_{a,s}(u,h,y)\, \lambda_{\mathcal{U}}(du).$$

This is a measurable nonnegative function. The stationary law of motion then gives, for every measurable $B \subseteq \mathcal{U} \times \mathcal{H}_{a+1} \times \mathcal{Y}$,

$$\begin{aligned}\bar\eta_{a+1}(B) &= \sum_{s\in\mathcal{S}_0} \int \int_{\mathcal{U}} \int_{\mathcal{Y}} \mathbf{1}\left\{(u', h\oplus(s,y), y') \in B\right\} \\ &\qquad \times q_{a,s}(u' \mid u)\, \lambda_{\mathcal{Y}}(dy')\, \lambda_{\mathcal{U}}(du')\, \nu_{a,s}(du,dh,dy) \\ &= \int_B g_{a+1}(u',h',y')\, \Lambda_{a+1}(du',dh',dy').\end{aligned}$$

Consequently, $\bar\eta_{a+1} \ll \Lambda_{a+1}$, completing the induction.

Assumption 4.2 implies that $\mathcal{I}_a^\varphi(\xi^\sigma)$ is Λ_a-null. By (55), it is therefore also $\bar\eta_a$-null. Equivalently,

$$\eta_a\left(\{(u,h,\rho,y) : (u,h,y) \in \mathcal{I}_a^\varphi(\xi^\sigma)\}\right) = 0.$$

Thus, the dynamically optimal school is unique for η_a-almost every state.

Under Axiom 1, the menu and the associated dynamic choice problem depend on (u,h,y) but not on ρ. Let $\bar\pi_a^\sigma(u,h,y)$ denote the unique dynamically optimal school whenever it is unique, and define it arbitrarily on the remaining null set. Since equilibrium assignments are dynamically optimal, we have

$$\pi_a^\sigma(u,h,\rho,y) = \bar\pi_a^\sigma(u,h,y)$$

for η_a-almost every (u,h,ρ,y). Hence, the equilibrium assignment is independent of the message record almost everywhere. □

Proof of Theorem 4.1. Fix a mechanism φ and a stationary equilibrium σ satisfying the hypotheses, and write (η,ξ,π) for the induced stationary objects. By Lemma D.1, after modifying the allocation on null sets if necessary, we may write $\pi_a(u,h,y)$ without the message-record argument. Let

$$\mu^s(\pi) := \frac{1}{A}\sum_{a=1}^{A} \int \mathbf{1}\{\pi_a(u,h,y) = s\}\, \eta_a(du,dh,d\rho,dy)$$

denote the stationary mass assigned to school s.

Menu representation. Because the mechanism has no message records, a current message affects the agent's future only through the school it induces. The equilibrium allocation therefore solves

$$\pi_a(u,h,y) \in \operatorname*{arg\,max}_{s\in\mathcal{M}_a(h,y|\xi)} \left\{ u^s + \int_{\mathcal{U}} \int_{\mathcal{Y}} U_{a+1}^\varphi(u', h\oplus(s,y), y' \mid \xi) \lambda_{\mathcal{Y}}(dy') Q_a^\star(du' \mid u,s) \right\}, \tag{56}$$

with the continuation term omitted at age A.

Encoding menus and physical states. For each age, choose a Borel injection

$$\kappa_a : \mathcal{H}_a \times \mathcal{Y} \longrightarrow (-1/4, 1/4)$$

whose inverse on its image is Borel, and define

$$g_a^s(h,y) := \kappa_a(h,y) + \begin{cases} \frac{1}{2}, & s \in \mathcal{M}_a(h,y \mid \xi), \\ -\frac{1}{2}, & s \notin \mathcal{M}_a(h,y \mid \xi), \end{cases} \qquad s \in \mathcal{S}. \tag{57}$$

For every school s,

$$s \in \mathcal{M}_a(h,y \mid \xi) \quad \Longleftrightarrow \quad g_a^s(h,y) \geq 0. \tag{58}$$

Moreover, the common residual is recovered from any coordinate by subtracting $1/2$ when that coordinate is positive and adding $1/2$ when it is negative. Since κ_a is injective, the score vector identifies (h,y) measurably.

Define the entry-score distribution by

$$\widehat{p}(B) := \lambda_{\mathcal{Y}}\big(\{y : g_1(\varnothing, y) \in B\}\big). \tag{59}$$

For a score r in the image of g_a, write $(h_a(r), y_a(r)) := g_a^{-1}(r)$ and set, for $a < A$,

$$\widehat{Q}_a(B \mid r, s) := \int_{\mathcal{Y}} \mathbf{1}\{g_{a+1}(h_a(r) \oplus (s, y_a(r)), y') \in B\} \lambda_{\mathcal{Y}}(dy'). \tag{60}$$

Define the kernel arbitrarily off the image of g_a. The Borel properties above imply that this is a measurable probability kernel.

Equivalence of the dynamic problems. Set $P = (0, \ldots, 0)$. Equation (58) implies

$$\mathcal{F}(g_a(h,y); P) = \mathcal{M}_a(h,y \mid \xi) \tag{61}$$

for every on-path physical state. Let $\widehat{V}_a(u, r \mid P)$ denote the stationary value in the designed score economy. We show by backward induction that

$$\widehat{V}_a(u, g_a(h,y) \mid P) = U_a^{\varphi}(u,h,y \mid \xi). \tag{62}$$

At age A, the two problems maximize the same flow payoff over the same feasible set.

Suppose that (62) holds at age $a+1$. Fix an on-path state (u,h,y) and a current assignment $s \in \mathcal{S}_0$. By the definition of the score-transition kernel in (60), for every bounded measurable function $\psi : \mathcal{R} \to \mathbb{R}$,

$$\int_{\mathcal{R}} \psi(r') \, \widehat{Q}_a(dr' \mid g_a(h,y), s) = \int_{\mathcal{Y}} \psi\big(g_{a+1}(h \oplus (s,y), y')\big) \, \lambda_{\mathcal{Y}}(dy').$$

Applying this identity to $\psi(r') = \widehat{V}_{a+1}(u', r' \mid P)$ and then using the induction hypothesis gives

$$\begin{aligned} &\int_{\mathcal{U}} \int_{\mathcal{R}} \widehat{V}_{a+1}(u', r' \mid P) \, \widehat{Q}_a(dr' \mid g_a(h,y), s) \, Q_a^{\star}(du' \mid u, s) \\ &\quad = \int_{\mathcal{U}} \int_{\mathcal{Y}} \widehat{V}_{a+1}\big(u', g_{a+1}(h \oplus (s,y), y') \mid P\big) \, \lambda_{\mathcal{Y}}(dy') \, Q_a^{\star}(du' \mid u, s) \\ &\quad = \int_{\mathcal{U}} \int_{\mathcal{Y}} U_{a+1}^{\varphi}\big(u', h \oplus (s,y), y' \mid \xi\big) \, \lambda_{\mathcal{Y}}(dy') \, Q_a^{\star}(du' \mid u, s). \end{aligned}$$

Thus, for every current assignment s, the current flow payoff and the expected continuation payoff are the same in the two problems. Since their feasible sets also coincide by (61), they maximize the same choice-specific values over the same set of assignments. This proves (62) at age a.

The two problems therefore have the same argmax correspondence at every on-path state. By Lemma D.1, this correspondence is a singleton almost everywhere. Hence,

$$d_a((g_a(h,y),u) \mid P) = \pi_a(u,h,y) \tag{63}$$

for η_a-almost every state.

Stationarity, market clearing, and DA. Define the age-specific distribution of designed types by

$$\widehat{\eta}_a(B) := \eta_a\big(\{(u,h,\rho,y) : (g_a(h,y),u) \in B\}\big). \tag{64}$$

Thus, $\widehat{\eta}_a$ is the image of η_a under the measurable map

$$(u,h,\rho,y) \longmapsto (g_a(h,y),u).$$

At age one, the payoff type has distribution $p^\star$, the public draw is uniform, and the two are independent. Hence, by the definition of $\widehat{p}$,

$$\widehat{\eta}_1 = \widehat{p} \otimes p^\star.$$

We next verify that $(\widehat{\eta}_a)_{a\in\mathcal{A}}$ satisfies the stationary law of motion in the designed score economy. Given $\widehat{\eta}_a$, let $\widetilde{\eta}_{a+1}$ denote the next-age distribution generated by the designed transition kernels and the demand rule at P. Thus, for every Borel set $C \subseteq \mathcal{R}\times\mathcal{U}$,

$$\begin{aligned}\widetilde{\eta}_{a+1}(C) := \int_{\mathcal{R}\times\mathcal{U}} \int_{\mathcal{R}} \int_{\mathcal{U}} \mathbf{1}\{(r',u') \in C\} \\ \times \widehat{Q}_a\big(dr' \mid r, d_a((r,u) \mid P)\big) Q_a^\star\big(du' \mid u, d_a((r,u) \mid P)\big) \widehat{\eta}_a(dr,du).\end{aligned}$$

We show that

$$\widehat{\eta}_{a+1} = \widetilde{\eta}_{a+1}.$$

Both sides are probability measures on $\mathcal{R}\times\mathcal{U}$. It is enough to establish equality on sets of the form $B_\mathcal{R} \times B_\mathcal{U}$, where $B_\mathcal{R} \subseteq \mathcal{R}$ and $B_\mathcal{U} \subseteq \mathcal{U}$ are Borel.

By the definition of $\widehat{\eta}_{a+1}$ in (64),

$$\begin{aligned}&\widehat{\eta}_{a+1}\big(B_\mathcal{R} \times B_\mathcal{U}\big) \\ &= \eta_{a+1}\left(\left\{(u',h',\rho',y') : \begin{array}{l} g_{a+1}(h',y') \in B_\mathcal{R}, \\ u' \in B_\mathcal{U} \end{array}\right\}\right).\end{aligned} \tag{65}$$

Under the stationary law of motion of the original mechanism, a current state (u,h,ρ,y), a next-period payoff u', and a next public draw y' generate the next state

$$\big(u', h \oplus (\pi_a(u,h,\rho,y),y), \rho \oplus \sigma_a(u,h,\rho,y), y'\big).$$

Substituting this expression into (65) gives

$$\begin{aligned}&\widehat{\eta}_{a+1}\big(B_\mathcal{R} \times B_\mathcal{U}\big) \\ &= \int_{\mathcal{U}\times\mathcal{H}_a\times\mathcal{R}_a^m\times\mathcal{Y}} \int_{\mathcal{U}} \int_{\mathcal{Y}} \mathbf{1}\left\{g_{a+1}\big(h \oplus (\pi_a(u,h,\rho,y),y), y'\big) \in B_\mathcal{R}\right\} \\ &\qquad \times \mathbf{1}\{u' \in B_\mathcal{U}\}\, \lambda_\mathcal{Y}(dy') Q_a^\star\big(du' \mid u, \pi_a(u,h,\rho,y)\big) \eta_a(du,dh,d\rho,dy).\end{aligned} \tag{66}$$

The updated message record does not appear in the indicator because the designed score $g_{a+1}(h', y')$ depends only on the next physical history and next public draw.

For notational convenience, define

$$s_a(u, h, y) := d_a\big((g_a(h, y), u) \mid P\big).$$

By (63),

$$\pi_a(u, h, \rho, y) = s_a(u, h, y)$$

for η_a-almost every (u, h, ρ, y). We may therefore replace $\pi_a(u, h, \rho, y)$ by $s_a(u, h, y)$ in (66) without changing the value of the integral. We obtain

$$\begin{aligned}
&\widehat{\eta}_{a+1}\big(B_{\mathcal{R}} \times B_{\mathcal{U}}\big) \\
&= \int \left[\int_{\mathcal{Y}} \mathbf{1}\left\{ g_{a+1}\big(h \oplus (s_a(u, h, y), y), y'\big) \in B_{\mathcal{R}} \right\} \lambda_{\mathcal{Y}}(dy') \right] \\
&\qquad \times \left[\int_{\mathcal{U}} \mathbf{1}\{u' \in B_{\mathcal{U}}\} Q_a^{\star}\big(du' \mid u, s_a(u, h, y)\big) \right] \eta_a(du, dh, d\rho, dy). \qquad (67)
\end{aligned}$$

The first term in brackets is, by (60),

$$\widehat{Q}_a\big(B_{\mathcal{R}} \mid g_a(h, y), s_a(u, h, y)\big),$$

while the second is

$$Q_a^{\star}\big(B_{\mathcal{U}} \mid u, s_a(u, h, y)\big).$$

Consequently,

$$\begin{aligned}
&\widehat{\eta}_{a+1}\big(B_{\mathcal{R}} \times B_{\mathcal{U}}\big) \\
&= \int \widehat{Q}_a\left(B_{\mathcal{R}} \mid g_a(h, y), d_a((g_a(h, y), u) \mid P)\right) \\
&\qquad \times Q_a^{\star}\left(B_{\mathcal{U}} \mid u, d_a((g_a(h, y), u) \mid P)\right) \eta_a(du, dh, d\rho, dy). \qquad (68)
\end{aligned}$$

To pass from the original state distribution η_a to the designed distribution $\widehat{\eta}_a$, define the measurable function

$$\begin{aligned}
\Psi_a(r, u) := \widehat{Q}_a\left(B_{\mathcal{R}} \mid r, d_a((r, u) \mid P)\right) \\
\times Q_a^{\star}\left(B_{\mathcal{U}} \mid u, d_a((r, u) \mid P)\right).
\end{aligned}$$

The integrand in (68) is precisely $\Psi_a(g_a(h, y), u)$. Moreover, by the definition of $\widehat{\eta}_a$ as the image of η_a under

$$(u, h, \rho, y) \longmapsto (g_a(h, y), u),$$

we have, for every bounded measurable function $\Psi : \mathcal{R} \times \mathcal{U} \to \mathbb{R}$,

$$\int_{\mathcal{R}\times\mathcal{U}} \Psi(r, u)\, \widehat{\eta}_a(dr, du) = \int \Psi(g_a(h, y), u)\, \eta_a(du, dh, d\rho, dy). \qquad (69)$$

Applying (69) to Ψ_a yields

$$\begin{aligned}
&\widehat{\eta}_{a+1}\big(B_{\mathcal{R}} \times B_{\mathcal{U}}\big) \\
&= \int_{\mathcal{R}\times\mathcal{U}} \widehat{Q}_a\left(B_{\mathcal{R}} \mid r, d_a((r, u) \mid P)\right) Q_a^{\star}\left(B_{\mathcal{U}} \mid u, d_a((r, u) \mid P)\right) \widehat{\eta}_a(dr, du). \qquad (70)
\end{aligned}$$

The right-hand side is exactly $\widetilde{\eta}_{a+1}(B_{\mathcal{R}} \times B_{\mathcal{U}})$. Hence $\widehat{\eta}_{a+1}$ and $\widetilde{\eta}_{a+1}$ agree on every measurable rectangle. Thus,

$$\widehat{\eta}_{a+1} = \widetilde{\eta}_{a+1}.$$

Together with $\widehat{\eta}_1 = \widehat{p} \otimes p^\star$, this proves inductively that $(\widehat{\eta}_a)_{a \in \mathcal{A}}$ is exactly the stationary family generated by P in the designed score economy.

For every school s, aggregate demand at P is consequently

$$\begin{aligned} D^s(P) &= \frac{1}{A} \sum_{a=1}^{A} \widehat{\eta}_a \left(\{(r,u) : d_a((r,u) \mid P) = s\} \right) \\ &= \frac{1}{A} \sum_{a=1}^{A} \eta_a \left(\{(u,h,\rho,y) : d_a((g_a(h,y),u) \mid P) = s\} \right) \\ &= \frac{1}{A} \sum_{a=1}^{A} \eta_a \left(\{(u,h,\rho,y) : \pi_a(u,h,\rho,y) = s\} \right) \\ &= \mu^s(\pi) \leq q^s, \end{aligned}$$

where the second equality follows from (64) and the third from (63). Because every component of P is zero, these capacity inequalities are sufficient for market clearing.

For each designed type, choose a report that ranks schools according to their dynamic choice-specific values, resolving ties with a fixed ordering, and let $\widehat{D}$ be the resulting static reported-demand function at the stationary distribution. Reported and dynamic demand agree at P, so

$$P \in \mathcal{C}(\widehat{D}).$$

Since the cutoff domain is $[0,1]^S$, the zero vector is the smallest admissible cutoff vector. It follows that

$$\mathrm{DA}(\widehat{D}) = P.$$

The constructed reports and distributions therefore form a stationary truthful equilibrium under repeated agent-proposing DA.

Let $\widehat{\pi}_a$ denote the allocation induced by this equilibrium. Because the equilibrium cutoff is P and the reports rank schools by their dynamic choice-specific values,

$$\widehat{\pi}_a(r,u) = d_a((r,u) \mid P)$$

for $\widehat{\eta}_a$-almost every (r,u). Hence, for every age a, Borel set $B \subseteq \mathcal{U}$, and school $s \in \mathcal{S}_0$,

$$\begin{aligned} &\widehat{\eta}_a \left(\{(r,u) : u \in B,\ \widehat{\pi}_a(r,u) = s\} \right) \\ &= \widehat{\eta}_a \left(\{(r,u) : u \in B,\ d_a((r,u) \mid P) = s\} \right) \\ &= \eta_a \left(\{(u,h,\rho,y) : u \in B,\ d_a((g_a(h,y),u) \mid P) = s\} \right) \\ &= \eta_a \left(\{(u,h,\rho,y) : u \in B,\ \pi_a(u,h,\rho,y) = s\} \right). \end{aligned}$$

The second equality follows from (64), and the last one from (63). This is precisely the equivalence condition in (32). □

D.2 Nonnegative Score

For an aggregate state $\xi = (\xi_a)_{a\in\mathcal{A}}$, define

$$\mu^s_\varphi(\xi) := \frac{1}{A}\sum_{a=1}^{A}\int_{\mathcal{X}\times\mathcal{H}_a\times\mathcal{R}_a^m\times\mathcal{Y}} \mathbf{1}\{\varphi_a(x,h,y\mid\xi)=s\}\,\xi_a(dx,dh,d\rho,dy). \tag{71}$$

Axiom 2 (Non-exclusion)**.** *For every aggregate state $\xi\in\Xi$ and every school $s\in\mathcal{S}$,*

$$\mu^s_\varphi(\xi) < q^s \quad\Longrightarrow\quad s\in\mathcal{M}_a(h,y\mid\xi)$$

for every age $a\in\mathcal{A}$, physical history $h\in\mathcal{H}_a$, and current draw $y\in\mathcal{Y}$.

Thus, an undersubscribed school must be obtainable by every agent.

Proposition D.1 (Nonnegative-score implementation by a stable selector)**.** *Suppose that Assumptions 4.2 and 4.1 hold. Let φ be a feasible mechanism satisfying Axioms 1 and 2, and let σ be a stationary equilibrium. Then there exist an entry-score distribution and score-transition kernels supported on $[0,1]^S$, together with a time-invariant stable cutoff selector, under which a stationary truthful equilibrium induces an allocation equivalent to π^σ.*

Proof. Fix a mechanism φ and a stationary equilibrium σ satisfying the hypotheses, and write (η,ξ,π) for the induced stationary objects. By Lemma D.1, after modifying the allocation on null sets if necessary, we may write $\pi_a(u,h,y)$ without the message-record argument. Let

$$\mu^s(\pi) := \frac{1}{A}\sum_{a=1}^{A}\int \mathbf{1}\{\pi_a(u,h,y)=s\}\,\eta_a(du,dh,d\rho,dy)$$

denote the stationary mass assigned to school s.

Menu representation. Because the mechanism has no message records, a current message affects the agent's future only through the school it induces. The equilibrium allocation therefore solves the dynamic menu problem

$$\pi_a(u,h,y)\in \operatorname*{arg\,max}_{s\in\mathcal{M}_a(h,y|\xi)}\left\{u^s + \int_{\mathcal{U}}\int_{\mathcal{Y}} U^\varphi_{a+1}(u',h\oplus(s,y),y'\mid\xi)\lambda_{\mathcal{Y}}(dy')Q^\star_a(du'\mid u,s)\right\}, \tag{72}$$

with the continuation term omitted at age A.

Encoding menus and physical states. For each age, choose a Borel injection

$$\kappa_a : \mathcal{H}_a\times\mathcal{Y}\longrightarrow(0,1/2)$$

whose inverse on its image is Borel, and define

$$f^s_a(h,y) := \kappa_a(h,y) + \frac{1}{2}\mathbf{1}\{s\in\mathcal{M}_a(h,y\mid\xi)\},\qquad s\in\mathcal{S}. \tag{73}$$

For every school s,

$$s\in\mathcal{M}_a(h,y\mid\xi)\quad\Longleftrightarrow\quad f^s_a(h,y)\geq\frac{1}{2}. \tag{74}$$

Moreover,

$$\kappa_a(h,y) = f_a^s(h,y) - \frac{1}{2}\mathbf{1}\{f_a^s(h,y) \geq 1/2\},$$

so the score vector identifies (h,y).

Define the entry-score distribution by

$$\widehat{p}(B) := \lambda_{\mathcal{Y}}\big(\{y : f_1(\varnothing, y) \in B\}\big). \tag{75}$$

For a score r in the image of f_a, write $(h_a(r), y_a(r)) := f_a^{-1}(r)$ and set, for $a < A$,

$$\widehat{Q}_a(B \mid r, s) := \int_{\mathcal{Y}} \mathbf{1}\{f_{a+1}(h_a(r) \oplus (s, y_a(r)), y') \in B\}\lambda_{\mathcal{Y}}(dy'). \tag{76}$$

Define the kernel arbitrarily off the image of f_a. The Borel properties above imply that this is a measurable probability kernel.

Equivalence of the dynamic problems. Define

$$P^s := \begin{cases} \frac{1}{2}, & \mu^s(\pi) = q^s, \\ 0, & \mu^s(\pi) < q^s. \end{cases} \tag{77}$$

If $P^s = 1/2$, (74) makes s feasible exactly when it belongs to the menu. If $P^s = 0$, the school is undersubscribed under π, so Axiom 2 places it in every menu; its score is then strictly positive and it is feasible at the zero cutoff. Hence

$$\mathcal{F}(f_a(h,y); P) = \mathcal{M}_a(h, y \mid \xi) \tag{78}$$

for every on-path physical state.

The remainder of the proof follows along the same lines as the proof of Theorem 4.1 and is therefore relatively brief. Let $\widehat{V}_a(u, r \mid P)$ be the stationary value in the designed score economy. Backward induction, using (76) and (78), gives

$$\widehat{V}_a(u, f_a(h,y) \mid P) = U_a^{\varphi}(u, h, y \mid \xi). \tag{79}$$

The two problems therefore have the same argmax correspondence. By Lemma D.1, their selected assignments coincide almost everywhere:

$$d_a((f_a(h,y), u) \mid P) = \pi_a(u, h, y). \tag{80}$$

Define the age-specific distribution of designed types by

$$\widehat{\eta}_a(B) := \eta_a\big(\{(u, h, \rho, y) : (f_a(h,y), u) \in B\}\big). \tag{81}$$

Equations (75), (76), and (80) imply inductively that $(\widehat{\eta}_a)_a$ is exactly the stationary family generated by P in the designed economy. Consequently,

$$D^s(P) = \mu^s(\pi) \leq q^s.$$

Whenever $P^s > 0$, definition (77) gives $\mu^s(\pi) = q^s$. Thus P is stationary market clearing, and (80) gives equivalence of the induced allocations.

Supporting the target cutoff. For every age a and designed type (r,u), let $\ell_a^P(r,u) \in \mathcal{L}$ be the strict rank-order list obtained by ranking all assignments in $\mathcal{S}_0$ according to their stationary dynamic choice-specific values in the designed economy at P, using the same fixed tie-breaking rule as in the definition of $d_a((r,u) \mid P)$. Let $\widehat{\sigma}_a^P(\cdot \mid r,u)$ be degenerate at $\ell_a^P(r,u)$, and let $\widehat{D}$ denote the static reported-demand function induced by the reporting strategy $\widehat{\sigma}^P$ and the stationary age-specific distributions $(\widehat{\eta}_a)_a$.

By construction, at the target cutoff P,

$$\widehat{d}\big(P \mid r, \ell_a^P(r,u)\big) = d_a((r,u) \mid P)$$

for $\widehat{\eta}_a$-almost every (r,u). Therefore, for every school $s \in \mathcal{S}$,

$$\begin{aligned}
\widehat{D}^s(P) &= \frac{1}{A}\sum_{a=1}^{A} \int_{\mathcal{R}\times\mathcal{U}} \mathbf{1}\left\{\widehat{d}\big(P \mid r, \ell_a^P(r,u)\big) = s\right\} \widehat{\eta}_a(dr,du) \\
&= \frac{1}{A}\sum_{a=1}^{A} \widehat{\eta}_a\left(\{(r,u) : d_a((r,u) \mid P) = s\}\right) \\
&= D^s(P).
\end{aligned}$$

Hence reported demand and dynamic demand coincide at P. Since P is stationary market clearing for dynamic demand, it follows that

$$P \in \mathcal{C}(\widehat{D}).$$

We may therefore choose a time-invariant stable cutoff selector Φ satisfying

$$\Phi(\widehat{D}) = P;$$

at every other reported-demand function, Φ may select any element of the corresponding set of static market-clearing cutoffs. Initialize the cohorts already present at date zero at the stationary distributions $(\widehat{\eta}_a)_{a=2}^{A}$ and use the time-invariant reporting strategy $\widehat{\sigma}^P$ at every date. At the resulting cross-sectional distribution, the reported-demand function is $\widehat{D}$, so Φ selects P. The induced assignments are $d_a((r,u) \mid P)$, and the law-of-motion argument above implies that they reproduce the stationary distributions $(\widehat{\eta}_a)_a$. Thus, the cutoff vector, assignments, and age-specific distributions are all time invariant.

Finally, because each report ranks assignments according to their dynamic choice-specific values at P, it induces a dynamically optimal assignment among the schools feasible at that cutoff. The strategy $\widehat{\sigma}^P$ is therefore a truthful stationary equilibrium under Φ. By the assignment equivalence established above, the allocation induced by this equilibrium is equivalent to π^σ.

□

E Additional Results and Examples for Section 4

E.1 Deferred Acceptance as a Particular Case of our General Mechanisms

Repeated agent-proposing Deferred Acceptance is a particular case of the mechanisms in (24). To see this, fix an entry-score distribution $\widehat{p}$ and score-transition kernels $(\widehat{Q}_a)_{a<A}$. By the universal-randomization property of $\mathcal{Y}$, the resulting score process can be represented by measurable maps

$$R_a : \mathcal{H}_a \times \mathcal{Y} \longrightarrow \mathcal{R}, \qquad a \in \mathcal{A},$$

where $R_a(h, y)$ is the current score vector generated by the physical history h and current draw y.

Take the message space to be the finite set $\mathcal{L}$ of strict rank-order lists over $\mathcal{S}_0$. For every aggregate state $\xi = (\xi_a)_{a \in \mathcal{A}}$, define

$$\widehat{D}^s_\xi(P) := \frac{1}{A} \sum_{a=1}^{A} \int \mathbf{1}\left\{ \widehat{d}(P \mid R_a(h, y), \ell) = s \right\} \xi_a(d\ell, dh, d\rho, dy),$$

and let

$$P^{DA}(\xi) := \min \mathcal{C}(\widehat{D}_\xi)$$

be the cutoff vector selected by agent-proposing DA. The associated allocation rule is

$$\varphi^{DA}_a(\ell, h, \rho, y \mid \xi) := \widehat{d}\big(P^{DA}(\xi) \mid R_a(h, y), \ell\big).$$

This rule is feasible because $P^{DA}(\xi)$ clears the reported demand generated by ξ. Obviously, it is anonymous, time invariant, and aggregative. Thus repeated agent-proposing DA belongs to the class of mechanisms considered in Section 4.

E.2 The Value of Message Records

The following example shows that the no-message-records axiom is not without loss. Retaining an agent's past message may allow the mechanism to elicit cardinal preference intensity and thereby improve welfare.

Example (Message records can improve welfare). Let $A = 2$ and $\mathcal{S} = \{1, 2\}$. The capacities are

$$q^1 = \frac{1}{8}, \qquad q^2 = \frac{3}{4}.$$

There are two equally likely payoff types, denoted by H and L. Their age-one payoff vectors are

$$u^H_Y = \left(\frac{1}{4}, \frac{1}{5}\right), \qquad u^L_Y = \left(\frac{3}{10}, \frac{1}{5}\right),$$

and their age-two payoff vectors are

$$u^H_O = \left(1, \frac{1}{10}\right), \qquad u^L_O = \left(\frac{1}{2}, \frac{2}{5}\right).$$

Payoff transitions are deterministic and independent of the age-one assignment:

$$u^H_Y \longmapsto u^H_O, \qquad u^L_Y \longmapsto u^L_O.$$

The particular age-one payoffs are inessential. They are chosen only so that the two entry types are distinct and all school payoffs are positive.

Because each age group has mass $1/2$, at every stationary date there is a mass $1/4$ of each of the four groups: young H, young L, old H, and old L.

A mechanism that records messages. Let the message space be $\mathcal{X} = \{H, L\}$. At age one, every agent is assigned to school 2, irrespective of their message, but the message is recorded. At age two, the current message is ignored. An agent whose age-one message was H receives school 1 with probability $1/2$ and the outside option otherwise, whereas an agent whose age-one message was L receives school 2 with probability one. The age-two draw y implements the first lottery:

$$\varphi_2(H\text{-record}, y) = \begin{cases} 1, & y \le 1/2, \\ \emptyset, & y > 1/2, \end{cases} \qquad \varphi_2(L\text{-record}, y) = 2.$$

The mechanism can be completed away from the equilibrium aggregate state so as to remain feasible. In particular, if $\gamma_H(\xi)$ and $\gamma_L(\xi)$ denote the masses of old agents with records H and L, respectively, replace the preceding assignment probabilities by

$$\alpha_H(\xi) := \min\left\{1, \frac{q^1}{\gamma_H(\xi)}\right\}, \qquad \alpha_L(\xi) := \min\left\{1, \frac{q^2 - 1/2}{\gamma_L(\xi)}\right\},$$

with an arbitrary convention when a denominator is zero. Agents with an H-record receive school 1 when $y \le \alpha_H(\xi)$, and agents with an L-record receive school 2 when $y \le \alpha_L(\xi)$; otherwise they receive the outside option. Since young agents use a mass $1/2$ of school 2, this rule respects both capacities at every aggregate state.

Consider the stationary strategy under which type H reports H when young and type L reports L when young. At the resulting aggregate state,

$$\gamma_H(\xi) = \gamma_L(\xi) = \frac{1}{4},$$

and hence

$$\alpha_H(\xi) = \frac{1}{2}, \qquad \alpha_L(\xi) = 1.$$

The age-one assignment is independent of the report, so only continuation payoffs matter. For type H, the continuation payoff from reporting H is

$$\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 0 = \frac{1}{2} > \frac{1}{10},$$

which is the continuation payoff from reporting L. For type L, the corresponding comparison is

$$\frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} \cdot 0 = \frac{1}{4} < \frac{2}{5}.$$

Thus, the proposed reports are strictly optimal at age one. The mechanism ignores current messages at age two, so the strategy constitutes a stationary equilibrium.

The stationary allocation is therefore

Group	Assignment
all young agents	2,
one half of old type H	1,
the other half of old type H	$\emptyset$,
all old type L	2.

It uses masses $1/8$ and $3/4$ at schools 1 and 2, respectively.

Efficiency. This allocation is the utilitarian first best. A direct comparison of the four groups shows that school 1 should be assigned to old type-H agents. Once its capacity is filled, the values of school 2 satisfy

$$\frac{2}{5} > \frac{1}{5} > \frac{1}{10},$$

for old type L, young agents, and the remaining old type-H agents, respectively. Hence school 2 should be assigned to all old type-L agents and all young agents, leaving the remaining old type-H agents unassigned. These comparisons are strict, so the first-best masses assigned to each type-age group are unique.

The resulting stationary welfare is

$$\frac{1}{2} \cdot \frac{1}{5} + \frac{1}{8} \cdot 1 + \frac{1}{4} \cdot \frac{2}{5} = \frac{13}{40}.$$

Why Axiom 1 prevents the first best. Suppose, toward a contradiction, that a stationary equilibrium of a mechanism satisfying Axiom 1 attained the same allocation. In the first best, every young agent is assigned to school 2. Consequently, an old agent's physical history consists of the same past assignment, together with a past draw whose distribution is independent of their payoff type.

Fix an old agent's physical history, current draw, and the stationary aggregate state. Under Axiom 1, old types H and L face the same set of schools obtainable through their current messages. Moreover, both types have the same strict ordinal ranking,

$$1 \succ 2 \succ \emptyset.$$

Because age two is terminal and assignments are deterministic conditional on the current draw and message, both types must therefore choose the same school from this common menu. Since their physical histories and draws also have the same distribution, their age-two assignment distributions must coincide.

This contradicts the first best, which assigns school 1 only to old type-H agents and school 2 only to old type-L agents. A mechanism satisfying Axiom 1 could distinguish the two types by assigning them differently when young and thereby encoding information in their physical histories, but that would depart from the unique first-best allocation. It would consequently yield strictly lower welfare.

Thus, a mechanism that records individual messages can achieve strictly higher stationary welfare than any stationary equilibrium of any mechanism satisfying Axiom 1.

F Additional results and proofs for Section 5

F.1 Framework and method to find a constrained efficient allocation

F.1.1 Complete Information

There are $A < \infty$ periods, indexed by $a = 1, \ldots, A$, and a finite set of schools S. Let

$$\bar{S} := S \cup \{\varnothing\},$$

where $\varnothing$ is an outside option. The stationary mass of agents of each generation a is $1/A$. An agent is born with a budget $B \in [0, 1]$ and a payoff path

$$\bar{u} = (\bar{u}_1, \ldots, \bar{u}_A) \in \overline{U}, \qquad \bar{u}_a \in [0, 1]^S \text{ for each } a = 1, \ldots, A.$$

While flow utilities over schools belong to $[0,1]^S$, the value of the outside option is normalized to zero. The budget B is drawn according to a probability measure ν on $[0,1]$, while the payoff path $\bar{u}$ is drawn according to a probability measure p on $\overline{U}$. We assume that these two draws are independent, so that the joint law of $(B,\bar{u})$ is $\nu \otimes p$. The entire payoff path $\bar{u}$ is deterministic and known to the agent at age $a=1$. The special case in which payoffs are time-invariant is obtained by restricting $\overline{U}$ to payoff paths satisfying $\bar{u}_1 = \cdots = \bar{u}_A$.

A pure allocation is a measurable map

$$y_x : [0,1] \times \overline{U} \to \bar{S}^A,$$

where

$$y_x(B,\bar{u}) = \big(y_{x,1}(B,\bar{u}), \ldots, y_{x,A}(B,\bar{u})\big)$$

is the lifetime assignment path of type $(B,\bar{u})$. Equivalently, one may note the assignment using the following variables

$$x_{a,B,\bar{u},s} = \mathbf{1}\{y_{x,a}(B,\bar{u}) = s\}, \qquad s \in \bar{S}$$

where $x_{a,B,\bar{u},s}$ takes value 1 when an agent of age a, with budget B and payoff path u is assigned to s. Conversely, consider any measurable array $(x_{a,B,\bar{u},s})$ satisfying $x_{a,B,\bar{u},s} \in \{0,1\}$ and

$$\sum_{s \in \bar{S}} x_{a,B,\bar{u},s} = 1$$

for every a and for $(\nu \otimes p)$-a.e. $(B,\bar{u})$. Letting $y_{x,a}(B,\bar{u})$ be the unique element $s \in \bar{S}$ such that $x_{a,B,\bar{u},s} = 1$ yeilds a lifetime path $y_x(B,\bar{u})$. Thus the path notation y_x and the indicator notation x are interchangeable.

Given an allocation x, define $M(x) := (M_{a,s}(x))_{a,s}$ the aggreagte assignment generated by x where

$$M_{a,s}(x) := \frac{1}{A} \int x_{a,B,\bar{u},s} \nu(dB) p(d\bar{u})$$

The objectives of the designer are described by a finite set of linear constraints on the aggregate assignment induced by an allocation. We collect them in a set $\mathcal{A} \subseteq \mathbb{R}_+^{A \times S}$ offeasible aggregate assignments, and assume throughout that $\mathcal{A}$ is a nonempty polytope. The set of feasible allocations is

$$A := \{x : \ M(x) \in \mathcal{A}\}.$$

Note that whether an allocation is feasible depends on it only through its aggregate: two allocations inducing the same aggregate assignment are indistinguishable for the designer, and, as will be clear below, for the mechanism as well.

Define the lower contour set of $\mathcal{A}$

$$\text{lcs}(\mathcal{A}) := \{x : \exists x' \in \mathcal{A} \text{ with } x \leq x'\}.$$

Lemma 1 in Echenique et al. 2021 gives us the following representation of the lower contour set.

Lemma F.1. *(Representation, Echnique, Miralles and Zhang, 2021, Lemma 1) There exists a finite set $\Omega^\star$ of pairs (β^c, b^c) with $\beta^c \in \mathbb{R}_+^{A \times S}$ and $b^c \geq 0$ such that*

$$\text{lcs}(\mathcal{A}) = \{x : \sum_{a,s} \beta^c_{a,s} M_{a,s} \leq b^c \text{ for all } c \in \Omega^\star\}$$

The inequalities in $\Omega^\star$ translate into constraints on allocations. For $c \in \Omega^\star$ set

$$\alpha^c_{a,s} := \beta^c_{a,s}/A, \text{ for } s \in S \quad \text{and} \quad \alpha^c_{a,\emptyset} := 0,$$

and define the use of constraint c by an allocation x as

$$\Phi_c(x) := \int \sum_{a=1}^{A} \alpha^c_{a,y_{x,a}(B,\bar{u})}\, \nu(dB)p(d\bar{u}).$$

By construction,

$$\Phi_c(x) \;=\; \sum_{a,s} \beta^c_{a,s}\, M_{a,s}(x), \tag{82}$$

so the set

$$C := \{x :\; \Phi_c(x) \le b^c \text{ for every } c \in \Omega^\star\}$$

is exactly the set of allocations whose aggregate lies in $\mathrm{lcs}(\mathcal{A})$. In particular $A \subseteq C$, since $\mathcal{A} \subseteq \mathrm{lcs}(\mathcal{A})$: the market below operates on the relaxed set C, in which all lower bounds have been dropped, and we show that every equilibrium allocation in fact lies in A.

Many constraints are naturally included in $\Omega^\star$. Capacity constraints for example can be included as follows. If school $s \in S$ has capacity q_s, add a constraint $c = s$ to Ω^* with

$$\alpha^s_{a,t} = \frac{1}{A}\mathbf{1}\{t = s\}, \qquad t \in \bar{S}, \qquad b^s = q_s.$$

Then $\Phi_s(x) \le q_s$ is exactly

$$\frac{1}{A}\int \sum_{a=1}^{A} \mathbf{1}\{y_{x,a}(B,\bar{u}) = s\}\, \nu(dB)p(d\bar{u}) \le q_s.$$

More generally, any weighted linear constraint of the form

$$\frac{1}{A}\int \sum_{a=1}^{A}\sum_{s\in\bar{S}} \beta^c_{a,s} x_{a,B,\bar{u},s}\, \nu(dB)p(d\bar{u}) \le b^c, \qquad \beta^c_{a,\varnothing} = 0,$$

is obtained by setting $\alpha^c_{a,s} = \beta^c_{a,s}/A$. Constraints outside this class such as floor constraints are replaced by the ceiling inequalities of Lemma F.1, which define the set C the market only needs to operate on.

Following Echenique et al. we will price the elements of $\Omega^\star$ to internalise the externalities an agent exerts on others. Given a vector of constraint prices $P = (P_c)_{c\in\Omega^*} \in \mathbb{R}_+^{\Omega^*}$, define generation-school prices by

$$P_{a,s} := \sum_{c\in\Omega^*} \alpha^c_{a,s} P_c, \qquad s \in \bar{S}.$$

In particular, $P_{a,\varnothing} = 0$. For a lifetime path $y = (y_1, \ldots, y_A) \in \bar{S}^A$, define its lifetime price and lifetime utility by

$$\Pi_P(y) := \sum_{a=1}^{A} P_{a,y_a},$$

and

$$U(y;\bar{u}) := \sum_{a=1}^{A} \delta^{a-1}\bar{u}_a(y_a),$$

where $\delta \in (0,1)$ is the discount factor. More generally, for each age a, define continuation utility

$$U_a(y;\bar{u}) := \sum_{t=a}^{A} \delta^{t-a} \bar{u}_t(y_t).$$

We impose the following conditions on distributions p and ν, ensuring indifferences have mass zero and budget richness.

Assumption F.1. *For every pair of distinct lifetime paths $y, z \in \bar{S}^A$,*

$$p\left(\{\bar{u} : U(y;\bar{u}) = U(z;\bar{u})\}\right) = 0.$$

Assumption F.2. *The budget law ν is atomless.*

We now define the equilibrium concept allowing us to pin down a constrained efficient allocation.

Definition F.1. *A pair (x^*, P^*) is a stationary market-clearing equilibrium if:*

1. *$x^* \in \mathcal{A}$;*
2. *for $(\nu \otimes p)$-a.e. $(B, \bar{u})$,*

$$y_{x^*}(B,\bar{u}) \in \arg_{y \in \bar{S}^A:\ \Pi_{P^*}(y) \leq B} U(y;\bar{u});$$

3. *for every $c \in \Omega^*$,*

$$P^*_c > 0 \quad \Longrightarrow \quad \Phi_c(x^*) = b^c.$$

Proposition F.1. *Suppose Assumptions F.1 and F.2 hold. Then a stationary market-clearing equilibrium exists.*

The proof is relegated in Appendix F.2.1 and relies on a standard walrasian tatonnement idea.

Definition F.2. *An allocation $x^* \in C$ is dynamically C-constrained efficient if there is no allocation $x \in C$ such that, for every age $a = 1, \ldots, A$,*

$$U_a(y_x(B,\bar{u});\bar{u}) \geq U_a(y_{x^*}(B,\bar{u});\bar{u})$$

for $(\nu \otimes p)$-a.e. $(B,\bar{u})$, with strict inequality for some age a on a set of positive $(\nu \otimes p)$-measure.

Theorem F.1. *Suppose Assumption F.1 holds. Every stationary market-clearing equilibrium allocation is dynamically C-constrained efficient.*

The proof is provided in Appendix F.3.1. Since $A \subseteq C$, dynamic C-constrained efficiency of the equilibrium allocation implies its dynamic A-constrained efficiency.

F.1.2 Incomplete Information

This section considers the case in which future payoffs are not known at age 1. There are $A < \infty$ periods. Let again

$$\bar{S} := S \cup \{\varnothing\},$$

where $\varnothing$ is an outside option. Flow payoff vectors over schools belong to $[0,1]^S$, and we set $u_a(\varnothing) = 0$.

A newborn agent is endowed with a budget $B \in [0,1]$ and a first-period payoff vector $u_1 \in [0,1]^S$. The budget B is drawn according to an atomless probability measure ν on $[0,1]$, while u_1 is drawn according to a probability measure p_1 on $[0,1]^S$. We assume that B and u_1 are independent. For each $a = 1, \ldots, A-1$, conditional on the current payoff vector u_a, the next payoff vector $u_{a+1} \in [0,1]^S$ is drawn from a transition kernel

$$Q_a(du_{a+1} \mid u_a).$$

Transitions are assignment-independent: conditional on u_a, the distribution of u_{a+1} does not depend on the age-a assignment. Let $u_{1:A} := (u_1, \ldots, u_A)$ be the history of realised payoffs across an agent's lifetime. The induced law over full histories $(B, u_{1:A})$ is

$$m(dB, du_{1:A}) := \nu(dB) p_1(du_1) \prod_{a=1}^{A-1} Q_a(du_{a+1} \mid u_a).$$

Let $h_a := (u_1, \ldots, u_a)$ be the history of realised flow payoffs up to age a and let m_a denote the corresponding marginal law of (B, h_a).

A pure allocation is a collection of measurable maps[29]

$$x_a : [0,1] \times ([0,1]^S)^a \to \bar{S}, \qquad a = 1, \ldots, A.$$

As in the complete-information case, we can pre-process a primitive polytope into a set of constraints $\Omega^\star$ that depends on x only through the aggregate assignment, now defined with respect to the marginal laws

$$M_{a,s}(x) := \frac{1}{A} m_a(\{(B, h_a) : x_a(B, h_a) = s\})$$

For every $c \in \Omega^*$ and every generation-school pair $(a, s) \in \{1, \ldots, A\} \times \bar{S}$, let $\alpha^c_{a,s} \geq 0$, with $\alpha^c_{a,\varnothing} = 0$.[30] For an allocation x, define

$$\Phi_c(x) := \int \sum_{a=1}^{A} \alpha^c_{a, x_a(B,h_a)} \, m(dB, du_{1:A}).$$

As before,

$$C := \{x : \Phi_c(x) \leq b^c \text{ for every } c \in \Omega^*\},$$

and prices induced by constraint multipliers are

$$P_{a,s} := \sum_{c \in \Omega^*} \alpha^c_{a,s} P_c, \qquad s \in \bar{S}.$$

[29] Thus the age-a assignment may depend on the realized payoff history up to age a, but not on future payoff realizations.

[30] By a slight abuse of notation, we sometimes regard the outside option $\emptyset$ as a school.

Thus $P_{a,\varnothing} = 0$.

For a price vector P, define the feasible choices at age a given residual budget $R \in \mathbb{R}$ by

$$A_a(R; P) := \{s \in \bar{S} : P_{a,s} \leq R\}.$$

Since $P_{a,\varnothing} = 0$, the set $A_a(R; P)$ is nonempty whenever $R \geq 0$.

Define value functions recursively. At age A,

$$V_A^P(R, u_A) := \max_{s \in A_A(R;P)} u_A(s).$$

For $a = A - 1, \ldots, 1$, define

$$V_a^P(R, u_a) := \max_{s \in A_a(R;P)} \left\{ u_a(s) + \delta \int V_{a+1}^P(R - P_{a,s}, u_{a+1}) Q_a(du_{a+1} \mid u_a) \right\}.$$

The functions $R \mapsto V_a^P(R, u_a)$ are weakly increasing because a larger residual budget expands the set of feasible continuation choices.

Given an allocation x and prices P, define the residual budget generated by x at age a by

$$R_a^x(B, h_a; P) := B - \sum_{t=1}^{a-1} P_{t,x_t(B,h_t)}.$$

Definition F.3. *A pair (x^*, P^*) is a stationary market-clearing equilibrium if:*

1. *$x^* \in \mathcal{A}$;*
2. *for every $a = 1, \ldots, A$, and for m_a-a.e. (B, h_a), writing*

$$R_a^* := R_a^{x^*}(B, h_a; P^*),$$

if $a < A$, then

$$x_a^*(B, h_a) \in \operatorname*{arg\,max}_{s \in A_a(R_a^*; P^*)} \left\{ u_a(s) + \delta \int V_{a+1}^{P^*}(R_a^* - P_{a,s}^*, u_{a+1}) Q_a(du_{a+1} \mid u_a) \right\};$$

if $a = A$, then

$$x_A^*(B, h_A) \in \operatorname*{arg\,max}_{s \in A_A(R_A^*; P^*)} u_A(s).$$

3. *for every $c \in \Omega^*$,*

$$P_c^* > 0 \quad \Longrightarrow \quad \Phi_c(x^*) = b^c.$$

Similarly to the complete information case, we impose a few assumptions guaranteeing richness in the payoff distribution.

Assumption F.3. *For every age a and every pair of distinct schools $s, t \in \bar{S}$,*

$$m_a\big(\{(B, h_a) : u_a(s) = u_a(t)\}\big) = 0,$$

where we recall $u_a(\varnothing) = 0$.

Assumption F.4. *For every age $a < A$, for m_a-a.e. history (B, h_a), and for every $\eta > 0$,*

$$\Pr\left(\max_{t=a+1,\dots,A} \max_{s\in S} u_t(s) < \eta \;\middle|\; h_a\right) > 0.$$

The conditional probability is computed under the assignment-independent Markov law generated by the kernels $Q_a, \dots, Q_{A-1}$.

Remark F.1. *Assumptions F.3 and F.4 is implied by full support of the transition kernels on $[0,1]^S$ as well as of distribution p_1 over initial flow payoffs.*

Assumption F.5. *The aggregate demand functions*

$$D_c(P) := \int \sum_{a=1}^{A} \alpha^c_{a,d_a^P(B,h_a)}\, m(dB, du_{1:A}), \qquad c \in \Omega^*,$$

are continuous in P over its domain, where $d^P = (d_1^P, \dots, d_A^P)$ is a measurable policy selected from dynamic program

$$d_a^P(B, h_a) \in \operatorname*{arg\,max}_{s\in A_a(R_a,P)} \left\{ u_a(s) + \delta \int V_{a+1}^P(R_a - P_{a,s}, u_{a+1}) Q_a(du_{a+1} \mid u_a) \right\},$$

along with a deterministic tie-breaking rule over the finite set $\bar{S}$, used to break ties in individual demands.

Proposition F.2. *Suppose Assumptions F.3 and F.5 hold. Then a stationary market-clearing equilibrium exists.*

The proof is similar in spirit to the proof of Proposition F.1 and is relegated to Appendix F.2.2.

For a full history $\omega = (B, u_{1:A})$, define the realized continuation payoff from age a by

$$U_a(x; \omega) := \sum_{t=a}^{A} \delta^{t-a} u_t(x_t(B, h_t)).$$

Definition F.4. *An allocation $x^* \in C$ is ex-post dynamically C-constrained efficient if there is no allocation $x \in C$ such that, for every $a = 1, \dots, A$,*

$$U_a(x; \omega) \geq U_a(x^*; \omega) \quad \text{for } m\text{-a.e. } \omega,$$

with strict inequality for some age a on a set of positive m-measure.

We say that an allocation x weakly ex-post dominates x^* if the weak continuation-payoff inequalities in the preceding definition hold for every age, without imposing the strict-improvement requirement.

Theorem F.2. *Suppose Assumptions F.3 and F.4 hold. Every stationary market-clearing equilibrium allocation is ex-post dynamically C-constrained efficient.*

As in the complete-information case, ex-post C-constrained efficiency implies ex-post A-constrained efficiency. The proof is relegated in Appendix F.3.2 and relies on two lemmas for which we give intuition here. The first lemma establishes that ex-post dominance implies dominance on the flow payoffs at each age i.e. if x dominates $s^\star$, at any age a, the flow payoff of the age-a assignment is higher under x than under $x^\star$. Assumption F.4 indeed implies that for any mass of agents of i-th higher flow payoff under $x^\star$ at a, a positive fraction of them will then enjoy small enough payoffs under x in the future that the realised continuation payoffs from age a under x will end up being lower than under $x^\star$, a contradiction to ex-post dominance. The second lemma establishes that given a stationary market-clearing equilibrium $(x^\star, P^\star)$ any allocation that dominates $x^\star$ must have a higher price at any age under $P^\star$. Indeed, the previous lemma implies that any allocation that dominates $x^\star$ yields higher flow payoffs at any age. The optimality of $x^\star$[31] thus implies that any preferred school at a must be infeasible. Equipped with this result on prices, the proof then follows standard welfare theorem arguments.

F.1.3 Implementation via Deferred Acceptance

The above results guarantee existence of a constrained efficient allocation and provide a method to identify such an allocation. While the budget mechanism we described yields a constrained efficient allocation, it may not appear appealing to a policy maker committed to using the Deferred Acceptance mechanism. Drawing on our implementation principle, we now show that the constrained efficient allocation identified by the budget mechanism can be implemented using DA provided that scores are chosen appropriately, and we provide a methodology to systematically construct those scores.

The budget mechanism we described is naturally anonymous, time-invariant and aggregative. Also note that in this mechanism the message sent by the agent is the (sequence of) school in his feasible set to which he wants to be assigned. In equilibrium, the message exactly coincides with the allocation and the mechanism naturally conditions only the history of previous allocation through gradual depletion of the budget. The budget mechanism thus naturally belongs to the general class of mechanisms we identified as being implementable through market-clearing cutoffs provided initials scores and their evolution are well-designed.[32]

We provide below a methodology to construct initial scores and their evolution from the budget mechanism. Such a construction can be achieved by following the proof of the implementation theorem which is constructive. We provide here a different construction that relies on the specific structure of the budget mechanism to build scores that evolve only through additive school-age dependent bonuses in each period.

For an agent with realised budget $b \in [0,1]$ drawn from the atomless law ν, note $\mathcal{M}_b$ the menu containing all his feasible lifetime paths $y = (y_1, \cdots, y_A)$

$$\mathcal{M}_b := \left\{ y \in \bar{S}^A \mid \Pi_P(y) \leq b \right\}$$

.

For any age a, write $y_{<a} := (y_1, \cdots, y_{a-1}) \in \bar{S}^{a-1}$ the partial school path up to age a. The residual budget after $y_{<a}$ is

$$R_a(y_{<a}, b) \;:=\; b - \sum_{t<a} P_{t,y_t}$$

[31] By definition of stationary market clearing equilibrium, $x^\star$ solves the age-a dynamic program at the equilibrium residual budget.

[32] Note that the budget mechanism can exclude agents while maintaining vacant seats in schools if constraints other than the capacity constraints are binding and generate a positive price for all schools. Recall however that our implementation result holds for exclusive mechanism if we allow scores to be negative, which is the case here.

and the feasible choice set at age a given residual budget R is, as in the incomplete-information section,

$$A_a(R,P) := \big\{ s \in \bar{S} : P_{a,s} \le R \big\}.$$

We now specify a vector of candidate market-clearing cutoffs $(\kappa^s)_s \in [0,1]^S$, and for any agent with realised budget b in the budget mechanism, an initial score vector $\{r^{(1),b}\}_{b\in[0,1]}$ and a bonus $\beta_{s'}(s,a) \in \mathbb{R}$ for each $(s,s',a) \in \bar{S} \times \bar{S} \times \{1,\cdots,A-1\}$ such that, for any partial school path $y_{<a}$,

$$r_s^{(a),b} = r_s^{(1),b} + \sum_{t<a} \beta_s(y_t,t), \quad \forall s \in \bar{S}.$$

The score-feasible set at age a is $\{s \in \bar{S} : r_s^{(a),b} \ge \kappa^s\}$. A lifetime path y is *accessible* to budget b iff $r_{y_a}^{(a),b} \ge \kappa_{y_a}$ for all a, the scores being updated along y; write $\widetilde{\mathcal{M}}_b$ for the set of accessible paths.

Given the price matrix P, define, for every $b \in [0,1]$:

$$\begin{aligned} r_s^{(1),b} &:= b - P_{1,s}, \\ \beta_{s'}(s,a) &:= P_{a,s'} - P_{a+1,s'} - P_{a,s}, \\ \kappa^s &:= 0. \end{aligned}$$

The cutoff age- and school-invariant, and the bonuses do not depend on the budget b: the budget enters the mechanism only through the initial scores, and it does so additively.

Proposition F.3. *With the construction above, for every budget $b \in [0,1]$:*

(i) For every age $a \in \{1,\dots,A\}$ and every partial school path $y_{<a} \in \bar{S}^{a-1}$,

$$\big\{ s \in \bar{S} : r_s^{(a),b} \ge 0 \big\} \;=\; \mathcal{A}_a\big(R_a(y_{<a},b);\, P\big).$$

This holds for arbitrary prices $P_{a,s} \in \mathbb{R}$; nonnegativity is not used.

(ii) Since prices are nonnegative, the set of accessible lifetime paths is exactly the menu:

$$\widetilde{\mathcal{M}}_b \;=\; \mathcal{M}_b.$$

Statement (i) says that the score mechanism reproduces the feasible choice correspondence $\mathcal{A}_a(\,\cdot\,;P)$ at every node of the tree of school paths; statement (ii) is the induced equality of ex-ante menus. The proof is in the appendix. Note that in the present setting the two statements are equivalent.

Remark F.2. *Because prices are nonnegative and the outside option is free, any partial school path can be patched with the outside option, and this makes (i) and (ii) equivalent: in either mechanism, s is feasible at age a after $y_{<a}$ iff the padded path $(y_{<a}, s, \varnothing, \dots, \varnothing)$ belongs to the corresponding ex-ante set, since padding adds no further spending. Node-wise choice sets are therefore the sections of the ex-ante sets in both mechanisms.*

Proposition F.3 equates opportunity sets. The result below shows it implies the equality of induced allocations. Fixing a stationary market-clearing equilibrium $(x^\star, P^\star)$ and building scores as described, we can show the following.

Proposition F.4. *(a)* (Complete Information) *Suppose Assumptions F.1 and F.2 hold. For* $(\nu \otimes p)$*-a.e.* $(B, \bar{u})$*, the maximizer of* $U(\,\cdot\,; \bar{u})$ *over the score-accessible set* $\widetilde{\mathcal{M}}_B$ *is unique and equals* $y_{x^\star}(B, \bar{u})$*. Hence the score mechanism also induces the allocation* $x^\star$*.*

(b) (Incomplete Information) *Define value functions in the score mechanism by*

$$\tilde{V}_a(r, u_a) := \max_{s\,:\, r_s \geq \kappa^s} \Big\{ u_a(s) + \delta \int \tilde{V}_{a+1}\big(r + \beta_\cdot(s,a),\, u_{a+1}\big)\, Q_a(du_{a+1} \mid u_a) \Big\}.$$

, for $a = A, \ldots, 1$*, with the convention that* $\tilde{V}_{A+1} \equiv 0$*. Then for every age* a*, every partial school path* $y_{<a} \in \bar{S}^{a-1}$ *and budget* b *such that* $R_a(y_{<a}, b) \geq 0$*, and every* u_a*,*

$$\tilde{V}_a(r^{(a),b}, u_a) \;=\; V_a^{P^\star}(R(y_{<a,b}), u_a),$$

and the two programs have identical argmax correspondences at every such state. Fixing a deterministic tie-breaking rule used to break ties in individual demands, $x_a^\star(B, h_a)$ *is optimal in the score program at the corresponding score state for* m_a*-a.e.* (B, h_a)*. The score mechanism thus also induces the allocation* $x^\star$*.*

The proof is in Appendix F.4.

Since the allocations coincide under both mechanisms, all constraints in $\Omega^\star$ are satisfied under the score mechanism with cutoff vector $(\kappa^s)_s$. In particular, capacity constraints are satisfies and $(\kappa^s)_s$ is a vector of stationary market-clearing cutoffs.

We can then apply arguments similar to those used in Section 4 and the proof of Theorem 4.1. Since $(\kappa^s)_s$ is the zero vector, it is the smallest admissible cutoff vector and therefore it is the cutoff vector selected under agent-proposing DA. We can then construct a truthful stationary equilibrium by choosing for each type a report that ranks schools according to their dynamic choice-specific values at κ. The constructed reports and distributions form a stationary truthful equilibrium under repeated agent-proposing DA.

F.2 Existence proofs

F.2.1 Proof of Proposition F.1

Let $Y := \bar{S}^A$. For each $y \in Y$, write

$$A^c(y) := \sum_{a=1}^{A} \alpha^c_{a,y_a}.$$

For each $c \in \Omega^*$, if there exists a generation-school pair (a, s) such that $\alpha^c_{a,s} > 0$, let

$$\rho_c := \min\{\alpha^c_{a,s} : \alpha^c_{a,s} > 0\},$$

and choose $\bar{P}_c > 1/\rho_c$. If no such pair exists, set $\bar{P}_c = 1$. Define

$$\mathcal{P} := \prod_{c \in \Omega^*} [0, \bar{P}_c].$$

For $P \in \mathcal{P}$, define demand by

$$d_{B,\bar{u}}(P) \in \operatorname*{arg\,max}_{y \in Y:\ \Pi_P(y) \leq B} U(y; \bar{u}),$$

using an arbitrary measurable tie-breaking rule on the null set of ties. The choice set is always nonempty because the path consisting of always taking the outside option has price zero. Define

$$D_c(P) := \int A^c(d_{B,\bar{u}}(P))\,\nu(dB)p(d\bar{u}).$$

We claim that D_c is continuous. Let $P^n \to P$. Consider a type $(B, \bar{u})$ outside the null set at which either two distinct lifetime paths give the same lifetime utility or $B = \Pi_P(y)$ for some $y \in Y$. For such a type, the feasible set induced by P^n coincides with the feasible set induced by P for all sufficiently large n since Y is finite. Moreover, since no two distinct lifetime paths give the same lifetime utility, the maximizer over this feasible set is unique. Therefore,

$$d_{B,\bar{u}}(P^n) = d_{B,\bar{u}}(P)$$

for all sufficiently large n. The set just excluded has $(\nu \otimes p)$-measure zero by Assumptions F.1 and F.2: the utility-tie part has zero p-measure, and the budget-threshold part has zero ν-measure because ν is atomless. Since Y is finite, $A^c(d_{B,\bar{u}}(P^n))$ is bounded. Dominated convergence implies $D_c(P^n) \to D_c(P)$.

Let $z_c(P) := D_c(P) - b^c$ and define $F : \mathcal{P} \to \mathcal{P}$ by

$$F_c(P) := \min\left\{\max\{P_c + z_c(P), 0\}, \bar{P}_c\right\}.$$

The map F is continuous and maps the nonempty compact convex set $\mathcal{P}$ into itself. By Brouwer's fixed-point theorem, there exists $P^* \in \mathcal{P}$ such that $F(P^*) = P^*$.

Build x^* from $d(P^*)$. Individual optimality is immediate. It remains to show feasibility and that the constraint is binding when the cutoff is strictly positive. Fix $c \in \Omega^*$. If $P^*_c = \bar{P}_c$, then any school s with $\alpha^c_{a,s} > 0$ for agents of generation a has a cutoff $P_{a,s}$ lower bounded by $\bar{P}_c\alpha^c_{a,s} > 1$, since by construction, $P_c > 1 : \rho_c$. Since $B \leq 1$, no such school s can be chosen in any affordable lifetime path. Hence $D_c(P^*) = 0 < b^c$, so $F_c(P^*) < \bar{P}_c = P^*_c$, contradicting the fixed-point property. Therefore $P^*_c < \bar{P}_c$.

If $P^*_c > 0$, then $P^*_c \in (0, \bar{P}_c)$, and the fixed-point equation implies $z_c(P^*) = 0$, i.e. $D_c(P^*) = b^c$. If $P^*_c = 0$, the fixed-point equation implies $z_c(P^*) \leq 0$, i.e. $D_c(P^*) \leq b^c$. Since $D_c(P^*) = \Phi_c(x^*)$, we obtain that (x^*, P^*) binds when the cutoff is strictly positive.

The fixed-point argument establishes conditions 2 and 3 together with $\Phi_c(x^\star) \leq b^c$ for all c, i.e., by (82), $M(x^\star) \in \text{lcs}(\mathcal{A})$. It remains to verify feasibility, i.e., $x^\star \in A$. Suppose, toward a contradiction, that $x^\star \notin A$, i.e. $M(x^\star) \notin \mathcal{A}$. By definition of the lower contour set there exists $M' \in \mathcal{A}$ with $M(x^\star) \leq M'$ and, since $M(x^\star) \neq M'$, a pair $(a, s) \in \{1, \ldots, A\} \times S$ with $M_{a,s}(x^\star) < M'_{a,s}$.

First, we show that every constraint involving (a, s) must be free. Fix $c \in \Omega^\star$ with $\beta^c_{a,s} > 0$. Since $\beta^c \geq 0$ and $M(x^\star) \leq M(x')$ with a strict inequality at (a, s),

$$\Phi_c(x^\star) = \sum_{a',s'} \beta^c_{a',s'} Ma', s'(x^\star) < \sum_{a',s'} \beta^c_{a',s'} Ma', s'(x') = \Phi_c(x') \leq b^c$$

where the last inequality holds because $x' \in \mathcal{A} \subseteq C$. By the fact slack constraints have zero price, $P^\star_c = 0$ and since this holds for all c with $\beta^c_{a,s} > 0$, $P^\star_{a,s} = \sum_{c\in\Omega^\star} \alpha^c_{a,s} P^\star_c = \frac{1}{A}\beta^c_{a,s} P^\star_c = 0$.

Additionnaly, a positive mass of age-a agents is unassigned under $x^\star$. Since $M(x^\star) \leq M(x')$ with a strict inequality at (a, s)

$$\sum_{s'\in S} M_{a,s'}(x^\star) < \sum_{s'\in S} M_{a,s'}(x') \leq \frac{1}{A}.$$

As the mass of generation a is $\frac{1}{A}$, the set $H := \{(B, \bar{u}) : y_{x^\star,a}(B, \bar{u}) = \emptyset\}$ has positive measure.

We conclude by showing that unmatched agents have a profitable deviation, which contradicts the optimality of $x^\star$. Let G be the intersection of the full-measure sets on which condition 2 of Definition F.3 holds and on which no two distinct lifetime paths yield the same lifetime utility (Assumption F.1); then $(\nu \otimes p)(H \cap G) > 0$. Fix $(B, \bar{u}) \in H \cap G$ and write $y^* := y_{x^*}(B, \bar{u})$, so $y^*_a = \emptyset$. Define y' by $y'_t := y^*_t$ for $t \neq a$ and $y'_a := s$. Then, using $P^*_{a,\emptyset} = 0$ and Step 1,

$$\Pi_{P^*}(y') = \Pi_{P^*}(y^*) + P^*_{a,s} - P^*_{a,\emptyset} \;=\; \Pi_{P^*}(y^*) \;\leq\; B,$$

so y' is affordable. Moreover $U(y'; \bar{u}) - U(y^*; \bar{u}) = \delta^{a-1}\bar{u}_a(s) \geq 0$, and since $y' \neq y^*$ and $(B, \bar{u}) \in G$, the two lifetime utilities differ. Hence $U(y'; \bar{u}) > U(y^*; \bar{u})$. This contradicts the optimality of y^* (condition 2) on the positive-measure set $H \cap G$. Therefore $x^* \in \mathcal{A}$.

F.2.2 Proof of Proposition F.2

For each $c \in \Omega^*$, if there exists a generation-school pair (a, s) such that $\alpha^c_{a,s} > 0$, let

$$\rho_c := \min\{\alpha^c_{a,s} : \alpha^c_{a,s} > 0\},$$

and choose $\bar{P}_c > 1/\rho_c$. If no such pair exists, set $\bar{P}_c = 1$. Define

$$\mathcal{P} := \prod_{c \in \Omega^*} [0, \bar{P}_c].$$

For $P \in \mathcal{P}$, let $d^P = (d^P_1, \ldots, d^P_A)$ be the measurable policy selected from the dynamic program in Assumption F.5. Define

$$D_c(P) := \int \sum_{a=1}^{A} \alpha^c_{a, d^P_a(B, h_a)} \, m(dB, du_{1:A}).$$

By Assumption F.5, each D_c is continuous.

Let $z_c(P) := D_c(P) - b^c$ and define $F : \mathcal{P} \to \mathcal{P}$ by

$$F_c(P) := \min\left\{\max\{P_c + z_c(P), 0\}, \bar{P}_c\right\}.$$

The map F is continuous and maps $\mathcal{P}$ into itself. Hence, by Brouwer's fixed-point theorem, there exists $P^* \in \mathcal{P}$ such that $F(P^*) = P^*$. Let $x^* := d^{P^*}$. By construction, x^* satisfies the sequential optimality conditions in the definition of stationary market-clearing equilibrium.

Fix $c \in \Omega^*$. If $P^*_c = \bar{P}_c$, then for any generation-school pair with $\alpha^c_{a,s} > 0$, the cutoff of school s for generation a is at least $\bar{P}_c \alpha^c_{a,s} > 1$. So no such school can be chosen. Thus $D_c(P^*) = 0 < b^c$, so $F_c(P^*) < \bar{P}_c = P^*_c$, contradicting the fixed point. Hence $P^*_c < \bar{P}_c$.

If $P^*_c > 0$, then $P^*_c \in (0, \bar{P}_c)$, and the fixed-point equation implies $D_c(P^*) = b^c$. If $P^*_c = 0$, the fixed-point equation implies $D_c(P^*) \leq b^c$. Since $D_c(P^*) = \Phi_c(x^*)$, the allocation (x^*, P^*) satisfies positive price for binding constraints.

The fixed-point argument above establishes conditions 2 and 3 together with $\Phi_c(x^\star) \leq b^c$ for all c, i.e., by (82), $M(x^\star) \in \mathrm{lcs}(\mathcal{A})$. It remains to verify feasibility, i.e., $x^\star \in A$.

Write $M^* := M(x^*) \in \mathrm{lcs}(\mathcal{A})$ and suppose $M^* \notin \mathcal{A}$. As in the complete information case, there is (a, s) with $M^*_{a,s} < M'_{a,s}$ for some $M' \in \mathcal{A}$. Every $c \in \Omega^*$ with $\beta^c_{a,s} > 0$ satisfies $\Phi_c(x^*) < b^c$ and hence $P^*_c = 0$ by condition 3, so $P^*_{a,s} = 0$. Similarly, we can show the set $H_a := \{(B, h_a) : x^*_a(B, h_a) = \emptyset\}$ has $m_a(H_a) > 0$.

Let G_a be the intersection of: the full-m_a-measure set on which the sequential optimality condition 2 of the equilibrium definition holds; and the set $\{(B, h_a) : u_a(s') \neq u_a(t') \; \forall s' \neq t' \in \bar{S}\}$, which has full m_a-measure by Assumption F.3. Then $m_a(H_a \cap G_a) > 0$. Fix $(B, h_a) \in H_a \cap G_a$ and write $R_a^* := R_a^{x^*}(B, h_a; P^*)$ for the equilibrium residual budget. Since $x_a^*(B, h_a) = \emptyset$ was chosen and $P_{a,\emptyset}^* = 0$, feasibility of the equilibrium choice gives $R_a^* \geq 0$; since $P_{a,s}^* = 0$, school s also belongs to the feasible set $A_a(R_a^*; P^*)$.

Because $u_a(s) \neq u_a(\emptyset) = 0$ on G_a and $u_a(s) \geq 0$, we have $u_a(s) > 0$. If $a = A$, choosing s strictly improves the age-A objective, contradicting optimality. If $a < A$, compare the age-a objective at s and at $\emptyset$:

$$u_a(s) + \delta \int V_{a+1}^{P^*}(R_a^* - P_{a,s}^*, u_{a+1})\, Q_a(du_{a+1} \mid u_a) > u_a(\emptyset) + \delta \int V_{a+1}^{P^*}(R_a^* - P_{a,\emptyset}^*, u_{a+1})\, Q_a(du_{a+1} \mid u_a)$$

where the continuation terms are equal because $P_{a,s}^* = P_{a,\emptyset}^* = 0$ leaves the same residual budget R_a^*, and the kernel $Q_a(\cdot \mid u_a)$ is assignment-independent. This contradicts the sequential optimality of $x_a^*(B, h_a) = \emptyset$ (condition 2) on the positive-measure set $H_a \cap G_a$. Therefore $M^* \in \mathcal{A}$.

F.3 Constrained efficiency proofs

F.3.1 Proof of Theorem F.1

Let (x^*, P^*) be a stationary market-clearing equilibrium. Suppose, toward a contradiction, that there is $x \in C$ that dominates x^* in the sense of the preceding definition.

First, the improvement must be strict from the perspective of age 1 on a set of positive measure. If not, then

$$U(y_x(B, \bar{u}); \bar{u}) = U(y_{x^*}(B, \bar{u}); \bar{u})$$

for $(\nu \otimes p)$-a.e. $(B, \bar{u})$. By Assumption F.1, this implies

$$y_x(B, \bar{u}) = y_{x^*}(B, \bar{u})$$

for $(\nu \otimes p)$-a.e. $(B, \bar{u})$, and hence all continuation utilities are equal a.e., contradicting strict improvement at some age on a positive-measure set. Therefore

$$(\nu \otimes p)\left(\{(B, \bar{u}) : U(y_x(B, \bar{u}); \bar{u}) > U(y_{x^*}(B, \bar{u}); \bar{u})\}\right) > 0.$$

Now fix a type $(B, \bar{u})$ outside the null set of lifetime utility ties and outside the null set on which individual optimality of x^* may fail. If

$$U(y_x(B, \bar{u}); \bar{u}) > U(y_{x^*}(B, \bar{u}); \bar{u}),$$

then the path $y_x(B, \bar{u})$ cannot be affordable at prices P^*; otherwise it would contradict the optimality of $y_{x^*}(B, \bar{u})$. Hence

$$\Pi_{P^*}(y_x(B, \bar{u})) > B \geq \Pi_{P^*}(y_{x^*}(B, \bar{u})),$$

and therefore

$$\Pi_{P^*}(y_x(B, \bar{u})) > \Pi_{P^*}(y_{x^*}(B, \bar{u})).$$

If instead the two lifetime utilities are equal, since, by assumption, type $(B, \bar{u})$ is outside the null set of lifetime utility ties, the two lifetime paths coincide, and hence their prices are equal. Thus

$$\Pi_{P^*}(y_x(B, \bar{u})) \geq \Pi_{P^*}(y_{x^*}(B, \bar{u}))$$

for $(\nu \otimes p)$-a.e. $(B, \bar{u})$, with strict inequality on a set of positive $(\nu \otimes p)$-measure. Integrating,

$$\int \Pi_{P^*}(y_x(B,\bar{u}))\,\nu(dB)p(d\bar{u}) > \int \Pi_{P^*}(y_{x^*}(B,\bar{u}))\,\nu(dB)p(d\bar{u}).$$

To see how this inequality is rewritten in terms of the aggregate constraints, first note that, by the definition of generation-school prices,

$$\begin{aligned}
\int \Pi_{P^*}(y_x(B,\bar{u}))\,\nu(dB)p(d\bar{u}) &= \int \sum_{a=1}^{A} P^*_{a,y_{x,a}(B,\bar{u})}\,\nu(dB)p(d\bar{u}) \\
&= \int \sum_{a=1}^{A} \sum_{c\in\Omega^*} \alpha^c_{a,y_{x,a}(B,\bar{u})} P^*_c\,\nu(dB)p(d\bar{u}) \\
&= \sum_{c\in\Omega^*} P^*_c \int \sum_{a=1}^{A} \alpha^c_{a,y_{x,a}(B,\bar{u})}\,\nu(dB)p(d\bar{u}) \\
&= \sum_{c\in\Omega^*} P^*_c \Phi_c(x).
\end{aligned}$$

The same calculation applied to x^* gives

$$\int \Pi_{P^*}(y_{x^*}(B,\bar{u}))\,\nu(dB)p(d\bar{u}) = \sum_{c\in\Omega^*} P^*_c \Phi_c(x^*).$$

Hence the strict price inequality above is equivalent to

$$\sum_{c\in\Omega^*} P^*_c \Phi_c(x) > \sum_{c\in\Omega^*} P^*_c \Phi_c(x^*).$$

Therefore there exists $c \in \Omega^*$ such that

$$P^*_c > 0 \quad \text{and} \quad \Phi_c(x) > \Phi_c(x^*).$$

However, by definition of a stationary market-clearing equilibrium, $P^*_c > 0$ implies $\Phi_c(x^*) = b^c$. Hence $\Phi_c(x) > b^c$, which contradicts $x \in C$.

F.3.2 Proof of Theorem F.2

We start with the derivation of two useful lemmas.

Lemma F.2. *Suppose Assumptions F.3 and F.4 hold. Let (x^*, P^*) be a stationary market-clearing equilibrium. If an allocation x weakly ex-post dominates x^*, then for every age $a = 1, \ldots, A$,*

$$u_a(x_a(B,h_a)) \geq u_a(x^*_a(B,h_a))$$

for m_a-a.e. (B, h_a).

Proof. Fix an age a. If $a = A$, then the claim follows directly from $U_A(x;\omega) \geq U_A(x^*;\omega)$, since U_A is just the age-A flow payoff.

Now let $a < A$ and suppose, toward a contradiction, that the set

$$H := \{(B,h_a) : u_a(x^*_a(B,h_a)) - u_a(x_a(B,h_a)) > 0\}$$

has positive m_a-measure. Since

$$H = \bigcup_{n=1}^{\infty} \left\{ (B, h_a) : u_a(x_a^*(B, h_a)) - u_a(x_a(B, h_a)) \geq \frac{1}{n} \right\},$$

there exists $\gamma > 0$ such that

$$H_\gamma := \{(B, h_a) : u_a(x_a^*(B, h_a)) - u_a(x_a(B, h_a)) \geq \gamma\}$$

has positive m_a-measure.

For an age-a history $h_a = (u_1, \ldots, u_a)$, let $Q^a(\cdot \mid h_a)$ denote the conditional distribution of the future payoff path $(u_{a+1}, \ldots, u_A)$ generated by the transition kernels $Q_a, \ldots, Q_{A-1}$.[33]

Let

$$D_a := \sum_{k=1}^{A-a} \delta^k > 0$$

and choose $\eta > 0$ such that

$$D_a \eta < \gamma.$$

For each age-a history (B, h_a), define

$$E(B, h_a) := \left\{ u_{a+1:A} : \max_{t=a+1,\ldots,A} \max_{s \in S} u_t(s) < \eta \right\}.$$

By Assumption F.4,

$$Q^a(E(B, h_a) \mid h_a) > 0$$

for m_a-almost every (B, h_a).

Now let N_a be the set of full histories at which the age-a ex-post dominance inequality fails:

$$N_a := \{\omega : U_a(x; \omega) < U_a(x^*; \omega)\}.$$

Since x ex-post dominates x^*, we have

$$m(N_a) = 0.$$

For each age-a history (B, h_a), define the section of N_a after (B, h_a) by

$$N_a(B, h_a) := \{u_{a+1:A} : (B, h_a, u_{a+1:A}) \in N_a\}.$$

By disintegration,

$$0 = m(N_a) = \int Q^a(N_a(B, h_a) \mid h_a)\, m_a(dB, dh_a).$$

Hence

$$Q^a(N_a(B, h_a) \mid h_a) = 0$$

for m_a-almost every (B, h_a).

[33] Thus, if $E_{a+1}, \ldots, E_A$ are measurable subsets of $[0, 1]^S$, then

$$Q^a(E_{a+1} \times \cdots \times E_A \mid h_a) := \int_{E_{a+1}} \cdots \int_{E_A} Q_{A-1}(du_A \mid u_{A-1}) \cdots Q_a(du_{a+1} \mid u_a).$$

Define

$$G_a := \Big\{(B,h_a) : Q^a(N_a(B,h_a) \mid h_a) = 0\Big\} \cap \Big\{(B,h_a) : Q^a(E(B,h_a) \mid h_a) > 0\Big\}.$$

Both sets have full m_a-measure.[34] Therefore

$$m_a(H_\gamma \cap G_a) > 0.$$

Fix $(B,h_a) \in H_\gamma \cap G_a$. For every continuation $u_{a+1:A} \in E(B,h_a)$, we have, for each $t > a$,

$$u_t(x_t(B,h_t)) - u_t(x_t^*(B,h_t)) \le \eta.$$

Indeed, on $E(B,h_a)$, every school payoff is below η, while the outside option has value zero. Hence the payoff under x_t is at most η, and the payoff under x_t^* is at least zero.

Since $(B,h_a) \in H_\gamma$, the age-a flow loss under x is at least γ. Therefore, for every continuation $u_{a+1:A} \in E(B,h_a)$,

$$\begin{aligned} U_a(x;\omega) - U_a(x^*;\omega) &= u_a(x_a(B,h_a)) - u_a(x_a^*(B,h_a)) \\ &\qquad + \sum_{t=a+1}^{A} \delta^{t-a}\left[u_t(x_t(B,h_t)) - u_t(x_t^*(B,h_t))\right] \\ &\le -\gamma + \sum_{t=a+1}^{A} \delta^{t-a}\eta \\ &= -\gamma + D_a\eta \\ &< 0. \end{aligned}$$

Thus

$$E(B,h_a) \subseteq N_a(B,h_a).$$

But $(B,h_a) \in G_a$, so

$$Q^a(E(B,h_a) \mid h_a) > 0 \quad \text{and} \quad Q^a(N_a(B,h_a) \mid h_a) = 0,$$

which is impossible. This contradiction shows that H has m_a-measure zero. Since a was arbitrary, the result follows. □

Lemma F.3. *Suppose Assumptions F.3 and F.4 hold. Let (x^*, P^*) be a stationary market-clearing equilibrium. If an allocation x weakly ex-post dominates x^*, then, for every age $a = 1,\dots,A$,*

$$P^*_{a,x_a(B,h_a)} \ge P^*_{a,x_a^*(B,h_a)}$$

for m_a-a.e. (B,h_a). Moreover, whenever the age-a flow payoff is strictly higher under x, the corresponding price inequality is strict outside a null set.

Proof. Fix an age a. We first define the full-measure set of age-a histories on which the argument is carried out.

Let G_a^{opt} be the set of histories (B,h_a) such that the equilibrium assignment $x_a^*(B,h_a)$ solves the age-a dynamic program at the equilibrium residual budget. That is, writing

$$R_a^* := R_a^{x^*}(B,h_a;P^*),$$

[34] By Assumption F.4, applied successively from age $a+1$ to age A, the event $E(B,h_a)$ has strictly positive conditional probability for m_a-almost every (B,h_a).

we have

$$x_a^*(B,h_a) \in \arg\max_{r\in A_a(R_a^*;P^*)} \left\{ u_a(r) + \delta \int V_{a+1}\big(R_a^* - P_{a,r}^*, h_a, u_{a+1}\big) Q_a(du_{a+1} \mid u_a) \right\},$$

with the convention that the continuation term is omitted when $a = A$. By the definition of market-clearing equilibrium,

$$m_a(G_a^{\text{opt}}) = 1.$$

Let

$$G_a^{\text{ns}} := \{(B,h_a) : u_a(x_a(B,h_a)) \geq u_a(x_a^*(B,h_a))\}.$$

By Lemma F.2,

$$m_a(G_a^{\text{ns}}) = 1.$$

Finally, let

$$G_a^{\text{tie}} := \big\{(B,h_a) : u_a(r) \neq u_a(r') \text{ for every distinct } r, r' \in \bar{S}\big\}.$$

By Assumption F.3,

$$m_a(G_a^{\text{tie}}) = 1.$$

Define

$$G_a := G_a^{\text{opt}} \cap G_a^{\text{ns}} \cap G_a^{\text{tie}}.$$

Then

$$m_a(G_a) = 1.$$

Now fix $(B,h_a) \in G_a$. Write

$$s = x_a(B,h_a), \qquad s^* = x_a^*(B,h_a),$$

and

$$R_a^* := R_a^{x^*}(B,h_a;P^*).$$

Since $(B,h_a) \in G_a^{\text{ns}}$,

$$u_a(s) \geq u_a(s^*).$$

Suppose first that

$$u_a(s) > u_a(s^*)$$

and, toward a contradiction, that

$$P_{a,s}^* \leq P_{a,s^*}^*.$$

Since s^* is feasible at residual budget R_a^*, we have

$$P_{a,s^*}^* \leq R_a^*.$$

Therefore

$$P_{a,s}^* \leq P_{a,s^*}^* \leq R_a^*,$$

so s is also feasible at residual budget R_a^*. Moreover,

$$R_a^* - P_{a,s}^* \geq R_a^* - P_{a,s^*}^*.$$

Continuation values are weakly increasing in residual budget, because a larger residual budget expands the set of feasible future choices. Also, the transition kernel is assignment-independent,

so replacing s^* by s at age a does not change the law of future payoff vectors. Hence choosing s at age a yields weakly higher continuation value than choosing s^*. Since $u_a(s) > u_a(s^*)$, choosing s yields strictly higher total dynamic payoff than choosing s^*. This contradicts the optimality of s^* on G_a^{opt}. Thus, for every $(B, h_a) \in G_a$,

$$u_a(s) > u_a(s^*) \quad \Longrightarrow \quad P^*_{a,s} > P^*_{a,s^*}.$$

Now suppose that

$$u_a(s) = u_a(s^*).$$

Since $(B, h_a) \in G_a^{\text{tie}}$, equality of flow payoffs implies

$$s = s^*.$$

Hence

$$P^*_{a,s} = P^*_{a,s^*}.$$

Combining the two cases, for every $(B, h_a) \in G_a$,

$$P^*_{a,x_a(B,h_a)} \geq P^*_{a,x_a^*(B,h_a)}.$$

Since $m_a(G_a) = 1$, this inequality holds m_a-almost everywhere. Moreover, on the set of histories where

$$u_a(x_a(B, h_a)) > u_a(x_a^*(B, h_a)),$$

the inequality is strict:

$$P^*_{a,x_a(B,h_a)} > P^*_{a,x_a^*(B,h_a)}.$$

□

Proof of Theorem F.2

Proof. Let (x^*, P^*) be a sequential market-clearing equilibrium. Suppose, toward a contradiction, that there exists an allocation $x \in C$ that ex-post dominates x^*, with strict improvement at some age on a set of positive m-measure.

By Lemma F.3, for every age a,

$$P^*_{a,x_a(B,h_a)} \geq P^*_{a,x_a^*(B,h_a)}$$

for m-a.e. full history. The inequality is strict for some age on a set of positive m-measure.

Therefore

$$\int \sum_{a=1}^{A} P^*_{a,x_a(B,h_a)}\, m(dB, du_{1:A}) > \int \sum_{a=1}^{A} P^*_{a,x_a^*(B,h_a)}\, m(dB, du_{1:A}).$$

Using $P^*_{a,s} = \sum_{c\in\Omega^*} \alpha^c_{a,s} P^*_c$, this becomes

$$\sum_{c\in\Omega^*} P^*_c \Phi_c(x) > \sum_{c\in\Omega^*} P^*_c \Phi_c(x^*).$$

Therefore there exists $c \in \Omega^*$ such that

$$P^*_c > 0 \quad \text{and} \quad \Phi_c(x) > \Phi_c(x^*).$$

By definition of stationary market-clearing equilibrium, $\Phi_c(x^*) = b^c$. Hence $\Phi_c(x) > b^c$, which contradicts $x \in C$. □

F.4 Implementation Proofs

Lemma F.4. *For any budget $b \in [0,1]$, any age $a \in \{1,\dots,A\}$, and any partial school path $y_{<a} \in \bar{S}^{a-1}$,*

$$r_s^{(a),b} \;=\; b \;-\; \sum_{t<a} P_{t,y_t} \;-\; P_{a,s} \;=\; R_a(y_{<a}, b) - P_{a,s} \qquad \forall\, s \in \bar{S}.$$

Proof. By induction on a, for fixed b. The base case $a = 1$ holds by definition: $r_s^{(1),b} = \frac{1}{2} + b - P_{1,s}$ and $R_1 = b$. For the inductive step, append $y_a \in \bar{S}$ to $y_{<a}$:

$$\begin{aligned} r_{s'}^{(a+1),b} &= r_{s'}^{(a),b} + \beta_{s'}(y_a, a) \\ &= \big[\, b - \textstyle\sum_{t<a} P_{t,y_t} - P_{a,s'} \,\big] + \big[\, P_{a,s'} - P_{a+1,s'} - P_{a,y_a} \,\big] \\ &= b - \textstyle\sum_{t\le a} P_{t,y_t} - P_{a+1,s'}, \end{aligned}$$

as claimed. □

Proof of Proposition F.3. (i) Fix $(a, y_{<a}, b)$ and write $R := R_a(y_{<a}, b)$. By Lemma F.4, for any $s \in \bar{S}$,

$$r_s^{(a),b} \ge 0 \iff P_{a,s} \le R,$$

i.e. the score-feasible set equals $\mathcal{A}_a(R; P)$.

(ii) Fix $y \in \bar{S}^A$. By (i) applied along y, accessibility is equivalent to the running-budget conditions

$$\sum_{t\le a} P_{t,y_t} \;\le\; b \qquad \forall\, a \in \{1,\dots,A\}. \tag{$\star$}$$

Since prices are nonnegative, the partial sums are nondecreasing in a, so $(\star)$ is equivalent to the single terminal condition $\Pi_P(y) \le b$, i.e. $y \in \mathcal{M}_b$. Hence $\widetilde{\mathcal{M}}_b = \mathcal{M}_b$. □

Proof of Proposition F.4. (a) By Proposition F.3(ii), $\widetilde{\mathcal{M}}_B = \mathcal{M}_B$. Since $\bar{S}^A$ is finite, Assumption F.1 implies that outside a null set of payoff paths no two distinct lifetime paths yield the same lifetime payoff. Hence, outside this null set the maximizer of $U(\,\cdot\,; \bar{u})$ over a fixed choice set in $\bar{S}^A$ is unique. Since the menus are the same under the budget and the score mechanism, the allocations must be the same.

(b) We proceed by backward induction on a. At $a = A$: by Proposition F.3(i), the two programs maximize the same objective $u_A(\cdot)$ over the same feasible set, giving the identity $\tilde{V}_A(r^{(A),b}, u_A) = V_A^{P^\star}(R_A(y_{<A}, b), u_A)$ and the equality of the argmax sets, for every u_A.

Assume the identity at $a+1$. At age a, by Proposition F.3, $\{s : r_s^{(a),b} \ge \kappa^s\} = \mathcal{A}_a(R(y_{<a}, b); P^\star)$, so the two programs maximize over the same set. Second, for any feasible s, the updated score vector is

$$r^{(a),b} + \beta_\cdot(s, a) \;=\; \big(R(y_{<a,b}) - P^\star_{a,s}\big)\mathbf{1} - P^\star_{a+1,\cdot}.$$

Third, the payoff kernels $Q_a(\,\cdot \mid u_a)$ are common to the two programs and assignment-independent. Hence, applying the inductive hypothesis inside the integral,

$$\tilde{V}_{a+1}\big(r^{(a),b} + \beta_\cdot(s,a), u_{a+1}\big) Q_a(du_{a+1} \mid u_a) = \int V_{a+1}^{P^\star}\big(R - P^\star_{a,s}, u_{a+1}\big) Q_a(du_{a+1} \mid u_a)$$

for every feasible s. Hence the objective functions of the two maximization problems coincide term by term, giving both the value identity and equality of the argmax correspondences.

□

F.5 Pre-processing of lower bounds on average age

A policy maker seeking to reduce inequalities in the distribution of experience across schools may want to impose a lower bound on the average age at the worst-off school i.e. the school with the youngest population.

Define age weights

$$w_a = \frac{a}{A}, \quad \forall a \in \{1, \cdots, A\},$$

and the convention $w_0 = 0$. Given an allocation x and induced aggregate assignment M, define the average age at school s

$$z^s(M) = \frac{1}{q^s} \sum_a w_a M_{a,s}.$$

The designer is interested in ensuring

$$z^s \geq Z, \forall s \in S$$

for some exogenous Z, which is equivalent to imposing the bound at the worst-off school, since $\min_s z^s \geq Z$ if and only if $z^s \geq Z$ for every s.

This definition treats every school as if it were staffed to capacity: it normalizes age content by capacity rather than by enrollment. We adopt this definition rather than the enrollment-weighted average for the following reason. A school that leaves seats vacant must cover them with substitute teachers, who are typically less experienced than newly hired teachers; we adopt the convention that a substitute contributes zero age content, so that a vacant seat is filled by a teacher of age $a = 0$, carrying weight $w_0 = 0$. Because the denominator q^s is then fixed, a school cannot meet the floor by shrinking its workforce and retaining only its most senior teachers: every vacancy pulls its measured experience down.

The set of feasible aggregate assignments for this problem is

$$\begin{aligned} \mathcal{A} = \{\, M \in \mathbb{R}_+^{A \times S} : & \sum_{s \in S} M_{a,s} \leq \frac{1}{A}, \forall a, && \text{(generation feasibility)} \\ & \sum_{a=1}^{A} M_{a,s} \leq q^s, \forall s, && \text{(capacity)} \\ & z^s(M) \geq Z, \forall s \in S^f \Big\} && \text{(age floor)} \end{aligned}$$

where $S^f \subseteq S$, in particular, in the application discussed above $S^f = S$.

Since $\mathcal{A}$ contains floor contraints, we need to derive its lower contour set before computing the equilibrium.

Proposition F.5. *Let $S^f \subseteq S$ be the set of schools subject to an average age requirement. For a nonempty $T \subseteq S^f$ and a pivot $\rho \in [0, 1]$, define the ceiling inequality*

$$C(T, \rho) : \sum_{a=1}^{A} \sum_{s \in T} (\rho - w_a)_+ M_{a,s} + \sum_{a=1}^{A} \sum_{s \in S \setminus T} (w_a - \rho)_+ M_{a,s} \leq (\rho - Z) \sum_{s \in T} q^s + \sum_{a=1}^{A} \frac{1}{A} (w_a - \rho)_+$$

where $(k)_+ := \max\{k, 0\}$. *Then*

$$\begin{aligned}\text{lcs}(\mathcal{A}) = \Big\{ M \in \mathbb{R}_+^{A\times S} : &\sum_{s\in S} M_{a,s} \leq \frac{1}{A}, \forall a, \\ &\sum_{a=1}^{A} M_{a,s} \leq q^s, \forall s, \\ &C(T,\rho), \forall T \subseteq S^f, T \neq \{\emptyset\}, \forall \rho \in \{w_0, w_1, \cdots, w_A\} \Big\}\end{aligned} \tag{83}$$

Note that given a generation $\bar{a}$ and associated pivot $\rho = w_{\bar{a}}$, the family of constraints $(C(T,\rho))_{T,\rho}$ assigns a positive coefficient (and thus a positive price when the constraint is binding) to: (i) agents younger than $\bar{a}$ enrolling inside the floor set T, at rate $(\rho - w_a)_+$, the age deficit of generation a relative to generation $\bar{a}$, and (ii) agents older than $\bar{a}$ enrolling outside of the floor set T, at a rate $(w_a - \rho)_+$ corresponding to the age premium drained from T. In equilibrium, the personalised price of age a at school s, $\sum_c \alpha_{a,s}^c P_c$, is such that a young teacher pays extra at floor schools while an old teacher pays extra at schools outside the floor set T and faces low prices at deficient floor schools.

Proof. Let $\mathcal{K}$ be the right-hand side of Equation (83). We will show that $\text{lcs}(\mathcal{A}) = \mathcal{K}$. Given an aggregate assignment $M \in \mathbb{R}_+^{A\times S}$, write

$$\begin{aligned} g_a(M) &:= \frac{1}{A} - \sum_{s\in S} M_{a,s} \text{generation slack} \\ \kappa^s(M) &:= q^s - \sum_{a=1}^{A} M_{a,s} \text{capacity slack} \\ \sigma^s(M) &:= \sum_{a=1}^{A} w_a M_{a,s} - Zq^s \text{floor surplus} \end{aligned}$$

so that $M \in \mathcal{A}$ if and only if $M \geq 0$, $g_a \geq 0$ for all a, $\kappa^s \geq 0$ for all $sinS$ and $\sigma^s \geq 0$ for all $s \in S^f$. We call s in deficit if $\sigma^s < 0$ and write $d^s = -\sigma^s > 0$ and $D : \{s \in S^f : \sigma_s < 0\}$

Step 1. $\text{lcs}(\mathcal{A}) \subseteq \mathcal{K}$ Given an aggregate assignment M, define

$$\text{slack}_{C(T,\rho)}(M) = (\rho - Z)\sum_{s\in T} q^s + \sum_{a=1}^{A} \frac{1}{A}(w_a - \rho)_+ - \left(\sum_{a=1}^{A}\sum_{s\in T} (\rho - w_a)_+ M_{a,s} + \sum_{a=1}^{A}\sum_{s\in S\setminus T} (w_a - \rho)_+ M_{a,s} \right)$$

Notice that $(k)_+ = (k + (-k)_+)$ and rearrange

$$\begin{aligned}\text{slack}_{C(T,\rho)}(M) &= \sum_{s\in T}\left[\rho\left(q^s - \sum_{a=1}^{A} M_{a,s}\right) + \sum_{a=1}^{A} w_a M_{a,s} - Zq^s\right] + \sum_{a=1}^{A}(w_a - \rho)_+(\frac{1}{A} - \sum_{s\in S} M_{a,s}) \\ &= \sum_{s\in T}\left[\rho\kappa^s(M) + \sigma^s(M)\right] + \sum_{a=1}^{A}(w_a - \rho)_+ g_a(M) \end{aligned} \tag{84}$$

with the convention that $\text{slack}_{C(T,\rho)}(M) = \sum_{a=1}^{A}(w_a - \rho)_+ g_a(M)$ if T is empty. Let $M \in \mathcal{A}$. By definition of $\mathcal{A}$, $\kappa^s(M) \geq 0$, $\sigma^s(M) \geq 0$, $g_a(M) \geq 0$ and $\rho \geq 0$. Hence, $\text{slack}_{C(T,\rho)}(M) \geq 0$ and $M \in \mathcal{K}$ which implies $\mathcal{A} \subseteq \mathcal{K}$. Moreover, let $\tilde{M} \in \text{lcs}(\mathcal{A})$. By definition, there exists $M' \in \mathcal{A}$ such

that $\tilde{M} \leq M'$. By $\mathcal{A} \subseteq \mathcal{K}$, $M' \in \mathcal{K}$ and since $\mathcal{K}$ is defined by ceiling inequalities with nonnegative coefficients, any M'' with $0 \leq M'' \leq M'$ satisfies them as well and hence lies in $\mathcal{K}$. Hence $\tilde{M} \in \mathcal{K}$ yielding $\text{lcs}(\mathcal{A}) \subseteq \mathcal{K}$.

Step 2. $\mathcal{K} \subseteq \text{lcs}(\mathcal{A})$ We start by deriving the following useful identity. Fix $T \in S^f$. If every $s \in T$ is in deficit, 84 rearranges to

$$C(T,\rho) \text{ holds at M} \quad \Leftrightarrow \quad \sum_{s\in T} d^s \leq \rho \sum_{s\in T} \kappa^s + \sum_{a=1}^{A} g_a(w_a - \rho)_+. \tag{85}$$

$\mathcal{K}$ is a compact convex subset of the orthant (note that by nonnegativity and capacity constraints, we have $0 \leq M_{a,s} \leq q^s$, for all $M \in \mathcal{K}$ and all a, s). We will show that every maximal point of $\mathcal{K}$ lies in $\mathcal{A}$, where $M^* \in \mathcal{K}$ is maximal if no $M'' \in \mathcal{K}$ satisfies $M'' \geq M^*$, $M'' \neq M^*$. Fix $M \in \mathcal{K}$. The set $\{M' \in \mathcal{K} : M' \geq M\}$ is nonempty and compact, so the continuous map $M' \mapsto \sum_{a,s} M'_{a,s}$ attains its maximum on it at some $M^* \geq M$. Showing $M^* \in \mathcal{A}$ is sufficient to prove the claim: $M^* \in \mathcal{A}$ and $M \leq M^*$ give $M \in \text{lcs}(\mathcal{A})$ by definition of the lower contour set; as $M \in \mathcal{K}$ was arbitrary, $\mathcal{K} \subseteq \text{lcs}(\mathcal{A})$.

Let M be maximal and suppose, toward a contradiction that $M \notin \mathcal{A}$. Since $M \in \mathcal{K}$, M satisfies generation feasibility and capacity constraints. Then, $M \notin \mathcal{A}$ if and only if the set of schools in deficit, D, is not empty. We exhibit a coordinate $(\hat{a}, t)$, with $t \in D$ such that every inequality defining $\mathcal{K}$ whose coefficient at $(\hat{a}, t)$ is non zero, is strictly slack at M. Then, $M + \varepsilon e_{(\hat{a},t)} \in \mathcal{K}$ for ε small, where $e_{(\hat{a},t)}$ is the matrix with coordinate 1 at $(\hat{a}, t)$ and 0 elsewhere. We thus contradict the maximality of M.

Choice of $\hat{a}$. Some generation must have generation slack. Otherwise $g = 0$ and 85 applied to $T = D$ at pivot $\rho = w_0 = 0$ gives $\sum_{s\in D} d_s g = leq 0$, a contradiction with D being non empty. Let $\hat{a} := \max\{a : g_a > 0\}$ and remark that $g_a = 0$ for all $a > \hat{a}$.

The inequalities involving coordinate $(\hat{a}, t)$ are the generation feasibility constraint at $\hat{a}$, capacity constraint at t, and the elements of the family $(C(T,\rho))_{T,\rho}$ with either $t \in T$ and $\rho > w_{\hat{a}}$ or $t \notin T$ and $\rho < w_{\hat{a}}$. Note that members with $t \in T$ and $\rho \leq w_{\hat{a}}$ or $t \notin T$ and $\rho \geq w_{\hat{a}}$ have zero coefficient at $(\hat{a}, t)$ in $C(T,\rho)$.

Population feasibility at $\hat{a}$ is directly implied by the choice of $\hat{a}$.

Capacity slack at t. Since $M \in \mathcal{K}$, $C(\{t\}, w_{\hat{a}})$ must hold at M and by 85 it implies $d_t \leq w_{\hat{a}}\kappa^t + \sum_{a=1}^{A} g_a(w_a - w_{\hat{a}})_+ = w_{\hat{a}}\kappa^t$. As $t \in D$, $d_t > 0$ and $w_{\hat{a}} > 0$ since $\hat{a} \geq 1$ and so $\kappa^t > 0$.

C slack when $t \in T$ *and* $\rho > w_{\hat{a}}$. Let $T_D := T \cap D$. $t \in T_D$ since $t \in D$ and $t \in T$. Split 84

$$\text{slack}_{C(T,\rho)} = \sum_{s\in T\setminus T_D} (\rho\kappa^s + \sigma^s) + \sum_{s\in T_D} (\rho\kappa^s - d^s) + \sum_{a=1}^{A} (w_a - \rho)_+ g_a.$$

The first sum is nonnegative because $\rho \geq 0$, $\kappa^s \geq 0$, and schools outside D have $\sigma^s \geq 0$. The last sum only involves weakly positive terms and is weakly positive. By 85, at $(T_D, w_{\hat{a}})$, $\sum_{s\in T_D} d^s \leq w_{\hat{a}} \sum_{s\in T_D} \kappa^s$ and hence $\sum_{s\in T_D} (\rho\kappa^s - d^s) \geq \sum_{s\in T_D} (\rho - w_{\hat{a}})\kappa^s \geq 0$. If $\sum_{s\in T_D} (\rho - w_{\hat{a}})\kappa^s = 0$, then $\sum_{s\in T_D} \kappa^s = 0$ and we have $\sum_{s\in T_D} d^s \leq 0$, a contradiction. Hence $\sum_{s\in T_D} (\rho\kappa^s - d^s) > 0$ and so is $\text{slack}_{C(T,\rho)}$.

C slack when $t \notin T$ *and* $\rho < w_{\hat{a}}$. The above decomposition of slack applies and the first sum is nonnegative. Moreover, $\rho < w_{\hat{a}}$ implies $\sum_{a=1}^{A} (w_a - \rho)_+ g_a > 0$ since $(w_{\hat{a}} - \rho)_+ g_{\hat{a}} > 0$. However this slack could be compensated by a deficit at schools in T_D, i.e. by a negative second sum.

Call $(\tilde{T}, \tilde{\rho})$ a tight member of the family $(C(T,\rho))_{T,\rho}$ at M if $\tilde{\rho} < w_{\hat{a}}$ and $\text{slack}_{C(\tilde{T},\tilde{\rho})}(M) = 0$. If the set of tight members is empty, the slack of every member with $t \notin T$ and $\rho < w_{\hat{a}}$ is strictly

positive for any choice of $t \in D$, and we are done. Otherwise, we can choose t so that the slack is positive at $(\hat{a}, t)$[35].

The set of tight members is closed under intersections. Indeed, for any two tight members (T_1, ρ_1) and (T_2, ρ_2) with $\rho_2 \leq \rho_1$, remark that

$$\begin{aligned}
\text{slack}_{C(T_1 \cup T_2, \rho_2)} + \text{slack}_{C(T_1 \cap T_2, \rho_1)} &= \sum_{s \in T_1 \cup T_2} [\rho_2 \kappa^s + \sigma^s] + \sum_{a=1}^{A} (w_a - \rho_2)_+ g_a \\
&+ \sum_{s \in T_1 \cap T_2} [\rho_1 \kappa^s + \sigma^s] + \sum_{a=1}^{A} (w_a - \rho_1)_+ g_a \\
&= \text{slack}_{C(T_1, \rho_1)} + \text{slack}_{C(T_2, \rho_2)} + (\rho_2 - \rho_1) \sum_{s \in T_1 \setminus T_2} \kappa^s \\
&= (\rho_2 - \rho_1) \sum_{s \in T_1 \setminus T_2} \kappa^s \leq 0
\end{aligned}$$

where the first equality is the definition, the second comes from remarking that $T_1 \cup T_2$ is equal to the disjoint union of T_2 and $T_1 \setminus T_2$ and T_1 is equal to the disjoint union of $T_1 \setminus T_2$ and $T_1 \cap T_2$ (the σ- and g-sums recombine directly; the κ-sums recombine after adding and subtracting $\rho_1 \sum_{s \in T_1 \setminus T_2} \kappa^s$, which produces the correction term); the third equality comes from $\text{slack}_{C(T_1, \rho_1)}$ and $\text{slack}_{C(T_2, \rho_2)}$ being 0 since (T_1, ρ_1) and (T_2, ρ_2) are tight members and the last inequality comes from $\rho_2 \leq \rho_1$ and $\kappa^s \geq 0$.

Since at M, $M \in \mathcal{K}$ implies $\text{slack}_{C(T,\rho)} \geq 0$ for any T and admissible ρ, $\text{slack}_{C(T_1 \cup T_2, \rho_2)}$ and $\text{slack}_{C(T_1 \cap T_2, \rho_1)}$ are both weakly positive. Hence, it must be that both terms are null. In particular $T_1 \cap T_2 \neq \emptyset$: otherwise $0 = \text{slack}_{C(\emptyset, \rho_1)} \geq (w_{\hat{a}} - \rho_1)_+ g_{\hat{a}} > 0$, a contradiction. Hence $C(T_1 \cap T_2, \rho_1)$ is a tight member.

Hence, if the set of tight members is nonempty, denote by $\bar{T}$ the intersection of all tight members. Since there are finitely many tight members, $\bar{T}$ is a nonempty tight member. Let $t \in \bar{T} \cap D$. t then belongs to every tight member and when evaluating $slack_{C(T,\rho)}$ with $t \notin T$ and $\rho < w_{\hat{a}}$, we are guaranteed that (T, ρ) is not a tight member and so $slack_{C(T,\rho)} > 0$.

Hence, all constraints involving coordinate $(\hat{a}, t)$ are strictly slack which yields the desired contradiction. It follows that $\mathcal{K} \subseteq \text{lcs}(\mathcal{A})$, and thus $\mathcal{K} = \text{lcs}(\mathcal{A})$ which proves the claim.

□

F.6 Simulation design and additional simulation results

F.6.1 Demand Engine

Demand computation is made tractable by the assumption of logit utilities which provide closed-form expressions for demands and value functions, conditional on a remaining budget (Rust, 1987).

All agents share the same discount factor $\delta = 0.95$, so a type is summarized by initial budget. We however keep δ as an argument below because the engine supports heterogeneous discounting, and we note where the common δ specification simplifies the computation.

Given a residual budget R at age a and feasible choices $A_a(R, P)$ at a given this residual budget and prices $(P_{a,s})_{a,s}$, the continuation values are given by

$$V_a(R, P, \delta) := \log \sum_{s' \in A_a(R,P)} \exp(\bar{u}_a^{s'} + \delta V_{a+1}(R - P_{s',a}), P, \delta) + \gamma,$$

[35] Note that all other slacks hold for any $t \in D$, and since we only need to exhibit a coordinate $(\hat{a}, t)$ we can freely choose the most adequate t.

where $V_{A+1} = 0$ and γ is the Euler constant. Given continuation values, the probability that an agent demands school s at age a, conditional on having remaining budget R is defined by

$$\tilde{d}_a(R, P, \delta) := \frac{\exp(\bar{u}_{s,a} + \delta V_{a+1}(R - P_{s,a}, P, \delta))}{\sum_{s' \in A_a(R,P)} \exp(\bar{u}_{s',a} + \delta V_{a+1}(R - P_{s',a}, P, \delta))}.$$

Note that if all prices at some age are positive, all agents with initial budget lower than threshold $\underline{B} := \sum_a \min_s P_{s,a}$ cannot avoid eventual depletion of their budget and will end up unmatched at some point of their trajectory. Since $u_a(\emptyset) = -\infty$ for all a, all their continuation paths are worth $-\infty$ and the above theory does not predict their behavior. We thus adopt the convention that these agents choose schools myopically: they choose by softmax over current systematic utilities among currently affordable schools

$$\bar{d}_a^s(R, P, \delta) := \frac{\exp(\bar{u}_{s,a})}{\sum_{s' \in A_a(R,P)} \exp(\bar{u}_{s',a})}.$$

The unconditional demands $(d_a^s(B, P, \delta))_s$ of an agent born with initial budget b and discount factor δ depend on the distribution of remaining budgets induced by the agents' previous choices. A partial path of assignments $y_{<a} = (y_1, \cdots, y_{a-1}) \in S^{a-1}$ with remaining budget $R_{a'}$ at each $a' < a$ along the path is reached with probability

$$w_{y_{<a}}(B, P, \delta) = \prod_{a'<a} \tilde{d}_{a'}^{y_{a'}}(R_{a'}, P, \delta),$$

and the unconditional demand for school s at prices P from an agent of age a born with budget B and discount factor δ is

$$d_a^s(B, P, \delta) = \sum_{y_{<A} \in \mathcal{Y}_{<A}} w_{y_{<a}}(B, P, \delta) \tilde{d}_a^s(R_a, P, \delta).$$

Computing exact demands relies on two steps: a backward pass to compute the value functions V_a, and a forward pass to track the distribution of budget histories. Both become intractable when A and N are large. In the backward pass, evaluating V_a at a single (b, δ) requires evaluating V_{a+1} at up to N new budget points, each of which requires N further evaluations, and so on: the cost of an exact evaluation grows as $O(N^{A-a})$. This is an instance of Bellman's curse of dimensionality (Bellman, 1957; Rust, 1997). In the forward pass, the number of budget histories grows as N^{a-1}.

The backward pass is handled by grid-based memorization with interpolation (Keane and Wolpin, 1994; Eisenhauer, 2019). The value functions are precomputed by backward induction on grids $G_b \times G_d$ and stored as A matrices of shape $(n_b \times n_d)$, where G_b and G_d are the grids discretizing respectively the budget and discount factor spaces and n_b and n_d the number of points in each grid. Since all agents share $\delta = 0.95$ we take $n_d = 1$. The budget grid is non-uniform, with nodes concentrated near zero where optimal spending accumulates mass: spacing 5×10^{-6} below 0.02, coarsening to 1.1×10^{-3} near $B = 1$. Three nested grids are used: 1,741 nodes while searching for the active set, 3,188 while solving and 6,383 for evaluation and certification. At the terminal age, the value function is computed at each point of the grid from its expression. At each earlier age $a < A$, computing the value function at a with remaining budget R from the grid might require evaluating V_{a+1} at $R - P_a^s$ that is not on the grid. V_{a+1} is thus interpolated in b. Interpolation error is small, as in Keane and Wolpin (1994) and Eisenhauer (2019), who show that evaluating the value function on a subset of points and interpolating the rest produces negligible approximation errors in practice.

The engine then computes the expected demand directly, by integrating the Gumbel and history randomness out in closed form. The population is represented as a mass distribution $m[j,k]$ over a fixed grid of budget nodes b_j and discount nodes δ_k — a histogram of the type distribution, which the algorithm then iterates forward age by age. The initial masses are trapezoid quadrature weights for the uniform density on $[0,1]^2$: node masses are chosen so that sums over nodes approximate integrals against the uniform density to second order, i.e. each interior node j carries weight $w_j = (b_{j+1} - b_{j-1})/2$ (half the distance to each neighbor), endpoints half their single interval, and $m[j,k] = w_j^b w_k^\delta$, which sums to one. Given a vector of prices P, at each age a:

1. the choice probabilities are computed at every node from $\tilde{d}_a^s(b_j, P\delta_k)$ for agents whose budget is sufficient to be matched in each subsequent period ($b_j \geq \underline{B}$), and from $\bar{d}_a^s(b_j, P\delta_k)$ for agents whose budget will eventually leave them unmatched;

2. the node mass is split across schools in the proportions prescribed by choice probabilities;

3. mass choosing school s moves to $b_j - P_{s,a}$, which generally falls off the grid and is split linearly between the two bracketing nodes. This is the standard non-stochastic simulation method of Young 2010 for propagating distributions on grids, with one refinement: the flow of agents with initial budget high enough to remain matched throughout their careers and whose lower bracket would fall below $\underline{B}$ is assigned entirely to the upper node;

4. mass with an empty choice set stays put and is tallied as unmatched;

Note that the engine was certified against a Monte Carlo simulator with 10^6 agents at fixed price vectors and the demand vectors compared entrywise: at $P = 0$ the largest entrywise discrepancy is 6×10^{-5} and the mean 2×10^{-5}; at positive price vectors, mean discrepancies are $1 - 3 \times 10^{-4}$.

F.6.2 Approximation of low prices

On the grid, demand is a step function of prices. Demand changes only when a price crosses a budget node, and then demand jumps by the mass at that node. Over most of the price range this staircase is harmless. The mass per node is small, and the steps are far finer than the amount by which demand needs to be adjusted. Newton-type methods, used to solve for the equilibrium vector of prices, behave as if demand were smooth. The difficulty is concentrated near zero. Optimal spending drives a large mass of budgets toward exhaustion, so near-zero steps are macroscopic, and the price adjustments that would clear the cheapest schools are smaller than a grid cell. Meaningful prices are too small to be represented by the grid.

We use an approximation directed at solving this issue. We assume prices live in $\{0\} \cup [\tau, \infty)$ with $\tau = 10^{-4}$, i.e. we approximate prices lower than τ by 0. Zero-price schools thus cannot discriminate based on budget, and potential excess demand is rationed by lottery, with unplaced mass reallocated to remaining free capacity. The computed prices are certified below as an ε-equilibrium of the original economy, and all reported quantities are stable under doubling and quadrupling τ.

F.6.3 Equilibrium solver

Given an instrument and a bound, we obtain the set of constraints we need to work with. The equilibrium is a nonlinear complementarity problem in the $2N$ constraint prices. Each price is at its floor with a slack constraint, or above it with an exactly binding one. We solve it by a damped Josephy–Newton method with Broyden updates on a finite-difference Jacobian, and a discrete outer

search over the set of zero-price schools. Certification of a point requires, at an independent finer evaluation grid: capacity violations below 1% of a school's capacity, distributional violations below 0.2% of the capped share.

F.6.4 Certification

Internal residuals The first layer of certification consists in checking residuals as described above to ensure convergence of the solver. A point is reported as certified when, at the independent and finer evaluation grid, the capacity violation is below 1% of a school's capacity, the distribution violation is below the instrument-specific tolerance: the residuals on distributional constraints are expressed in enrollment share for the novice cap and years of experience for the ceiling on average age. The tolerance used for the novice cap is 5×10^{-4} and 2×10^{-2} for the ceiling on average age. Certification also requires that the mass left unplaced by the lottery be below 5×10^{-4}.

Distance to the exact equilibrium concept We evaluate unrationed demand at the computed prices. Every school with a positive price clears to within 0.9% of capacity. At the capacity-only benchmark, the mass allocated by the lottery rather than the equilibrium concept is 0.17% of the population.

Tick robustness We recompute the benchmark and the tightest constrained equilibrium at τ, 2τ and 4τ. The statistics used in the main text (dispersion of mean experience and its compression under a constraint) move by less than three percentage points across a fourfold change in the resolution, while the free-block size moves from five to eleven schools.

F.6.5 Additional simulation results and figures

A family of instruments The novice cap belongs to a family of age-weighted composition constraints,

$$\sum_a w_a \, D_{a,s} \;\le\; A\kappa \, q^s, \qquad \forall s, \tag{86}$$

indexed by the weights $(w_a)_a$ placed on successive generations. The novice cap is its step-function member, $w_a = \mathbf{1}\{a \le a_0\}$ and $b_s = A\kappa$: it distinguishes only two blocks of agents, those in their first a_0 years and everyone else. The graded member $w_a = a$ places weight proportional to experience and, since a full school satisfies $\sum_a D_{a,s} = Aq^s$, the constraint

$$\frac{1}{Aq^s} \sum_a a \, D_{a,s} \;\le\; \bar{\kappa}, \qquad \forall s,$$

is a ceiling on the average experience of every school. It is the age-weighted aggregate of one novice cap per generation, and it is the instrument we use below.

The ceiling on average experience We sweep $\bar{z}$ from 22.23 years, just below the capacity-only maximum, down to 21.16 years, in sixteen steps. Fifteen of the sixteen points are certified, and no point excludes any agent. Dispersion falls from 0.697 to 0.471 years, a compression of 32%, against the at most 2% that the novice cap achieves before it starts excluding entrants. The reallocation remains modest in ladder terms: at the tightest bound the average novice gains 0.28 rungs of the fifty-school ladder and the average agent in her last decade loses 0.30.

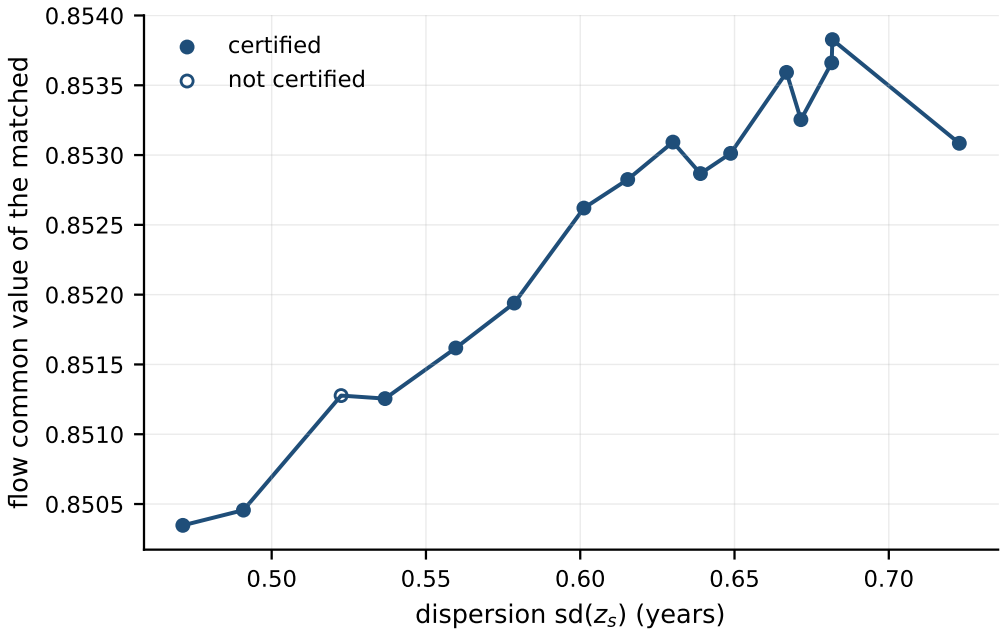


Figure 6: The equality–welfare locus for the ceiling on average experience. Each point is a constrained equilibrium; the horizontal axis is the capacity-weighted standard deviation of mean experience across schools and the vertical axis the flow common value of the matched population. Filled markers are certified in the sense of Appendix F.6.4, open markers are not and are excluded from every reported statistic.

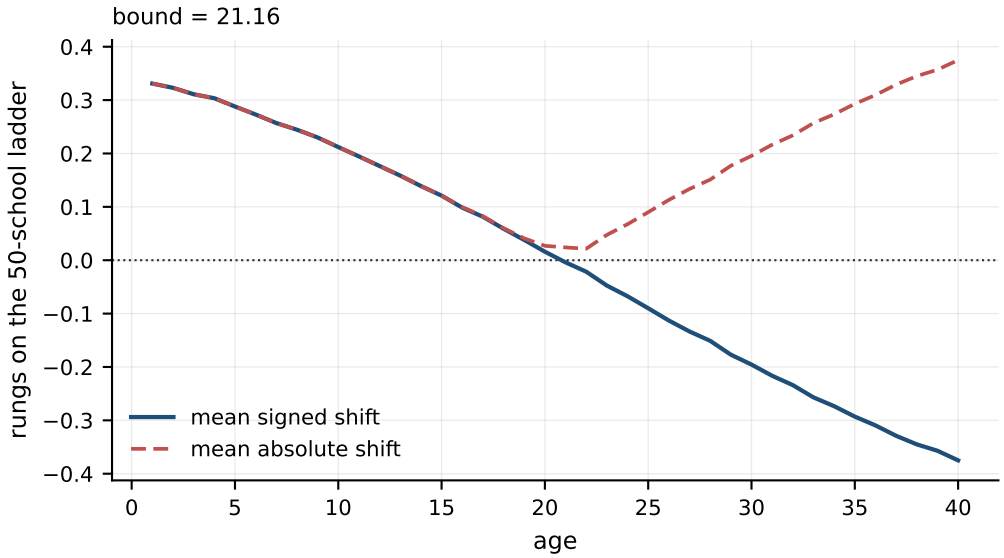


Figure 7: Displacement in ladder position by age at the tightest ceiling bound, relative to the capacity-only allocation. The mean signed shift is the average change in assigned school, in rungs of the fifty-school ladder, under the monotone coupling of the two allocations at each age; the mean absolute shift takes the same average over $|\cdot|$. The constraint transfers position from the oldest generations to the youngest, and does so by less than half a rung for every generation.

Exclusion under the novice cap Figure 8 reports, for three bounds below the exclusion threshold, the mass of agents left unmatched at each age. Exclusion is confined to the first a_0 generations: entrants deplete their budgets progressively through their novice years, and the mass unmatched rises until it reaches the doomed mass $\underline{B}$, after which it is exactly zero at every age above a_0. This is the pattern the theory predicts, since the novice cap places no distributional price for older age seats at the worst schools and the worst schools always remain free for these agents.

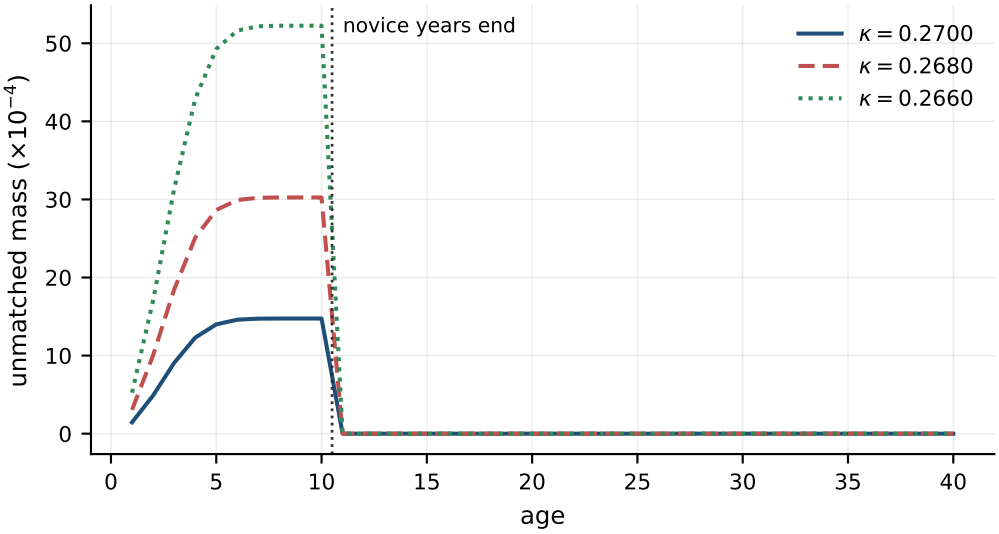


Figure 8: Unmatched mass by age under the novice cap, at three bounds below the exclusion threshold. The vertical line marks the end of the novice years.