Search and Matching for Adoption from Foster Care

Ludwig Dierks — UI Chicago
Nils Olberg — Zurich
Sven Seuken — Zurich
Vince W. Slaugh — Cornell
M. Utku Ünver — **Boston College**

Hong Kong University of Science and Technology
June 5, 2026

Children Waiting for Adoptive Families

- U.S. foster care serves about $525{,}000$ children each year.
- On Sept. 30, 2024: $70{,}421$ out of $328{,}947$ children in foster care were waiting for adoptive placements.
- About $20{,}000$ youth age out of the system each year without a permanent family.
- Persistent negative outcomes for those who age out: incarceration, homelessness, teen pregnancy (Triseliotis, 2002; Kushel et al., 2007; Gypen et al., 2017).

Where Search Is Actually Hard

- Most foster children leave quickly through reunification, relatives, or foster parents.
- The hard residual is children whose case goal is adoption *and* who have no relative or foster resource at TPR.
- For these children (Florida 2014–2021):
    - 16.1% adopted within 1 year of TPR,
    - 38.0% within 2 years,
    - 49.0% within 3 years.
- Florida alone spends $>\$20$M / year on adoption-search support for $\sim 3{,}800$ such children (Florida Department of Children and Families, 2022, 2024).

The Adoption Process

- A judge issues a Termination of Parental Rights (TPR) order.
- A caseworker is appointed to represent the child.
- Separately, prospective families complete a home study: training, references, home visit, written evaluation.
- Once approved, families enter the agency's pool.
- How the caseworker reaches them defines the search discipline.

Family-driven Search (FS): How It Works

The traditional approach used by most U.S. agencies.

- When a child c becomes eligible, the agency emails all approved families with a brief profile of c.
- Interested families respond by email — self-selecting into the case.
- The caseworker compiles the list of responders and reviews each one (home study, follow-ups, screening visits).
- If no family is suitable, the child is re-advertised after a waiting period.

FS in Florida: A Binding Legal Constraint

Florida Administrative Code, Rule 65C-16.003:

> ``Once the potential adoptive families have been identified, the staffing team will rate each family based on the family's ability to meet the identified needs of the child... The documentation must include a key of the rating scale used by the team.''

If the caseworker rejects a family for non-formal reasons, Rule 65C-16.005(9) requires review by a five-person Adoption Applicant Review Committee.

- Practical consequence. Every responding family must be considered, documented, and individually scored.

Caseworker-driven Search (CS): How It Works

An alternative some agencies have begun adopting (Riley, 2019; Schaefer Riley, 2019).

- No mass email announcement.
- For child c, the caseworker sequentially contacts specific families from the pool.
- Selection uses caseworker judgment, informal networks, and — optionally — a recommender tool.
- If a family expresses interest, the caseworker investigates (home-study review, interviews).
- If the eligible pool is exhausted with no match, the search is suspended until the pool refreshes.

The Rise of CS

- Florida's child-welfare system is split into 20 circuits, each administered by a non-profit Community-Based Care agency.
- One multi-county agency switched to CS on July 1, 2018, supported by a new technology platform.
- It was a managerial decision — the agency mandated CS for every case.
- Other circuits had access to the same platform but used it only occasionally, mostly for hard-to-place children.
- Caseworker engagement was real: average 25 family-profile views per child.
- This natural experiment will form our empirical evidence.

Why FS Strains the System

- A "cute" young child can draw hundreds of responses — each requiring documented review.
- Caseworkers carry roughly twice the recommended caseload (Yamatani et al., 2009; Lushin et al., 2023).
- Failure to respond to a family invites complaints to the governor's office or social-media backlash.
- Older children, sibling sets, and children with disabilities draw very few responses — and may quietly remain in the pool until aging out.

Families Feel It Too

From an adoptive father describing FS:

> ``You put your life on hold and have your hopes set on this one particular child... You dream every day and night that this kid is the one, and every time, 30-45 days later, I would find out I wasn't chosen.... I went through that over and over for about two years.''

Wasted search is paid by *both* sides:

- the caseworker pays document/screening costs on every responder;
- the family pays emotional and time costs on every application.

What This Paper Does

- Build a search-and-matching model that
  - is dynamic (entry/exit in a search-style model);
  - has two-sided heterogeneous preferences;
  - embeds a centralized search protocol (FS or CS);
  - allows uncertain compatibility of a candidate pair.
- Characterize Nash equilibria of the induced stochastic game in each regime.
- Derive welfare comparisons between FS and CS.
- Illustrate with numerics on $12{,}800$ instances.
- Close with empirical evidence from a Florida agency that switched its search discipline.

Overview of Results

- In each regime, equilibria exist and form a complete lattice.
- No FS equilibrium can Pareto-improve a CS equilibrium; the converse can hold strictly.
- When families are sufficiently impatient, all children weakly prefer CS.
- Adding family supply ($\lambda \uparrow$) helps children in CS — but can hurt them in FS.
- Numerics on $12{,}800$ instances: CS dominates almost everywhere in the parameter space.
- Florida agency that switched: $+44.9\%$ three-year adoption rate; $+54\%$ hazard ratio.

Primitives: Agents and Values

- Agents. $C = \{c_1, \ldots, c_n\}$ child types; $F = \{f_1, \ldots, f_m\}$ family types.
    - Each child is represented by a benevolent caseworker.
    - Sibling sets that should be placed together count as a single "child."
- Values. $v_c(f) \in \mathbb{R}$ and $v_f(c) \in \mathbb{R}$, strict on each side; outside option $= 0$.
- Compatibility. A specific pair of types (c, f) is suitable with probability $p \in (0, 1)$, i.i.d. across pairs and time.

Primitives: Time, Costs, and Arrivals

- Time & discounting. Discrete $t = 0, 1, \ldots;$ discount factors $\delta_C, \delta_F \in [0, 1)$.
- Search cost. Each evaluation costs $\kappa_C \geq 0$ for the child, $\kappa_F \geq 0$ for the family.
- Arrival. At each t: *exactly one* child active (uniform over C); each family type active independently with probability $\lambda \in (0, 1]$ — the market thickness parameter.
- Stationarity. The instance $(v, \delta_C, \delta_F, \kappa_C, \kappa_F, p, \lambda)$ is fixed.

Search Disciplines

We consider two search disciplines:

1. Family-driven Search (FS)
2. Caseworker-driven Search (CS)

Family-driven Search (FS)

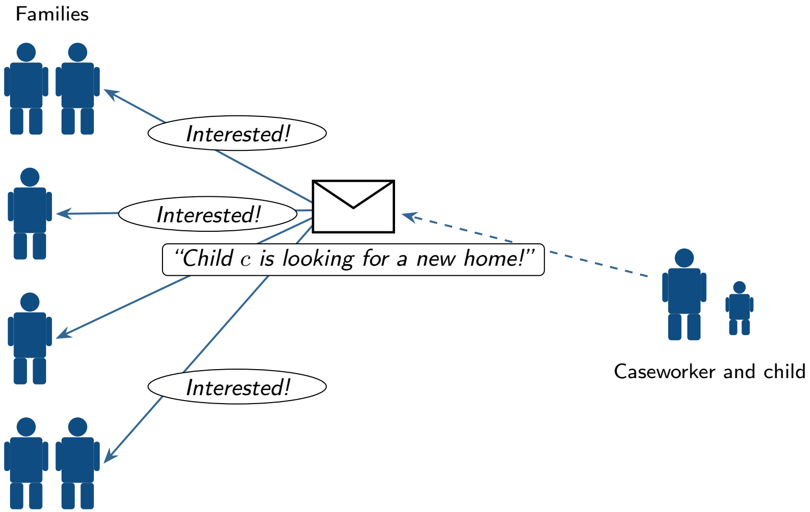

Broadcast: caseworker mass-emails the child's profile; *every* interested family writes back and must be evaluated.

Accounting for Costs in FS

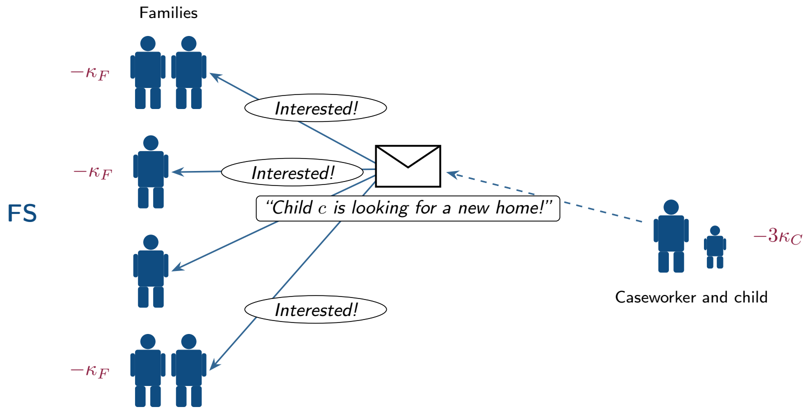

Cost accounting. Caseworker pays κ_C for *each* interested family she must evaluate; every interested family pays κ_F for the application — *whether or not* the pair turns out to be suitable.

Family-driven Search (FS)

At time t, with active child c chosen and active families realized:

1. Each *active, mutually interested* family f writes to c's caseworker.
2. Caseworker evaluates all of them simultaneously: each is suitable indep. with prob. p, costing κ_C (child) and κ_F (family).
3. c matches with the *highest-*v_c suitable family (if any); otherwise both wait one period.

- Key takeaway: search costs are incurred for every mutually interested pair, suitable or not.

Caseworker-driven Search (CS)

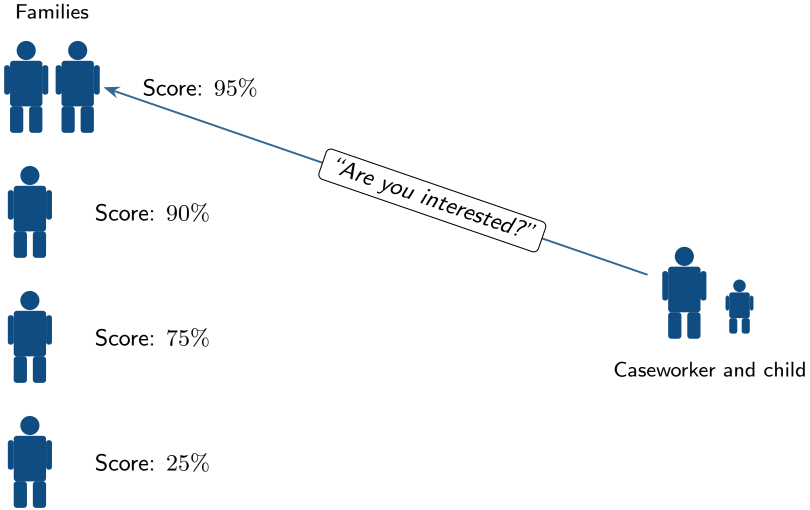

Sequential. Caseworker ranks the pool (e.g. via a recommender), contacts the *top* family first; only if she's not suitable, move on to the next.

Accounting for Costs in CS

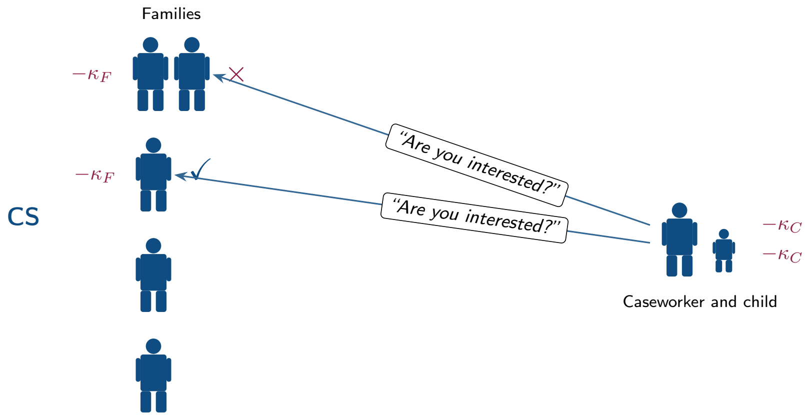

Cost accounting. Caseworker pays κ_C *per attempted candidate*, families pay κ_F *only if contacted* — search stops as soon as a suitable family is found.

Caseworker-driven Search (CS)

At time t, with active child c:

1. Caseworker enumerates *mutually interested, active* families in decreasing order of $v_c(f)$.
2. Investigate the first candidate f (costs κ_C and κ_F):
   - if suitable, match and exit;
   - if not, move to next candidate.
3. If list is exhausted with no match, both wait one period.

- Key takeaway: costs are only paid until a suitable match is found.
- Sequential, optimal stopping.

Timeline of One Period

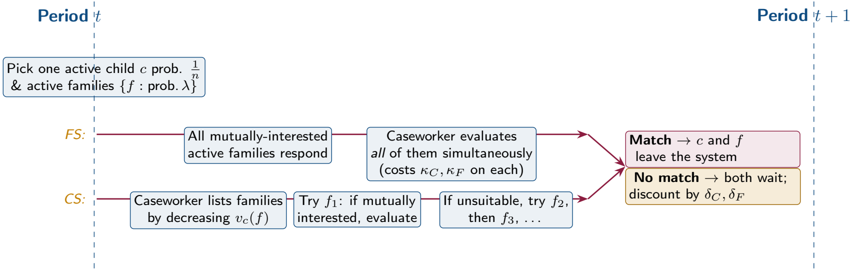

Key contrast. The two games share the same start and end; what differs is *which evaluations the caseworker performs* — and therefore *which costs are paid*.

Strategies

- A strategy for child c: $s_c \in \{0,1\}^m$, where $s_c(f) = 1$ iff c is *interested* in f.
- A strategy for family f: $s_f \in \{0,1\}^n$, where $s_f(c) = 1$ iff f is interested in c.
- Two agents are mutually interested iff $s_c(f) = s_f(c) = 1$.
- The mutual interest correspondence of s:

$$M(s) := \{(c,f) : s_c(f) = s_f(c) = 1\}.$$

- Agents condition on *types*, not on histories: stationary Markov strategies.

One-Period Expected Utilities

Let

$$b_{cf}(s) := |\{f' \in M_c(s) \colon v_c(f') > v_c(f)\}|,$$

and prob. c *does not* match someone better than f at this period

$$\beta_{cf}(s) := (1 - \lambda p)^{b_{cf}(s)}.$$

One-period expected utilities:

FS, child c:

$$\bar{u}_c^{\mathrm{FS}}(s) := \lambda \sum_{f \in M_c(s)} \Big(\beta_{cf}(s)\, p\, v_c(f) - \kappa_C \Big).$$

CS, child c:

$$\bar{u}_c^{\mathrm{CS}}(s) := \lambda \sum_{f \in M_c(s)} \beta_{cf}(s) \Big(p\, v_c(f) - \kappa_C \Big).$$

- Note. κ_C moves *inside* the β-weighting under CS but *outside* under FS.
- This single asymmetry drives all welfare differences.

Recursive (Bellman-type) Utilities

$u_c^{\mathrm{FS}}(s)$ is the unique fixed point of

$$u_c^{\mathrm{FS}}(s) = \delta_C u_c^{\mathrm{FS}}(s) + \underbrace{\lambda \sum_{f \in M_c(s)} \Big(\beta_{cf}(s)\, p\big(v_c(f) - \delta_C u_c^{\mathrm{FS}}(s)\big) - \kappa_C \Big)}_{= \bar{u}_c^{\mathrm{FS}}(s) - \delta_C u_c^{\mathrm{FS}}(s) \lambda p \sum_{f \in M_c(s)} \beta_{cf}(s)}.$$

$u_c^{\mathrm{CS}}(s)$ is the unique fixed point of

$$u_c^{\mathrm{CS}}(s) = \delta_C u_c^{\mathrm{CS}}(s) + \underbrace{\lambda \sum_{f \in M_c(s)} \beta_{cf}(s) \Big(p\big(v_c(f) - \delta_C u_c^{\mathrm{CS}}(s)\big) - \kappa_C \Big)}_{= \bar{u}_c^{\mathrm{CS}}(s) - \delta_C u_c^{\mathrm{CS}}(s) \lambda p \sum_{f \in M_c(s)} \beta_{cf}(s)}.$$

Analogous expressions hold for families (replace λ with $1/n$, δ_C with δ_F, etc.).

- Existence & uniqueness. Each map is a contraction in u.

Where the Two Regimes Diverge

Crucial difference, in one line:

In FS, search costs are paid whenever there is mutual interest.
In CS, they are paid only when prior attempts at t failed.

- This makes "expressing interest in a long-shot match" cheap under CS but expensive under FS.
- Under FS, agents trim their interest list to avoid wasted evaluations.
- Under CS, agents are willing to declare interest in more types — giving the market more match opportunities.

Tie-breaking Assumption

Assumption (Tie-breaking)

If i's utility weakly rises from mutual interest with j, then i is interested in j. If it strictly falls, i is not.

- Purpose. Excludes degenerate equilibria such as "nobody is interested in anybody," without distorting incentives away from utility maximization.

Threshold Strategies (CS)

Definition (CS threshold strategy)

A child of type c plays a CS-TS with threshold y_c if

$$s_c(f) = \mathbf{1}\left[\, p\big(v_c(f) - \delta_C y_c\big) \geq \kappa_C \right] \quad \forall f.$$

- A family of type f plays a CS-TS with threshold y_f analogously.
- Meaning. y_c is c's reservation utility: declare interest in f when the expected one-period gain over continuation $\delta_C y_c$ covers the cost κ_C.
- Analogously, y_f is f's reservation utility.

Threshold Strategies (FS)

Definition (FS threshold strategy)

A child of type c plays an FS-TS with threshold y_c if

$$s_c(f) = \mathbf{1}\left[\, \beta_{cf}(s)\, p\big(v_c(f) - \delta_C y_c\big) \geq \kappa_C \right] \quad \forall f.$$

- The extra factor $\beta_{cf}(s) = (1 - \lambda p)^{b_{cf}(s)}$ counts the families c prefers to f and that are mutually interested.
- A family of type f plays an FS-TS with threshold y_f analogously.
- Implication. Under FS, thresholds are *simultaneously determined* — one agent's interest in f depends on *others'* interest in higher-ranked families.
- y_i is i's reservation utility.

Comment on Threshold Strategies

- A simple reservation-utility threshold "accept iff $p\,v \geq y$" *fails* here: agents also weigh cost and (in FS) the probability another preferred match arrives at the same t.
- Under CS, the threshold is a function of $(\delta_C y_c)$ — one's continuation value.
- Under FS, the threshold is endogenous to others' strategies via $\beta_{cf}(s)$, so FS-thresholds are simultaneously determined.

Best Responses as Threshold Strategies

Proposition (threshold best responses)

For each agent i and any opponent profile s_{-i}, best response of i in FS (resp. CS) takes the threshold form with $y_i = u_i^*(s_{-i})$, i.e., i's best-response utility.

- Intuition. The marginal value of declaring interest in f equals

  (prob. of suitable match) $\times$ (value gain over continuation) $-$ (cost),

- which is monotone in $v_i(j) \Rightarrow$ accept exactly the top of one's preference list above a cutoff.

Example: Why Simple Thresholds Fail in FS

- Setup. 2 family types f_1, f_2 with $v_c(f_1) = 1$, $v_c(f_2) = 0.6$; common $p = \frac{1}{2}$, $\lambda = 1$, $\delta_C = 0.9$, $\kappa_C = 0.05$.
- CS: c keeps f_2 iff $p(v_c(f_2) - \delta_C y_c) \geq \kappa_C$ — cheap to declare interest, *independent* of whether f_1 is also active.
- FS: c keeps f_2 iff $\beta_{cf_2}(s)\, p(v_c(f_2) - \delta_C y_c) \geq \kappa_C$, where $\beta_{cf_2}(s) = (1 - \lambda p)^{|\{f' : v_c(f') > v_c(f_2),\ \text{mutual}\}|}$.
- Consequence. Whether c wants f_2 depends on whether f_1 is mutually interested — a strategic complementarity through β. A naive reservation-utility cutoff is not enough; one needs the full $\beta_{cf}(s)$-weighted threshold.

Existence and Lattice Structure

Define a partial order on threshold profiles:
$y \leq_C y'$ iff $y_c \leq y'_c \; \forall c \in C$ and $y_f \geq y'_f \; \forall f \in F$.

Theorem (equilibrium existence & lattice)

Both the FS-equilibrium-threshold-vector set and the CS-equilibrium-threshold-vector set are nonempty, and each is a complete lattice under $\leq_C$.

- How. Best-response operator T is monotone in $\leq_C$; apply Tarski (1955)'s fixed-point theorem (cf. Adachi, 2003).
- Two natural extreme equilibria in each regime:
  - child-optimal $s^{co\text{-FS}}, s^{co\text{-CS}}$;
  - family-optimal $s^{fo\text{-FS}}, s^{fo\text{-CS}}$.
- Resonance with Gale–Shapley: child-/family-optimality *within* each regime, exactly as in standard two-sided matching (Knuth, 1997).

An Operator Lemma

The operator $T : Y \to Y$ implicitly defined by mapping a threshold profile y to the best-response utility of every agent against $s(y)$ is

- monotone in $\leq_C$;
- maps the lattice $[0, \bar{v}]^{n+m}$ into itself.

Consequence. By Tarski's fixed-point theorem:

- iterating T from the $\leq_C$-min point $(y_c = 0, y_f = \bar{v})$ converges to the family-optimal equilibrium;
- iterating from the $\leq_C$-max point converges to the child-optimal equilibrium.
- this gives a constructive algorithm we use in the numerics.

Results

Lemma (1) (extra FS-only match)

Let $s^{\mathrm{FS}} \in S^{\mathrm{FS}}$, $s^{\mathrm{CS}} \in S^{\mathrm{CS}}$. *If* $(c, f) \in M(s^{\mathrm{FS}}) \setminus M(s^{\mathrm{CS}})$, *then either* $u_c(s^{\mathrm{FS}}) < u_c(s^{\mathrm{CS}})$ *or* $u_f(s^{\mathrm{FS}}) < u_f(s^{\mathrm{CS}})$.

- Why. An agent who declares interest in FS but not in CS has a *lower* reservation utility in FS than in CS; reservation utility = continuation u, so this agent is worse off.

Lemma (2) (when same-or-fewer matches)

If $M_c(s^{\mathrm{FS}}) \subseteq M_c(s^{\mathrm{CS}})$ *then* $u_c(s^{\mathrm{FS}}) \leq u_c(s^{\mathrm{CS}})$.

- Why. The child sees at least as many matching opportunities under CS, but pays strictly fewer search costs.

Main Theorem: CS vs. FS

Theorem (Pareto comparison)

(i) No FS equilibrium is a Pareto improvement over any CS equilibrium.
(ii) There exist instances in which every CS equilibrium Pareto-dominates every FS equilibrium.

Intuition.

- For (i): if anyone is better off in FS, by Lemma 1 someone else must have lowered their reservation utility — which by Lemma 2 means *that* agent is strictly worse off.
- For (ii): CS removes wasted-search costs, so agents can "afford" to be interested in more candidates; the market gets thicker on both sides simultaneously.

Skip the worse-off examples

A Single Child Can Be Worse Off Under CS

Proposition

There is an instance with unique FSE and unique CSE in which a child c is strictly worse off in CS.

Example. Two child types c_1, c_2 (all families prefer c_1); two family types f_1, f_2 (all children weakly prefer f_1).

- Under FS: f_2 pays too much to wait for c_1, so settles for c_2 – c_2 gets a match.
- Under CS: f_2 patiently waits for the rare c_1 slot (cheap to wait under CS); c_2 is left with no interested family.

A symmetric statement applies to families.

Impatient Families ⇒ Children Prefer CS

Proposition (impatient families)

For every instance there exists $\bar{\delta}_F < 1$ *such that, for all* $\delta_F \in [0, \bar{\delta}_F]$*,*

$$u_c(s^{\mathrm{FS}}) \leq u_c(s^{\mathrm{CS}}) \quad \forall c \in C,\ s^{\mathrm{FS}} \in S^{\mathrm{FS}},\ s^{\mathrm{CS}} \in S^{\mathrm{CS}}.$$

- Intuition. Impatient families are unwilling to wait for the perfect child under either regime; under FS they additionally bear repeated investigation costs. Their interest sets thus shrink *further* under FS, harming children.
- Asymmetry. The analogous result for families w.r.t. child patience *fails* — another family can "hog" the high-type child cheaply under CS.

Market Thickness λ

Let $s^{co\text{-CS},\lambda}$, $s^{co\text{-FS},\lambda}$ denote child-optimal equilibria parameterized by λ.

Proposition (λ in CS)

$\lambda \leq \lambda' \Rightarrow u_c(s^{co\text{-CS},\lambda}) \leq u_c(s^{co\text{-CS},\lambda'})$ *for all* c.

Proposition (λ in FS)

There exists an instance with $\lambda < \lambda'$ *and a child* c *s.t.* $u_c(s^{co\text{-FS},\lambda}) > u_c(s^{co\text{-FS},\lambda'})$.

- Why. In FS, adding a (worse) family raises competition for limited evaluation slots, lifting expected search costs for the families c likes most — they may drop out. In CS, no such crowding.
- Compare to Gale and Sotomayor (1985): in stable matching, adding agents on one side weakly helps the other. Here that result only survives in CS, and only for children.

Numerical Setup

Skip to numerical takeaways

- $n = m = 50$ types per side.
- Values: $v'_c(f) = \alpha\, q_f + (1-\alpha)\, \eta_c(f)$, $q_i \sim U[0,1]$, $\eta \sim U[0,1]$, then normalized.
  - $\alpha = 0$: idiosyncratic (horizontal) preferences.
  - $\alpha = 1$: common ranking (vertical) preferences.
- Parameters varied: $\alpha \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$, $\delta \in \{0.8, 0.9, 0.975, 0.99\}$, $\kappa \in \{0.01, 0.02, 0.05, 0.1\}$, $p = \lambda = 0.5$.
- $200 \times 4^3 = 12{,}800$ instances; $51{,}200$ equilibria.
- Solver: iterate T from the lattice extremes; tolerance ϵ.
- For $\sim 97\%$ of instances, FSE *unique* and CSE *unique*; FSE and CSE coincide in just one instance.

When Does CS Pareto-dominate FS?

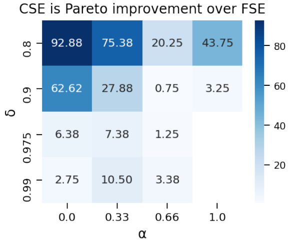

Two takeaways.

- CS Pareto-dominates more often when agents are impatient.
- CS rarely dominates when preferences are near-vertical ($\alpha \to 1$) and agents are *very* patient — the regime where some children gain from FS.

Who Is Better Off, Numerically?

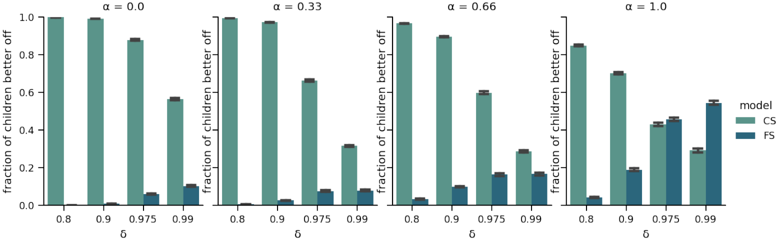

Fraction of children strictly better off under CS vs. FS in family-optimal equilibrium. CS wins for all $\delta \leq 0.9$, regardless of α. Only $\alpha = 1$, $\delta = 0.99$ tips the balance the other way. Families are almost always better off under CS (Appendix figure).

Match Values vs. Net Utility

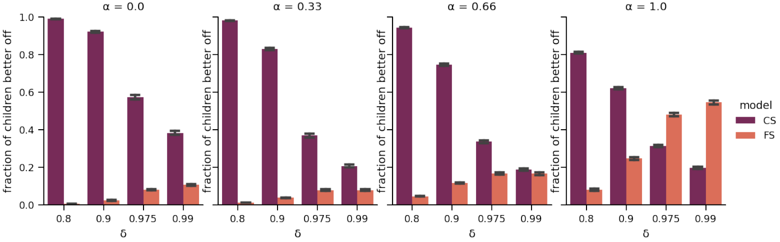

A child's match value = utility *net of* search costs (the welfare metric a child welfare agency may care about). CS doesn't only save costs; it also produces higher-quality matches in most parameter cells.

Numerical Takeaways

- Families are essentially always better off in CS, regardless of (α, δ, κ).
- Children are better off in CS unless preferences are near-vertical *and* agents are near-fully patient.
- These conditions are unlikely to bind in practice:
  - adoption agencies see significant idiosyncratic preference dispersion;
  - home-study certifications expire after ~ 1 year, putting an effective discount on families.
- Bottom line: CS is the welfare-preferred search technology in essentially all realistic parameter regions.

Related Literature: Matching and Search

- Two-sided matching (centralized). Gale and Shapley (1962); Roth (1984, 1991); Roth and Sotomayor (1992); Knuth (1997); Gale and Sotomayor (1985).
- Search with frictions. TU: Shimer and Smith (2000); Atakan (2006); NTU: Eeckhout (1999); Smith (2006); Adachi (2003); directed search: Lauermann et al. (2020); Cheremukhin et al. (2020); simultaneous vs. sequential: Stigler (1961); Weitzman (1979); Chade and Smith (2006); labor: Albrecht et al. (2006); Kircher (2009); platforms: Honka and Chintagunta (2017); Auster et al. (2025); dating: Hitsch et al. (2010); Lee and Niederle (2015); Kanoria and Saban (2021); surveys: Chade et al. (2017); Wright et al. (2021).
- Closest antecedents. Adachi (2003); Lauermann and Nöldeke (2014): single candidate per period; Immorlica et al. (2023): platform-guided search with symmetric values; Shi (2023): which side should drive the search?

Related Literature: Design and Applications

- Market design for vulnerable populations. Refugees (Andersson et al., 2018; Bansak et al., 2018; Delacrétaz et al., 2023); teachers (Combe et al., 2022, 2025); housing (Arnosti and Shi, 2020; Kawasaki et al., 2021; You et al., 2022); child welfare (Slaugh et al., 2016; Robinson-Cortés, 2019; MacDonald, 2019; Baccara et al., 2014; Slaugh et al., 2025).
- Dynamic matching. Congestion (Arnosti et al., 2021; Leshno, 2022); policy design (Ünver, 2010; Akbarpour et al., 2020b; Sönmez et al., 2020; Akbarpour et al., 2020a; Kerimov et al., 2023); platform design (Fradkin, 2017; Altinok and MacDonald, 2023; Dierks et al., 2026).

The Natural Experiment

- Florida is divided into 20 circuits, each administered by a nonprofit "CBC" agency.
- One CBC switched from FS to CS on July 1, 2018, adopting a technology platform supporting CS workflows.
- Other circuits continued primarily under FS, sometimes touching the platform for hard-to-place children ($< 4\%$ of statewide adoptions).
- Cleanly distinguishes the search discipline (managerial choice) from the technology platform (tool).
- Outcome of interest: time from TPR to finalized adoption.

Skip data details

Data

- AFCARS (statewide, federal): $9{,}544$ Florida children with TPR after Oct 2014, requiring adoptive search.
- Platform / focal-agency dataset: 335 children listed for CS by the focal agency before Oct 2021.
- Focal-agency children *also* appear in AFCARS (no record linkage) — any estimated effect understates the truth.
- Covariates available in both: sex, Black, Hispanic, clinical disability, age at TPR, age^2, FY of TPR.

Sample Characteristics

Focal agency vs. AFCARS sample:

	AFCARS ($N = 9{,}544$)	Platform ($N = 335$)
Case duration (yr, mean)	1.47	1.60
Adopted before horizon	27.2%	59.0%
Age at TPR (yr, mean)	7.71	8.39
Female	48.9%	42.3%
Black	36.1%	23.3%
Hispanic	14.6%	3.6%
Clinical disability	32.7%	37.6%

Focal-agency children skew older, more male, and more disabled — all of which make placement *harder*.

Empirical Strategy

Cox proportional hazards (Cox, 1972): hazard of adoption finalization at t since TPR,

$$h(t|X_i) = h_0(t)\exp\Big(\sum_k \beta_k X_{i,k}\Big).$$

Two specifications:

1. Benchmark (Model 1) — estimated on AFCARS only; *predicts* expected number of agency adoptions if agency cases followed statewide hazards.

2. Combined (Model 2) — AFCARS $\cup$ agency platform data; add $\mathit{FocalAgency}_i(t)$ — time-varying, turns on 3 months after platform listing.

Why time-varying? TPR date $\neq$ platform-listing date; we want to credit CS only for the post-listing period.

Benchmark: AFCARS-Predicted vs. Actual

- Predicted adoption probability at t for agency child i: $\tilde{\pi}(X_i, t, \tau_i^e, \tau_i^d)$ (conditioning on having survived until platform listing).
- Sum over agency children gives $\hat{\mu}(t)$ = expected adoptions by t.
- At the 2-year mark: 138 actual vs. 95.3 predicted — $+44.8\%$.
- At the 3-year mark: $+44.9\%$.
- Initial under-performance ($t < 1$ year) reflects the conservative 3-month delay built into τ_i^d.

Benchmark Comparison

Actual agency adoptions vs. AFCARS-predicted benchmark.

Dashed bands: 95% Poisson-binomial CI.

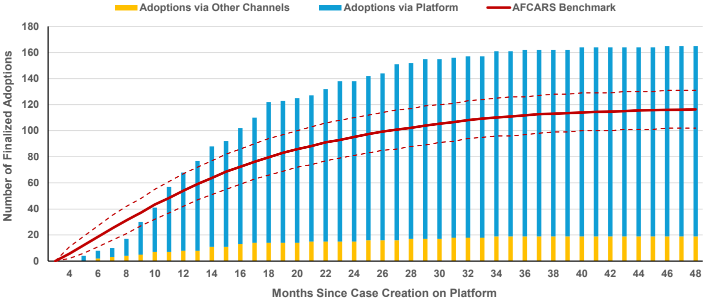

CS pulls ahead from ~ 12 months and stays ahead.

Combined-Sample Hazard: 54% Faster Adoption

Model 2 estimated on 9,544 (AFCARS) and 335 (agency) cases with a time-varying focal-agency indicator (lit 3 months post-listing):

	Model 1 (benchmark)	Model 2 (combined)
Female	~ 1.00 (ns)	~ 1.00 (ns)
Black	0.77^{***}	0.77^{***}
Hispanic	0.83^{***}	0.84^{***}
Age at TPR	0.89^{***}	0.89^{***}
Disability	0.88^{***}	0.87^{***}
FocalAgency	–	1.540^{***}
N	9,544	9,879

$+54\%$ adoption hazard for focal-agency cases, despite the double-counting bias.

Bias *Against* Finding an Effect

- Focal-agency children also sit *inside* the AFCARS dataset (no circuit linkage).
- About 3% of children in other circuits had CS-assisted matches via the platform but cannot be flagged — they show up as "FS" in the data.
- Some non-relative pre-identified placements remain in the benchmark, biasing it upward.
- The estimated $+54\%$ is therefore a lower bound on the true CS effect.

Connecting Theory to the Data

- Theory prediction. CS helps when preferences are heterogeneous and families are moderately impatient.
- Data check. Stated preferences on the platform vary widely (App. A); home-study certification expires after ~ 1 year (induced impatience).
- Pattern in benchmark figure. Early-period under-performance, later out-performance is consistent with CS supporting persistent search for hard-to-place children — which under FS would accumulate the highest wasted-search costs.
- Caveat. No circuit fixed effects available; cannot run a clean diff-in-diff. Empirics are *suggestive*, not causal.

Summary

- We model adoption from foster care as a dynamic search-and-matching game with two-sided heterogeneous preferences and a designed search technology.
- Equilibria exist, form a complete lattice, and are characterized by (non-simple) threshold strategies.
- Efficiency comparison theorem: FS equilibria cannot Pareto-dominate CS equilibria; instances exist where CS strictly dominates.
- Comparative statics: children weakly prefer CS when families are impatient; CS is monotone in market thickness; FS is not.
- Numerics: CS dominates outside the (vertical $\times$ very-patient) corner of parameter space.
- Florida empirics: agency that switched to CS shows $+44.9\%$ 3-year adoption and $+54\%$ hazard ratio.

Caveats and Open Questions

- Bias and fairness. CS concentrates discretion in caseworkers — a channel for explicit/implicit bias. FS broadens self-selection but may favor resourceful families.
- Causal identification. AFCARS lacks circuit-level identifiers; we cannot do diff-in-diff. Privacy laws prevent further casual analysis.
- Endogenous arrivals. Stationary arrival rates are a tractable approximation; long-run interactions between λ and recruitment are open.
- Recommender systems. Theory treats CS ordering as exogenous (decreasing v_c). Designing the recommender is the natural next step.
- Beyond binary suitability. A continuous suitability signal would change the optimal CS search order.

Appendix

Pareto Dominance: (κ, δ) Slice

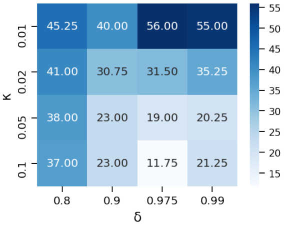

At higher search costs, CS Pareto-dominance is more frequent — consistent with the channel that CS is precisely the regime that economizes on wasted evaluation costs.

When Do FSE and CSE Coincide?

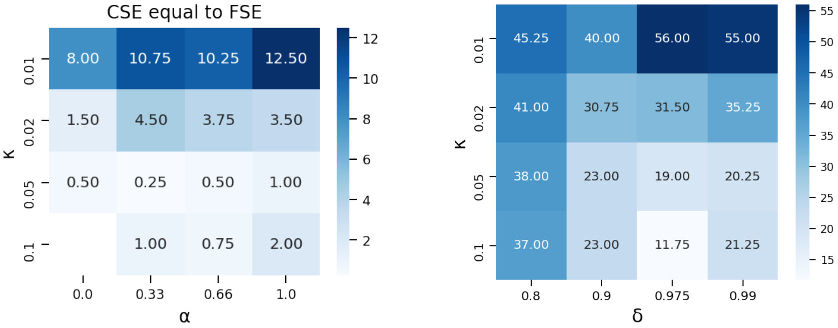

Coincidence is essentially never — the regimes induce different equilibrium correspondences in nearly every instance.

Families: Almost Always Better Off in CS

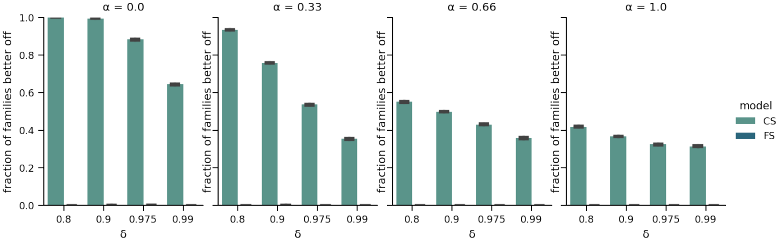

Average Child Utilities and Costs

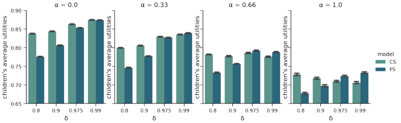

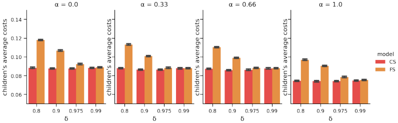

Match Probabilities

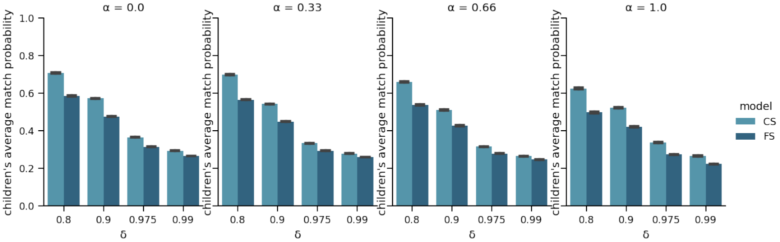

Match probabilities at each time step are higher in CS across (α, δ).

Robustness: Restricting the Black Variable

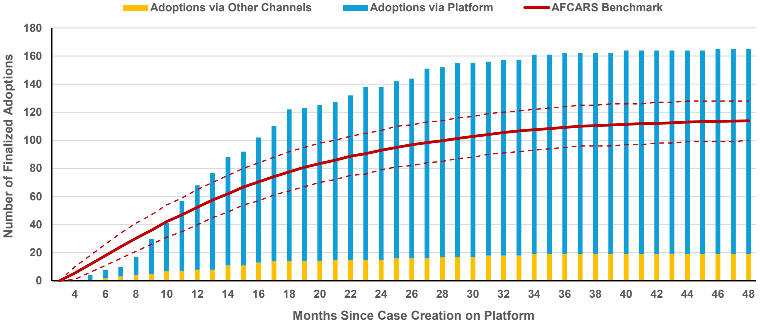

Robustness: Truncation and Other Channels

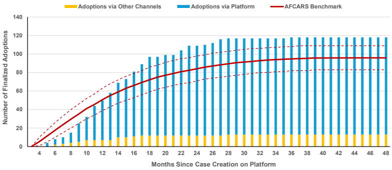

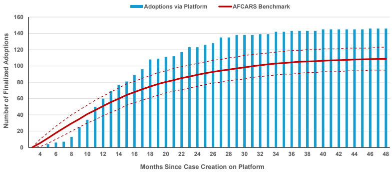

Cumulative Adoptions on Platform

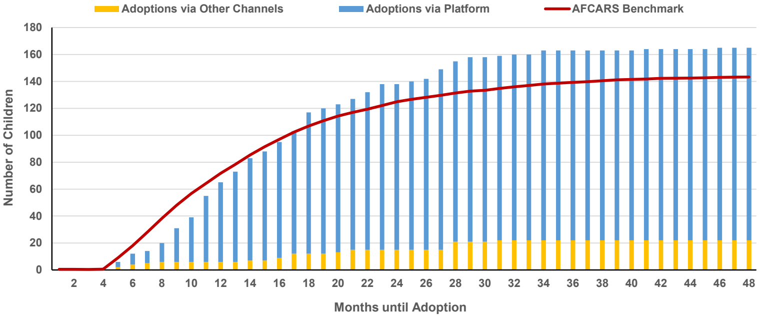

Family Preferences: Age Distribution

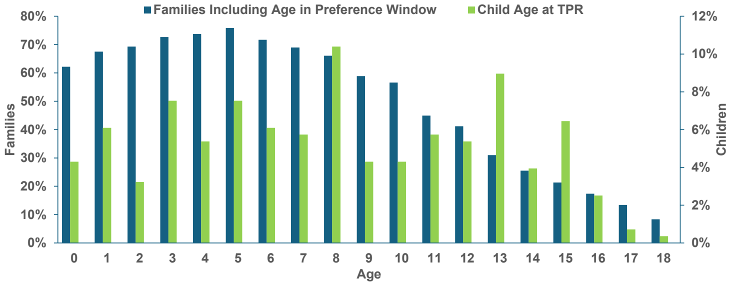

Evidence of idiosyncratic (not vertical) preferences — the regime where CS dominates theoretically.