

Search and Matching for Adoption from Foster Care

Ludwig Dierks

Department of Information and Decision Sciences, University of Illinois at Chicago, USA, ldierks@uic.edu

Nils Olberg

Department of Informatics, University of Zurich, Switzerland, olberg@ifi.uzh.ch

Sven Seuken

Department of Informatics, University of Zurich, Switzerland, seuken@ifi.uzh.ch

Vincent W. Slaughter

SC Johnson College of Business, Cornell University, USA, vslaughter@cornell.edu

M. Utku Ünver

Department of Economics, Boston College, USA, unv@bc.edu

To find families for the more than 100,000 children in need of adoptive placements, most United States child welfare agencies have employed a *family-driven search* approach in which prospective families respond to announcements made by the agency. However, some agencies have switched to a *caseworker-driven search* approach in which the caseworker directly contacts families recommended for a child. We introduce a novel search-and-matching model that captures the key features of the adoption process and compare family-driven with caseworker-driven search in a game-theoretical framework. Under either approach, the equilibria are generated by threshold strategies and form a lattice structure. Our main theoretical finding then shows that no family-driven equilibrium can Pareto dominate any caseworker-driven outcome, whereas it is possible that each caseworker-driven equilibrium Pareto dominates every equilibrium attainable under family-driven search. We also find that when families are sufficiently impatient, caseworker-driven search is better for all children. We illustrate numerically that most agents are better off under caseworker-driven search for a wide range of parameters. Finally, we provide empirical evidence from an agency that switched to caseworker-driven search and achieved a three-year adoption probability that outperformed a statewide benchmark by 24%, as well as a statistically significant 27% improvement in adoption hazard rates.

Key words: child adoption, search and matching, market design, game theory

Subject classifications: Government/Services: Child-Welfare System; Government/Programs: Children, Youth, and Families; Population/Family planning: Adoption

1. Introduction

Child welfare systems worldwide face the challenge of finding families for children in need of adoption. For example, the United States foster care system serves over 525,000 children annually. While the goal for most of these children is to reunite them with their parents or relatives, 77,809

out of 343,077 children in foster care on September 30, 2023, were waiting for permanent adoptive placements (U.S. Administration for Children and Families 2024). Finding an adoptive family for these children has become a public policy priority due to high levels of incarceration, homelessness, unemployment, and teen pregnancy observed in the population of children “aging out” of the child welfare system without a permanent family relationship (Triseliotis 2002, Kushel et al. 2007, Gypen et al. 2017). In this paper, we study the search and matching process for children waiting in the child welfare system for adoptive placement and compare two prominent search methods that differ in whether families or the children’s caseworkers drive the search process. We provide structural insights by analytically and numerically analyzing a game-theoretic model, and empirically validate our findings by examining the outcomes achieved by a Florida child welfare agency that switched its search strategy in 2018.

The search for an adoptive placement officially begins once a judge issues a termination of parental rights order. A *caseworker* represents the child’s interests throughout this process to find an adoptive family if adoption by a relative or foster care family is not possible. Identifying a family willing to adopt the child and capable of caring for the child can be a difficult task, and the challenge varies greatly according to the child’s demographic characteristics and special needs. While relatively little research has studied best practices for search operations — either through empirical studies of how caseworkers find families or prescriptive studies for how search should be conducted — states invest significant resources in trying to help the most vulnerable children find permanency; for example, the Florida Department of Children and Families (2022, 2024) reports spending over \$20 million annually to promote and support searches for approximately 3,800 children with a goal of adoption at any point in time.

Different states, counties, and agencies adhere to different paradigms for the practice of search and matching. Families first register with an adoption agency and provide a home study evaluation. In the predominant approach, caseworkers then announce children via email to a set of registered families, each of whom has the opportunity to express interest in a child. The caseworker receives these inquiries and works to identify the family that best fits the child’s needs. We label this approach *family-driven search*, as families direct the search process by expressing interest in available children. Hanna and McRoy (2011) demonstrates how the child welfare literature has almost exclusively focused on what we view as family-driven search processes. One important downside of this search process is that some children may attract hundreds of interested families, all of which the caseworker has to consider simultaneously.

These interactions can be time-consuming and emotionally stressful for all parties involved, and some states require agencies to engage with every family that has expressed interest. For example, Florida Administrative Code Rule 65C-16.003 states:

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 based on information documented in [the Florida Safe Families Network information system], the Child Study and the adoptive parent's home study. The documentation must include a key of the rating scale used by the team.

An adoptions manager for a multi-county agency in Florida directed us to this rule to emphasize the imperative that the agency must respond to every family that expresses interest in a child. Failure to respond to families can result in complaints to the governor's office or negative comments on social media. The manager also commented that the agency dreaded announcing the availability of a "cute" young child, who would attract dozens or even hundreds of responses from families that required thorough consideration and individual responses. This is especially problematic because caseworkers often deal with twice as many cases as they should ideally handle (Yamatani et al. 2009), an issue intensifying since the COVID-19 pandemic (Lushin et al. 2023). On the other hand, other children may attract very few interested families, and nearly 20,000 children age out of foster care each year (Children's Bureau 2022).

We have also heard families share their frustration with the emotional and time costs that they incurred while participating in a family-driven search system. In a video shared on social media, one adoptive father described his journey that began as a prospective parent engaged with an adoption agency that announced available children and had families respond to express interest in particular children:¹

Basically, you put your life on hold, and you have your hopes set on this one particular child that you've fallen in love with their profile and their picture. You dream every day and night and go to sleep hoping that this kid is the one, and every time that happens for me, 30-45 days later I would find out that I wasn't chosen. So, there's this tremendous sense of disappointment, rejection, and "why didn't they choose me?" I went through that process over and over and over again for about 2+ years until I was finally connected with [a caseworker-driven search platform]. What I loved about the concept was that I wouldn't have to go through that process. This time it was my profile, and I'd just have to wait for the caseworker or my forever match to find me. I'm happy to report that a couple weeks later I did find my forever match, and I have my son now.

In reaction to these challenges, some child welfare agencies have recently sought to improve outcomes by switching to *caseworker-driven search* (Riley 2019). In this approach, caseworkers sequentially contact specific prospective families to share details about the child. Caseworkers use

¹ https://www.linkedin.com/posts/adoption-share_by-inverting-the-outreach-for-adoption-matching-activity-7131265613947138048-YCM1 (Retrieved 06/17/2025)

their expertise to decide which families to contact based on the child’s and the families’ characteristics. Optionally, technological tools may also aid their decisions. This approach removes the burden of engaging a large number of families simultaneously and allows caseworkers to target compatible families for children with very specific needs.

However, even though these are clear advantages of caseworker-driven search, it is important to note that changing the search paradigm may change families’ incentives that influence their interest in different types of children. For a match to occur under either approach, both the caseworker and the prospective family need to agree. This can, for example, mean that if caseworker-driven search leads to a reduction in search costs for families, they may stop being interested in certain children, leaving those children unmatched. As such, it is unclear whether caseworker-driven search leads to more desirable outcomes for the entire population of children requiring adoptive placements.

To assess how different search disciplines affect outcomes, we pursue two complementary approaches in this paper. First, we analyze the strategic behavior of agents under both search disciplines in a game-theoretic search-and-matching model. Second, we provide empirical evidence for the performance of caseworker-driven search compared to classic search approaches by studying outcomes of a real-world child welfare agency that switched its approach in 2018.

Our analytical contributions begin with the introduction of a new search-and-matching model (Section 4), which captures critical features. First, as both approaches to adoption matching are inherently dynamic, we assume that children and families (hereafter referred to as *agents*) may enter and depart the system at any time. However, to keep the model tractable, we assume the distribution of agent types in the system remains stable over time. Second, we allow for uncertainty regarding whether a child-family pair is compatible. Third, our model captures the heterogeneous preferences of agents, which is a key distinction between our paper and most earlier literature. While we focus specifically on analyzing adoption systems, our work is also relevant to more general search-and-matching theory.

We perform a game-theoretic comparison of the two approaches within our model. In Section 5.2, we establish that (pure strategy) equilibria are guaranteed to exist under both *search technologies* (i.e., family-driven search and caseworker-driven search). Furthermore, we find that equilibria form a lattice reminiscent of the structure of the set of stable matchings in standard two-sided matching markets. We then present our main theoretical result: Family-driven search equilibrium outcomes can never be Pareto improvements over caseworker-driven search equilibrium outcomes, but there are instances where each caseworker-driven search equilibrium outcome is a Pareto improvement over all family-driven search equilibrium outcomes (Section 6.1). This holds because caseworker-driven search can reduce wasted search efforts. Thus, agents can worry less about accumulating search costs when they express interest in a child. However, because of multiplicity of equilibria and

the lattice structure over these equilibria, all children can be strictly better off in either system. The same holds for families. Even when both family-driven search and caseworker-driven search only admit a unique equilibrium each, an agent can be better off in either setting (Section 6.2). This may be surprising, given that caseworker-driven search reduces wasteful search efforts on both sides of the market. We therefore explore the conditions under which caseworker-driven search usually leads to more agents being better off. In Section 7.1, we show that all children will be better off in caseworker-driven search if families are sufficiently impatient. Furthermore, increasing family supply in family-driven search can have a negative effect on children’s utilities, but not in caseworker-driven search (Section 7.2). We find numerically that caseworker-driven search leads to more desirable outcomes for a wide variety of model parameter choices (Section 8).

To supplement these analytical and numerical insights, we empirically analyze case-level data from a technology platform that a multi-county child welfare agency in Florida has been using for the majority of its search efforts since 2018 (Section 9). We compare the outcomes for hundreds of children in need of adoptive placements to a benchmark from statewide case-level data that accounts for a child’s demographic information and disabilities. To measure how child- and case-level factors influence the timing of adoptions in Florida, we fit a Cox proportional hazards model (Cox 1972) to statewide records from the federal Adoption and Foster Care Analysis and Reporting System (AFCARS). We then use this model to predict the adoption outcomes of children listed on the platform if traditional search disciplines were used instead of the agency’s caseworker-driven search approach. This number can be compared to the actual outcome data from the platform. We show that the probability of adoption within three years was 24% higher than the benchmark for the agency’s children. Then, we extend our statistical analysis to directly measure the platform effect using a time-varying Cox model in a conservatively combined dataset. We find a statistically significant 27% improvement in adoption hazard rates under caseworker-driven search compared to statewide case data.

2. Related Literature

Recently, market designers and operations researchers have shown increased interest in how to best serve historically disadvantaged communities. Researchers have, for example, studied refugee resettlement (Andersson et al. 2018, Bansak et al. 2018, Delacrétaz et al. 2020), the improvement of teacher quality at disadvantaged schools (Combe et al. 2022, 2025), the management of volunteer workforces for non-profit organizations (Berenguer et al. 2023), and allocation of public housing (Arnosti and Shi 2020, Kawasaki et al. 2021, You et al. 2022).

Our paper contributes to the limited literature that studies child welfare systems from operations or market design perspectives — a challenge first articulated by Spindler (1970) and renewed by

Slaugh et al. (2025). Slaugh et al. (2016) investigated how the Pennsylvania Statewide Adoption & Permanency Network could utilize a match recommendation tool to improve their process of matching children to prospective parents. Their spreadsheet-based tool can be seen as a simplified version of the previously mentioned data-driven software system. Robinson-Cortés (2019) worked with a foster care data set to analyze placements of children in foster homes. His model predicts that allowing placements across administrative regions would be beneficial for children. MacDonald (2019) studied a dynamic matching problem where children and families can either form reversible matches (foster placements) or irreversible matches (adoptions). In her model, children are heterogeneous in the sense that there are children with disabilities and children without disabilities, while families are homogeneous. To the best of our knowledge, we are the first to formally analyze the economic effects of search and matching in the child welfare domain, while taking into account the full heterogeneity of preferences. Baccara et al. (2014) estimated families' preferences over children available for adoption from a data set documenting the operations of adoption agencies.

Within the child welfare literature, relatively little research has investigated the effectiveness of search disciplines for children in need of adoptive placements. Some research has reported positive impacts from intensive multi-faceted search efforts by caseworkers: Vandivere et al. (2015) show that children served by skilled recruiters from the Wendy's Wonderful Kids organization were 1.7 times as likely to have an adoptive placement than a control group in an experiment with over 1,000 children. In a similar context focusing on hard-to-place youth in New York, Feldman et al. (2016) show that a program of enhanced casework improved outcomes for children. The program utilized a variety of channels to promote 88 children and conduct searches. The search methods in both experiments require extensive work from skilled caseworkers funded by grants, while the platform we study provides an example of technology assisting caseworkers. Avery et al. (2009) study national photolisting service AdoptUSKids and use a hazards model to show better outcomes for children based on activity on AdoptUSKids. However, photolistings have drawn increased scrutiny since the early 2000s; Roby and White (2010) describe risks for exploitation and bullying for children publicly listed online.

Even though child adoption matching markets bear similarities to other two-sided matching markets such as centralized labor markets (Roth 1984, 1991) or ride-sharing platforms (Ma et al. 2020), there are important differences that necessitate new models and analyses. Adoption matching is inherently dynamic, and there is no centralized clearinghouse that determines final matches. Matches are only ever proposed, and both sides of the market have to invest *search efforts* to identify a match candidate.

One matching market similar to adoption from foster care is online dating (Hitsch et al. 2010a,b, Lee and Niederle 2015, Halaburda et al. 2018, Kanoria and Saban 2021). However, most online

dating markets are *completely* decentralized despite their dynamic and recommendation-based features. Individuals in search of a romantic partner can, at any time, decide to browse a dating platform and reach out to other individuals who appeal to them. In contrast, the approaches we analyze in this paper follow a *centralized protocol* despite the dynamic decentralized search component: Caseworkers perform repeat searches for a family on behalf of children, and they do so in approximately regular time intervals. Caseworkers, therefore, play an essential role throughout the process since they act as an intermediary to protect vulnerable children. Crucially, this introduces an asymmetry, both in the number of agents simultaneously active and the market power of agents on the two sides, that we have not observed in any other previously studied matching market. As a consequence, we develop a new model that allows us to capture the features of the two different search technologies in one model.

Purely random decentralized matching models have been widely studied under search frictions and homogeneous preferences, with transferable utility (Shimer and Smith 2000, 2001, Atakan 2006), and more relevant to our paper, with non-transferable utility (Eeckhout 1999, Chade 2001, Smith 2006). In all of these studies, unlike in our paper, preferences are aligned on each side of the market following a strict order of quality common for each individual on the same side. More recent studies have combined directional search rather than random search as an important feature (Lauermann et al. 2020, Cheremukhin et al. 2020).

Besides the directed versus random search distinction, simultaneous (Stigler 1961) versus sequential (Weitzman 1979) search have also been studied by classic search theory, with Chade and Smith (2006) bridging the gap by characterizing optimal hybrid strategies. Other work extends this analysis to competitive environments, such as labor markets. Albrecht et al. (2006) and Kircher (2009) study workers applying to multiple firms, yielding approximately efficient outcomes through endogenous wage dispersion that coordinates search. More recent work on online platforms also focuses on the trade-off between the two approaches. Honka and Chintagunta (2017) documents that simultaneous consumer search intensifies price competition on insurance aggregators, while Auster et al. (2025) cautions that frictionless simultaneous contacting can exacerbate adverse selection. Chade et al. (2017) and Wright et al. (2021) provide comprehensive surveys of various search models.

The closest to our model are the search theory papers by Adachi (2003) and Lauermann and Nöldeke (2014) on marriage markets, which consider non-symmetric preferences on both sides of the market. However, they only consider a single randomly chosen potential match in each period, while our problem requires that multiple matches with uncertain suitability may be investigated using either caseworker-driven or family-driven search in each period. Similarly, while Immorlica et al. (2023) considers the problem of a centralized platform guiding search through a sequence of

match proposals to agents in a two-sided market, they also only consider a single potential match per period. Additionally, they assume that the value for a match is symmetric between both sides. This is unrealistic for an adoption setting, where families’ and children’s desires can be at odds with each other. Consequently, we allow for non-symmetric values.

Our work is also related to the literature on dynamic matching regarding the effects of congestion (Arnosti et al. 2021, Leshno 2022) and different practical policies (Ünver 2010, Akbarpour et al. 2020b, Sönmez et al. 2020, Akbarpour et al. 2020a, Kerimov et al. 2023) in various market-design environments, including settings intended to help vulnerable populations (Baccara et al. 2020, Kasy and Teytelboym 2020). Different studies investigate how matching platforms should be designed so that desirable outcomes can be achieved (see, e.g., Lee and Niederle 2015, Fradkin 2017, Akbarpour et al. 2020a, Altinok and MacDonald 2023, Dierks et al. 2024). The research in this area most closely related to our work is Shi (2023), in which the author explores which side of the market should drive the search process depending on which side’s preferences can easily be expressed or satisfied. The setting, however, is quite different from our work since we allow agents on *both* sides of the market to arrive over time and let them face an optimal stopping problem as they can decide to remain unmatched until better future match opportunities arise.

3. Descriptions of Approaches to Adoption Matching

Based on conversations with caseworkers and managers of a Florida agency that transitioned from a family-driven search to a caseworker-driven search approach, we provide a more detailed introduction to how both approaches function in practice. Although very little has been published about search approaches and their prevalence, subject-matter experts indicate that similar approaches are commonly used nationally, with some variation among states and agencies. In any system, a prospective adoptive family first undergoes an extensive vetting and training process, which usually entails a written *home study*. Through this report, a caseworker evaluates the family on various dimensions, gathering information from home visits, interviews with family members, third-party sources, and the caseworker’s own judgment. State regulations determine minimum requirements for home studies to assess the suitability of the intended adoptive parents for different types of children. Typically, home studies also include additional information about the family’s environment and preferences. More details on home studies can be found in online Appendix A.

Regardless of the search approach, agencies may then employ scoring rules to assess family suitability for individual children, sometimes generating and using these scores as part of a recommender system (Slaugh et al. 2016). These tools and others, such as those described by Hanna and McRoy (2011), provide suggestions but are not determinative; all investigation and matching decisions are in the purview of a committee comprising the caseworker, supervisors, and other agency staff.

Once a family has gone through this process, they can apply for adoptive placements for any child for whom the parental rights of the birth family have been terminated. If an interested family is denied a requested placement for non-formal reasons — including the intent to place the child with a different family — Florida Administrative Code 65C-16.005(9) requires that the case be additionally reviewed by a five-person Adoption Applicant Review Committee to ensure a fair evaluation.

The core challenge is making the most suitable families aware of children who need adoptive placements. This is where the agency’s search discipline becomes relevant: identifying and informing suitable families. It should be noted that although there is typically a single agency responsible for finding placements for a child, their search is non-exclusive: any qualified prospective family can, in theory, apply for an adoptive placement and be considered by the agency in the same way as families identified through the agency’s own search. Consequently, a small percentage of children get adopted through other channels, e.g., by friends of their foster (or birth) family.

3.1. Family-driven Search

Traditionally, most agency lies with the families. In a typical family-driven search protocol, the agency emails all approved families when a child becomes eligible for adoption after a judge’s termination of parental rights (TPR) order. The email provides brief details about the child and invites families to express interest. After families have a chance to respond, the caseworker responsible for the child compiles a list of interested families. In most jurisdictions, caseworkers must give serious consideration to every responding family. The caseworker then begins reviewing those families to determine which family is the best suitable match. Obvious mismatches (e.g., a wheelchair-dependent child and a family living in a house not designed to accommodate wheelchairs) can be quickly identified and screened out, but most candidates require a careful review of the home study and follow-up interviews. This process is time-consuming and emotionally taxing for families. If multiple families are suitable, or the decision is complex, the case is referred to a committee, such as the Adoption Applicant Review Committee in Florida, for a final determination. When no family is deemed an adequate fit, the child remains in foster care, and the child’s availability is re-advertised after some time.

3.2. Caseworker-driven Search

In caseworker-driven search, agencies do not issue broad announcements. Instead, once a given child becomes eligible, the child’s caseworker sequentially informs specific families from the approved pool and invites them to apply to adopt the child. Family selection strategies vary across caseworkers. Some rely more on scoring rules and recommender systems, while others may rely on informal networks to prioritize families they know or apply their own heuristics to identify families to

consider. If an invited family expresses interest, the investigation mirrors the family-driven search process, including potential committee review. If the caseworker (or the review committee) feels that a family is reasonable but not ideal, they may also invite additional families to apply before coming to a final decision. If the caseworker — who might be managing dozens of cases at different stages of the child welfare system — has exhausted the pool of eligible families without a suitable match emerging, the search is suspended until the family pool has sufficiently refreshed. This means that the child remains in foster care until the search is restarted after some time.

4. Preliminaries

In this section, we develop an analytical model to contrast the two search disciplines and derive characterizations of agents' utilities.

4.1. Model

In our model, agents (children and families) have observable characteristics. Agents with the same characteristics are said to be of the same *type*. We treat sibling sets of children who should be placed together as a single child. Furthermore, we view the caseworker as a direct representative of the child; thus, we consider a child-caseworker pair as a single child agent. We let $C = \{c_1, \dots, c_n\}$ and $F = \{f_1, \dots, f_m\}$ denote the set of all n child types and the set of all m family types, respectively. Individual agents are indifferent between agents of the same type, i.e., their preferences are over agent types: A child of type c has a value $v_c(f) \in \mathbb{R}$ for family type f , and a family of type f has a value $v_f(c) \in \mathbb{R}$ for child type c . Preferences are assumed to be strict, i.e., $v_c(f) \neq v_c(f')$ if $f \neq f'$ and $v_f(c) \neq v_f(c')$ if $c \neq c'$. Agents' valuations are summarized by a list of vectors $v = (v_{c_1}, \dots, v_{c_n}, v_{f_1}, \dots, v_{f_m})$, and each agent has a value of 0 for remaining unmatched. Given valuations v , we let \bar{v} denote the maximum value of all $v_c(f)$ and $v_f(c)$.

There are infinitely many discrete time steps. At any time step, there is at most one agent of each type present in the system. Thus, we can use $c \in C$ and $f \in F$ to refer to either individual agents and agent types without ambiguity. We refer to an agent (or agent type) from either set as $i \in A := C \cup F$. From now on, we will simply say that agent $i \in A$ is *active* if they are present at the current time step. To account for multiple agents from the same type being present, we can introduce multiple agent types that are arbitrarily close in value, effectively creating a tie-breaker between agents of “almost” the same type.

At the beginning of each time step, all active agents are determined as follows: For each family type, one family of that type is active with probability $\lambda \in (0, 1]$. This is determined independently for each family type. We call parameter λ the *market thickness indicator*, as it determines the expected number of active families at each time step: For small values of λ , there will be few active families in expectation. For large values of λ , it is quite likely that a family of each type will be

active. Further, exactly one child (type) is selected uniformly at random to be active. This feature is motivated by search processes used by adoption agencies: A caseworker works on the case of one child at a time, which we assume she selects randomly.²

There is uncertainty regarding whether a specific child c and a specific family f are compatible: With probability $p \in (0, 1)$, f is a *suitable* match for c and *unsuitable* with probability $1 - p$. Whether a match is suitable or not is determined independently at random for child-family pairs. We refer to p as the *match success probability*. Parameter p captures the following aspect of adoption markets: When a family shows interest in a child, the family's decision is based on limited reported information, such as the sex, ethnicity, age of the child, and known disabilities. However, there are many other important characteristics of a child that determine whether the child is actually a good match for the family and whether there is mutual attraction. The same holds for a child (or his caseworker) showing interest in a family. Only if a family f is a suitable match for child c can a match between c and f be formed. If c and f form a match, they both obtain a value of $v_c(f)$ and $v_f(c)$, respectively.³ Determining the suitability of a match is costly for both sides, as it is a time-consuming process. Children and families incur search costs $\kappa_C \in \mathbb{R}_+$ and $\kappa_F \in \mathbb{R}_+$, respectively, each time the suitability of a match including them is determined. All agents discount the future; however, they only discount time steps in which they are active. Children's and families' discount factors are $\delta_C \in [0, 1)$ and $\delta_F \in [0, 1)$, respectively.⁴ Although we keep homogeneous costs and discount factors to keep the analysis and notation concise, they can be made type-specific without loss of generality, and all our results go through.

We model adoption matching as a dynamic process that we assume to be stationary. An instance $(v, \delta_C, \delta_F, \kappa_C, \kappa_F, p, \lambda)$ together with a search technology (which we will introduce) induce a game. To reduce notation, we assume that instance $(v, \delta_C, \delta_F, \kappa_C, \kappa_F, p, \lambda)$ is fixed unless stated otherwise. Agents' strategies in this game are captured as follows: Child c is either *interested* in a family of type f or not. Similarly, each family f is either *interested* in a child of type c or not. We assume that all agents of the same type play the same strategy, and agents don't change their strategies in different time steps. Therefore, we can represent a strategy for a child c as a vector $s_c \in \{0, 1\}^m$, where $s_c(f) = 1$ if c is interested in matching with a family of type f . Similarly, a strategy for a family f is given by a vector $s_f \in \{0, 1\}^n$, where $s_f(c)$ indicates whether f is interested in children of type c .

² Our model does not endogenize the number of agents present in the system. This kind of instant replacement is a standard large-market assumption in the search-and-matching literature, and it is necessary to keep our model tractable.

³ Note that a uniform, binary suitability probability p is an abstraction made to keep the model tractable. See online Appendix B.

⁴ An alternative interpretation is to think of children and families leaving the process before the next time step when they are active with probability $1 - \delta_C$ and $1 - \delta_F$, respectively.

A strategy profile is a tuple of vectors $s = (s_{c_1}, \dots, s_{c_n}, s_{f_1}, \dots, s_{f_m})$, while we let S denote the finite set of all possible strategy profiles. For $i \in A$ we let s_{-i} denote the tuple of all agents' strategies in s , except that agent i is excluded. We say that c and f are *mutually interested* in each other under strategy profile s if $s_c(f) = 1$ and $s_f(c) = 1$. The set $M(s) = \{(c, f) \in C \times F \mid s_c(f) = s_f(c) = 1\}$ is called the *matching correspondence of s* . We use $M_i(s)$ to denote the set of agents that agent i is mutually interested in under s .

In the remainder of this section, we describe the two search technologies. We analyze the dynamic stochastic games induced by an instance and a search technology in a full information environment.

Family-driven Search (FS): At the beginning of each time step, after a child c is randomly chosen to be active, each family that is active and interested in c lets the child's caseworker know of their interest. This corresponds to families responding to an email announcement made by a caseworker. The caseworker immediately discards any families without mutual interest in c , i.e., where $s_c(f) = 0$ or $s_f(c) = 0$. The caseworker then investigates all remaining families to determine whether they would actually be a suitable match. Recall that each investigated family is a suitable match for c with probability p —which is determined independently for each family—and that each agent incurs cost for each investigation in which they are involved. After all families have been processed by the caseworker, c either matches with the most preferred choice from among those families identified as suitable matches or remains unmatched if no such family exists. We move on to the next time step.

Caseworker-driven Search (CS): After a child c is randomly chosen to be active at the beginning of a time step, the child's caseworker can sequentially inform any of the active families. We assume families are informed in decreasing order of $v_c(f)$.⁵ If there is mutual interest between c and f , the suitability of a match between c and f is investigated. If the match turns out to be suitable, c 's search is over, c and f are matched and leave.⁶ We then move on to the next time step. Otherwise, the caseworker continues the search by selecting the next family in the list. If all families have been processed, the child remains unmatched and we move on to the next time step. Proposition 11 in online Appendix C.1 shows that c 's utility is maximized if the caseworker processes families in decreasing order of $v_c(f)$.

⁵ While this is trivially optimal in our model, higher degrees of uncertainty may warrant contacting families in a different order. See online Appendix B for a discussion.

⁶ In practice, both sides still need to agree to the match. However, given that it would not be rational to investigate a match you are not willing to accept and that the remaining families all have lower values, refusing a suitable match at this point is never rational.

4.2. Utilities

We define agents' utilities at a time step and characterize their flow utilities in both FS- and CS-induced stochastic games. Assume child c is active at the current time step and that active families are yet to be determined. Let f be an arbitrary family. For any $s \in S$, let $b_{cf}(s) = |\{f' \in M_c(s) \mid v_c(f') > v_c(f)\}|$ denote the number of families in $M_c(s)$ that c likes better than f . Further, let $\beta_{cf}(s)$ denote the probability that c will *not* match with any other family that c prefers over f at the current time step. Noting that for any child c the probability that a mutually interested family f' is active at the current time step and a suitable match is λp , it follows immediately that $\beta_{cf}(s) = (1 - \lambda p)^{b_{cf}(s)}$ for both FS and CS.

In FS, since investigations are conducted simultaneously, search costs are incurred for all of them, and therefore, the *expected immediate utility* for the active child c at an arbitrary time step follows as

$$\bar{u}_c^{FS}(s) = \lambda \sum_{f' \in M_c(s)} \left(\beta_{cf'}(s) p v_c(f') - \kappa_C \right). \quad (1)$$

In CS, as investigations are conducted sequentially, search costs κ_C are only incurred if the child has not successfully matched with a higher-valued family. Consequently, the *expected immediate utility* of the active child c follows as

$$\bar{u}_c^{CS}(s) = \lambda \sum_{f' \in M_c(s)} \beta_{cf'}(s) (p v_c(f') - \kappa_C). \quad (2)$$

In both cases above, λ is the probability that a family is present to be investigated when the child is active.

Similarly, for a family f , we can express the expected immediate utilities for FS and CS at an arbitrary time step (conditional on f being active) by

$$\bar{u}_f^{FS}(s) = \frac{1}{n} \sum_{c' \in M_f(s)} \left(\beta_{c'f}(s) p v_{c'}(f) - \kappa_F \right) \quad (3)$$

$$\bar{u}_f^{CS}(s) = \frac{1}{n} \sum_{c' \in M_f(s)} \beta_{c'f}(s) (p v_{c'}(f) - \kappa_F), \quad (4)$$

where $1/n$ is the probability that a given child is active in a timestep.

From this, the crucial difference between FS and CS in our model becomes apparent: In FS, search costs are always incurred if there is mutual interest between agents. In CS, however, the sequential nature of the search means that search costs are only incurred if there is mutual interest and all previous match attempts at the time step have been unsuccessful.

We assume that each agent is risk-neutral and maximizes their expected (overall) utility, which is the expected discounted value of their eventual match minus the total discounted search costs

they incur.⁷ We use $(z)^+$ as shorthand notation for $\max\{z, 0\}$. $\mathbb{1}[\cdot]$ is the indicator function, which has value 1 if its argument is true and value 0 otherwise. We denote the expected (overall) utility of agent i under strategy profile s by $u_i^{FS}(s)$ in FS and $u_i^{CS}(s)$ in CS. Whenever it is clear from context whether we are referring to FS or CS we will simply write $u_i(s)$. Proposition 1 characterizes children's and families' utilities in FS and CS via balance equations.

PROPOSITION 1. *Given strategy profile s , child c 's utility in FS is the unique value $u_c^{FS}(s)$ that satisfies*

$$u_c^{FS}(s) = \delta_C u_c^{FS}(s) + \lambda \sum_{f \in M_c(s)} \left(\beta_{cf}(s) p(v_c(f) - \delta_C u_c^{FS}(s)) - \kappa_C \right). \quad (5)$$

Similarly, family f 's utility in FS is the unique value $u_f^{FS}(s)$ that satisfies

$$u_f^{FS}(s) = \delta_F u_f^{FS}(s) + \frac{1}{n} \sum_{c \in M_f(s)} \left(\beta_{cf}(s) p(v_f(c) - \delta_F u_f^{FS}(s)) - \kappa_F \right). \quad (6)$$

In CS, child c 's utility in CS is the unique value $u_c^{CS}(s)$ that satisfies

$$u_c^{CS}(s) = \delta_C u_c^{CS}(s) + \lambda \sum_{f \in M_c(s)} \beta_{cf}(s) \left(p(v_c(f) - \delta_C u_c^{CS}(s)) - \kappa_C \right). \quad (7)$$

Similarly, family f 's utility in CS is the unique value $u_f^{CS}(s)$ that satisfies

$$u_f^{CS}(s) = \delta_F u_f^{CS}(s) + \frac{1}{n} \sum_{c \in M_f(s)} \beta_{cf}(s) \left(p(v_f(c) - \delta_F u_f^{CS}(s)) - \kappa_F \right). \quad (8)$$

A formal proof can be found in online Appendix D.1. One difference between FS and CS is immediately apparent from the above balance equations. In CS, search costs only incur if previous match attempts have been unsuccessful. In FS, however, costs are always incurred if a family is present and mutually interested.

5. Equilibria

We use a tie-breaking assumption, which allows us to exclude degenerate equilibria later on. After stating this assumption, we introduce two classes of strategies — one for CS and one for FS — and show that these classes capture agents' best responses. This will be helpful for obtaining results later on. We show that for both search technologies equilibria always exist and that equilibria form a lattice.

⁷ In practice, while caseworkers have multiple cases assigned to them, their goal with each case is to maximize the utility of the child belonging to that case. That is, based on our interviews with domain experts, the caseworker considers all assigned cases separately, and each child-caseworker pair is considered a separate self-interested agent.

5.1. Threshold Strategies

For both search technologies, we make the following tie-breaking assumption: If agent i 's utility would (weakly) increase from mutual interest with agent j , then i will be interested in j —even if j is not interested in i . Similarly, if agent i 's utility would decrease from mutual interest with agent j , then i will not be interested in j . This assumption allows us to exclude degenerate equilibria (e.g., no agent being interested in any other agent) later on without restricting agents in their endeavor to maximize their utility.

We now introduce *threshold strategies* for FS and CS. As we will see, our tie-breaking assumption implies that agents' best responses belong to the class of threshold strategies. Note that a best response always exists, because S is finite.

DEFINITION 1. Child c plays a *CS threshold strategy* (CS-TS) with threshold $y_c \in \mathbb{R}$ in s , if

$$s_c(f) = \mathbb{1}[p(v_c(f) - \delta_C y_c) \geq \kappa_C] \text{ for all } f \in F. \quad (9)$$

Family f plays a *CS-TS* with threshold y_f in s , if

$$s_f(c) = \mathbb{1}[p(v_f(c) - \delta_F y_f) \geq \kappa_F] \text{ for all } c \in C. \quad (10)$$

Child c plays an *FS threshold strategy* (FS-TS) with threshold y_c in s , if

$$s_c(f) = \mathbb{1}[\beta_{cf}(s)p(v_c(f) - \delta_C y_c) \geq \kappa_C] \text{ for all } f \in F. \quad (11)$$

Family f plays an *FS-TS* with threshold y_f in s , if

$$s_f(c) = \mathbb{1}[\beta_{cf}(s)p(v_f(c) - \delta_F y_f) \geq \kappa_F] \text{ for all } c \in C. \quad (12)$$

In a CS-TS or an FS-TS, the threshold y_i can be interpreted as i 's reservation utility. However, note that these threshold strategies are more involved than standard *simple threshold strategies* (Adachi 2003, Immorlica et al. 2023). While in a simple threshold strategy, an agent would be interested if the expected value is above their reserve utility, i.e., $pv_i(j) \geq y_i$, this does not suffice in this case. Instead, agents also have to account for their costs and for the likelihood that the child may find another match with a higher value during the same period.⁸

Let $u_i^{FS*}(s_{-i})$ and $u_i^{CS*}(s_{-i})$ denote i 's utility from a best response to s_{-i} in FS and CS, respectively. As before, we simply write $u_i^*(s_{-i})$ if there is no ambiguity. Proposition 2 shows that agents' best responses always have the form of a threshold strategy.

PROPOSITION 2. Let $i \in A$ and s_{-i} be an arbitrary strategy profile of all agents excluding i . In both FS and CS, a best response of i to s_{-i} corresponds to a threshold strategy with threshold $u_i^*(s_{-i})$.

⁸ See Proposition 12 in online Appendix C.2 for an illustration of the non-existence of simple threshold best responses.

A formal proof can be found in online Appendix D.2. To derive results later on, it will prove useful to switch between thresholds and strategies. We therefore provide the following definition.

DEFINITION 2. In CS, a *strategy profile induced by threshold profile* $y \in \mathbb{R}^{n+m}$ is denoted by $s^{CS}(y)$ and satisfies for each $i \in A$, $s_i^{CS}(y)$ is a CS-TS with threshold y_i in $s^{CS}(y)$. In FS, a *strategy profile induced by threshold profile* $y \in \mathbb{R}^{n+m}$ is denoted by $s^{FS}(y)$ and satisfies for each $i \in A$, $s_i^{FS}(y)$ is a FS-TS with threshold y_i in $s^{FS}(y)$.

We again omit the superscript if this does not lead to ambiguity. It is trivial to obtain $s^{CS}(y)$ by inserting y in the corresponding equations in Definition 1. Algorithm 1 from online Appendix E can be used to compute $s^{FS}(y)$. From now on, we use $\beta_{cf}^{FS}(y)$ and $\beta_{cf}^{CS}(y)$ as shorthand-notation for $\beta_{cf}(s^{FS}(y))$ and $\beta_{cf}(s^{CS}(y))$, respectively. Whenever it is clear from the context, we simply write $\beta_{cf}(y)$.

5.2. Equilibrium Existence and Lattice Structure

In this section, we show that Nash equilibria always exist under both search technologies.⁹ We say that strategy profile s is an *equilibrium in FS (FSE)* if s is a Nash equilibrium in the game induced by FS. Analogously, strategy profile s is an *equilibrium in CS (CSE)* if s is a Nash equilibrium in the game induced by CS. We use S^{FS} to denote the set of FSE, and let $Y^{FS} = \{(u_i(s))_{i \in A}\}_{s \in S^{FS}}$ be the corresponding set of *equilibrium threshold profiles in FS*. For CS, those sets are defined analogously. Before we can prove that these sets are never empty, we need to define a partial order \leq_C on $Y = [0, \bar{v}]^{n+m}$. Note that if agents only play individually rational strategies, their utility is always lower bounded by 0 and upper bounded by \bar{v} .

DEFINITION 3. Let \leq_C be the partial order on Y , where for all $y, y' \in Y$ it holds that $y \leq_C y'$ if and only if $y_c \leq y'_c$ for all $c \in C$ and $y_f \geq y'_f$ for all $f \in F$.

Having defined partial order \leq_C , we can now prove that equilibria always exist in both settings. We state this result in Proposition 3.

PROPOSITION 3. *The set of FS and CS equilibrium threshold profiles Y^{FS} and Y^{CS} is non-empty and both (Y^{FS}, \leq_C) and (Y^{CS}, \leq_C) form complete lattices.*

A formal proof can be found in online Appendix D.3. The proof proceeds by defining a best-response mapping and showing fixed-point existence using Tarski's fixed-point theorem (Tarski 1955). Such fixed points coincide with pure-strategy equilibria.

In general, there can be more than one FSE or CSE for a fixed instance $(v, \delta_C, \delta_F, \kappa_C, \kappa_F, p, \lambda)$. Proposition 3 not only guarantees that equilibria always exist in both settings, but also highlights

⁹ Technically, we have a stochastic game model, and therefore, these are Nash equilibria of stochastic games. They correspond to the Markov-perfect Nash-equilibrium selection among subgame-perfect Nash equilibria when considered as a repeated game.

that there is a special ordering over equilibria: there exists a child-optimal equilibrium that children unanimously prefer over all other equilibria; i.e., their utility is weakly higher compared to any other equilibrium. Similarly, there exists a family-optimal equilibrium that families prefer. From now on, we let s^{co-CS} denote the *child-optimal CSE (co-CSE)*, s^{co-FS} the *child-optimal FSE (co-FSE)*, s^{fo-CS} the *family-optimal CSE (fo-CSE)*, and s^{fo-FS} the *family-optimal FSE (fo-FSE)*. This is reminiscent of the structure of the set of stable matchings in standard two-sided matching markets (Knuth 1997).

6. Comparison of Family-driven Search and Caseworker-driven Search

In this section, we investigate the impact of the two search technologies on equilibrium outcomes. Here, we present our main theoretical result: An FSE can never Pareto dominate a CSE, as any increase in utility for one agent can only arise if another agent lowers their interest threshold, corresponding to a decrease in that agent's utility. There exist instances, however, where each CSE is a Pareto improvement over all FSEs. We further find that no approach is always preferable for either children or families.

6.1. Pareto Comparison

A natural way to determine which equilibrium outcomes are preferable is to check whether one equilibrium is a Pareto improvement over the other. We first formalize the Pareto dominance relationship for strategy profiles in our model.

DEFINITION 4. Strategy profile $s \in S$ is a Pareto improvement over strategy profile $s' \in S$ if $u_i(s') \leq u_i(s)$ for all $i \in A$ and there exists $j \in A$, such that $u_j(s') < u_j(s)$.

Note that $u_i(s)$ either denotes $u_i^{CS}(s)$ or $u_i^{FS}(s)$, depending on whether we refer to s as a CSE or an FSE. We find that FSEs can never Pareto dominate CSEs, but there are instances where each CSE Pareto dominates all FSEs. Before we can formally show this, we need to state two lemmas. The first lemma is useful for understanding why FSEs cannot Pareto dominate CSEs. Lemma 1 shows that if there is a pair with mutual interest in FS that is not present in CS, then at least one of the two agents in the pair must be strictly worse off in FS compared to CS.

LEMMA 1. Let $s^{FS} \in S^{FS}$ and $s^{CS} \in S^{CS}$. If there exists $c \in C$ and $f \in F$, such that $(c, f) \in M(s^{FS})$ and $(c, f) \notin M(s^{CS})$, then either $u_c(s^{FS}) < u_c(s^{CS})$ or $u_f(s^{FS}) < u_f(s^{CS})$.

A formal proof can be found in online Appendix D.4. Intuitively, if two agents are not mutually interested in each other in CS but are in FS, then at least one of them had to lower their interest threshold in FS—which means that their optimal reservation utility is lower in FS. However, since the optimal reservation utility corresponds to the once-discounted utility, this agent must be strictly worse off.

The next lemma almost immediately follows from Lemma 1 and is used for the proof of Theorem 1 as well as for later results.

LEMMA 2. *Let $s^{FS} \in S^{FS}$, $s^{CS} \in S^{CS}$, and $c \in C$. If $M_c(s^{FS}) \subseteq M_c(s^{CS})$, then $u_c(s^{FS}) \leq u_c(s^{CS})$.*

A formal proof can be found in online Appendix D.5.

It is quite intuitive that a child c cannot be worse off under CS if all families that are mutually interested in c under FS are also interested in c under CS. This allows us to show the following theorem as our main result.

THEOREM 1. *An FSE can never be a Pareto improvement over a CSE. On the other hand, there exists an instance where all CSEs are Pareto improvements over all FSEs.*

A formal proof can be found in online Appendix D.6. The main intuition for why an FSE can never be a Pareto improvement over a CSE is that the only way an agent can be better off in FS compared to CS is to have a higher chance of matching with someone they like. But with Lemma 1, this implies that some other agent had to lower their interest threshold, which means that their utility decreased. There are two reasons why a CSE can Pareto dominate an FSE. First, CS can save agents' search costs. Second, because search costs are only incurred in CS if previous match attempts at the current time step have been unsuccessful, agents do not have to worry about accumulating search costs as much in CS as in FS. Therefore, agents are incentivized to express interest in more potential match candidates in CS compared to FS.

6.2. No Approach Dominates the Other

Even though CSEs can be Pareto improvements over FSEs, we find that CS is not always better for everyone compared to FS. In fact, a CSE might yield arbitrarily higher (or lower) utility for all children or families compared to an FSE.

PROPOSITION 4. *For any $L > 0$ and $0 < \epsilon < L$, there exists an instance where*

1. *the child-optimal equilibrium, which we denote as s^{co} , is the same in both CS and FS, and similarly, the family-optimal equilibrium, which we denote as s^{fo} , is the same in both CS and FS,*
2. *$u_c(s^{co}) = L$ for all $c \in C$ and $u_f(s^{co}) \leq \epsilon$ for all $f \in F$, and*
3. *$u_c(s^{fo}) \leq \epsilon$ for all $c \in C$ and $u_f(s^{fo}) = L$ for all $f \in F$.*

A formal proof can be found in online Appendix D.7. It proceeds by constructing examples where child and family utilities are mismatched, causing the utility gap between child-optimal and family-optimal equilibria to be arbitrarily large in both FS and CS. Thus, depending on which equilibria are realized, both approaches can be arbitrarily better for either side of the market.

Both a single child and a single family can be arbitrarily worse off in equilibrium under CS, even if FS and CS admit only one equilibrium each. This is highlighted by the following two propositions.

PROPOSITION 5. *There exists an instance where a child is strictly worse off under the unique CSE compared to the unique FSE.*

A formal proof can be found in online Appendix D.8. It proceeds by constructing an example with two child and family types where child c_1 is significantly preferred over c_2 by all families, while all children slightly prefer f_1 over f_2 . This implies that family f_2 is only matched with child c_1 if no f_1 family is currently present. The higher search costs in FS can then make f_2 lose interest in c_1 , regardless of patience levels, causing f_2 to settle for c_2 . This allows c_2 to be matched. Conversely, in CS, family f_2 does not incur high search costs for waiting until they are matched with a c_1 . If they are patient enough, f_2 therefore prefers waiting for their preferred choice c_1 . This leaves c_2 without any interested family and, therefore, unmatched.

Similarly, a family can be worse off in CS when FS and CS each only admit one equilibrium.

PROPOSITION 6. *There exists an instance where a family is strictly worse off under the unique CSE compared to the unique FSE.*

A formal proof can be found in online Appendix D.9. Just as children can benefit from families that decide to settle for a less preferred child, so can other families. It can be the case that a family f is interested in a child c in a CSE but f is not interested in c in an FSE, because the associated expected costs would be too high. Not having f as competition might be enough incentive for another family f' to be interested in c under the FSE. As a result, f' can be strictly better off in FS.

7. Effects of Model Parameters

We showed that FSEs cannot be Pareto improvements over CSEs, but CSEs can be Pareto improvements over FSEs. Additionally, we found that some agents can be better off under an FSE compared to a CSE. In order to better understand the conditions under which one of the two approaches might be preferable, we explore the effects that different parameters have on equilibrium outcomes in FS and CS. We provide two more results in favor of CS: First, we show that as families' patience decreases, at some point all children will be weakly better off in any CSE compared to any FSE. Second, increasing supply on the family side, i.e., increasing the market thickness indicator λ , can negatively affect children's utilities in FS but not in CS. Finally, as a sanity check, we investigate the effect of certain parameters or parameter combinations in the limit; all of these latter results can be found in Section F.

7.1. Discount Factors

For this subsection, let $S^{CS}(\delta'_F)$ and $S^{FS}(\delta'_F)$ denote the set of CSEs and FSEs when $\delta_F = \delta'_F$, respectively. We now demonstrate that as families' patience decreases below a certain threshold, all children will always be better off in CS than in FS.

PROPOSITION 7. *For each instance there exists $\bar{\delta}_F \in [0, 1)$, such that for all $\delta'_F \in [0, \bar{\delta}_F]$ it holds that $u_c(s^{FS}) \leq u_c(s^{CS})$ for all $c \in C$, $s^{CS} \in S^{CS}(\delta'_F)$, $s^{FS} \in S^{FS}(\delta'_F)$.*

A formal proof can be found in online Appendix D.10. Intuitively, the statement follows because in CS, families' interest in a very unlikely match incurs lower search costs than in FS. While patient families may still not be interested in some children in CS that they are interested in under FS (which drives Proposition 5), any sufficiently impatient family will be unwilling to wait.

However, as can be seen in the proof of Proposition 6, an analogous statement for families' utilities and children's patience level does not hold. Intuitively, a family f might be worse off in CS, because another family f' is not shying away from c , as f' does not have to worry about wasted search efforts in CS.

7.2. Market Thickness

Adoption agencies might intuitively prefer to have a larger pool of available families to choose from. Here, we present another result that suggests this might be generally beneficial in CS but not always in FS when it comes to children's utilities. Increasing supply on the family side, i.e., increasing the market thickness indicator λ , can negatively affect children's utilities in FS but not in CS. For the remainder of Section 7.2, assume that all instance parameters are fixed except for λ . Let $s^{co-CS, \lambda}$ denote the child-optimal CSE given market thickness indicator λ . Definitions for $s^{fo-CS, \lambda}$, $s^{co-FS, \lambda}$, and $s^{fo-FS, \lambda}$ are analogous. Proposition 8 shows that increasing λ can lead to some children being worse off in FS.

PROPOSITION 8. *There exists an instance with a child $c \in C$ and $\lambda, \lambda' \in (0, 1]$ with $\lambda < \lambda'$, such that $u_c(s^{co-FS, \lambda}) > u_c(s^{co-FS, \lambda'})$.*

A formal proof can be found in online Appendix D.11. Effectively, what is happening is that if multiple families are interested in a child, then increased market thickness λ increases competition and, therefore, the search costs for less preferred families. If the child is close to being indifferent between families and some families lose interest due to the higher cost, then the resulting decrease in the child's utility can be larger than the increase caused by a higher chance to match with a slightly more preferred family.

In CS, on the other hand, increasing λ can only have a positive effect on children's utilities in equilibrium.

PROPOSITION 9. *Let $\lambda, \lambda' \in (0, 1]$, such that $\lambda \leq \lambda'$. Then $u_c(s^{co-CS, \lambda}) \leq u_c(s^{co-CS, \lambda'})$ and $u_c(s^{fo-CS, \lambda}) \leq u_c(s^{fo-CS, \lambda'})$ for all $c \in C$.*

A formal proof can be found in online Appendix D.12. The reason why this holds in CS is that, unlike in FS, families will not shy away from children in whom they are interested just because the probability of matching with them decreases. This result is reminiscent of a similar result in standard two-sided matching markets, as the number of agents in one side increases, the other side agents become all unambiguously better off under side-optimal stable matchings (Gale and Sotomayor 1985). However, it only holds for CS and only for the children’s welfare.

8. Numerical Evaluation

We previously established that CSEs can be Pareto improvements over FSEs while FSEs cannot be Pareto improvements over CSEs, and that agents can be better off in either approach (see Theorem 1 and Proposition 4). Additionally, we have shown that all children will be better off in CS compared to FS if families are sufficiently impatient. Here, we present numerical results to further investigate the conditions under which children and families will be better off in CS or FS. Our results suggest that CS is almost always preferable for both sides of the market. Only when agents’ preferences are perfectly correlated and families are very patient do we find that, on average, there are more child types better off in FS than in CS in equilibrium.

Section 8.1 describes how our numerical experiments are set up. We then explain how equilibria are computed in Section 8.2. In Section 8.3, we compare FS and CS in terms of their Pareto dominance relationship. We further quantify how many agents are typically better off in either approach.

8.1. Setup

We now describe the setup of our numerical evaluation. We set the number of agent types on each side to be $n = m = 50$.

Valuations: For the generation of agents’ valuations, we follow other approaches from the matching literature (Abdulkadiroğlu et al. 2015, Mennle et al. 2015). Each agent type $i \in A$ is uniformly assigned a “quality” q_i at random from $[0, 1]$. Then, for each child-family pair $(c, f) \in C \times F$, idiosyncratic values $\eta_c(f)$ and $\eta_f(c)$ are randomly drawn from $[0, 1]$. For a given value $\alpha \in [0, 1]$, we obtain the preliminary valuations $v'_c(f) = \alpha q_f + (1 - \alpha)\eta_c(f)$ and $v'_f(c) = \alpha q_c + (1 - \alpha)\eta_f(c)$. Note that as α increases, agents’ preferences become more similar and end up being identical (vertical) for $\alpha = 1$. Final valuations v are obtained by normalizing v' , such that the minimal and maximal value that each agent has for a match is 0 and 1, respectively. Although we consider various values for α , note that the practical level of verticality in child welfare tends to be relatively low. For

example, data provided by the platform that we examine in Section 9 shows notable variation in families' stated preferences on several dimensions, which, taken together, suggests that preferences are not very aligned. See online Appendix G for details.

Data: We generated 200 quality-value pairs $(q^{(1)}, \eta^{(1)}), \dots, (q^{(200)}, \eta^{(200)})$ as described above. Parameters p and λ are chosen to be $p = \lambda = 0.5$, and we let $\delta := \delta_C = \delta_F$ and $\kappa := \kappa_C = \kappa_F$. For each pair $(q^{(k)}, \eta^{(k)})$, we computed the child-optimal CSE/FSE and the family-optimal CSE/FSE for each combination of α , δ , and κ , where $\alpha \in \{0, 1/3, 2/3, 1.0\}$, $\delta \in \{0.8, 0.9, 0.975, 0.99\}$, and $\kappa \in \{0.01, 0.02, 0.05, 0.1\}$. Thus, we consider $200 \cdot 4 \cdot 4 \cdot 4 = 12800$ different instances and compute a total of 51200 equilibria. Unless specified otherwise, results are averaged over all instances.

8.2. Equilibrium Computation

The mapping T defined in the proof of Proposition 3 can be used to find the child-optimal and family-optimal equilibria. The following procedure converges to an equilibrium threshold profile: Start from the \leq_C -minimal element in Y , i.e., the minimum point of the lattice spanned by the partial order \leq_C over the threshold vectors in the game, and recursively apply T to it. The \leq_C -minimal element in Y is the threshold vector \underline{y} where $\underline{y}_c = 0$ for each child c and $\underline{y}_f = \bar{v}$ for each family f . This produces a sequence y^0, y^1, y^2, \dots of threshold profiles, which converges to the FSE. Starting from the \leq_C -maximal element in Y , i.e., the maximum point of the lattice spanned by the partial order \leq_C over the threshold vectors in the game, yields the co-FSE. The \leq_C -maximal element is the threshold vector \bar{y} where $\bar{y}_c = \bar{v}$ for each child c and $\bar{y}_f = 0$ for each family f . In order to terminate after a finite number of steps, we force the procedure to stop once $|y_i^k - y_i^{k+1}| \leq \epsilon$ for all $i \in A$ for some previously chosen small parameter $\epsilon > 0$. The threshold profiles obtained by this procedure can then be mapped to the corresponding strategy profiles. We have performed additional checks to validate that the computed strategy profiles are indeed equilibria.

8.3. Results

Before comparing FS and CS, we first note that family- and child-optimal equilibria in FS coincide roughly 97% of the time. The same holds for CS. For simplicity, we only consider family-optimal equilibria in our analysis. As equilibria are almost always unique within each search technology, results for child-optimal equilibria do not differ markedly. We do, however, observe substantial differences between FS and CS. Of all cases considered, the CSE and the FSE only coincide once in the sense that the same agents are mutually interested in each other.

Consistent with Theorem 1, the family-optimal FSE never represents a Pareto improvement upon the corresponding family-optimal CSE. However, for approximately 22% of all instances, the CSE Pareto dominates the corresponding FSE. Figure 1 shows the distribution of cases in which the CSE dominates an FSE for different discount factors and levels of correlation among preferences. Two

insights emerge from this analysis: First, as agents become more impatient, CSEs more frequently constitute Pareto improvements over FSEs. As indicated by Proposition 7, once families become sufficiently impatient, any CSE will Pareto dominate all FSEs. Second, when preferences exhibit high correlation, CSEs rarely Pareto dominate FSEs. The case of vertical preferences — i.e., $\alpha = 1$ — helps to explain this effect. If agents are patient enough, a family f in the CS regime might wait for an opportunity to match with a high-type child c , even if f is not c 's first choice and f must wait a long time until getting matched. In FS, however, if there are enough other families that c prefers over f , f or c might shy away from being interested in order to avoid accumulating search costs for such an “unlikely” match. In that case, f might settle for another low-type child or multiple low-type children instead (see the example from the proof of Proposition 5). These low-type children now benefit from FS, while f will be worse off in FS compared to CS.




Figure 1 Percentage of instances in which the CSE is a Pareto improvement over the corresponding FSE for different combinations of agents' patience levels and degree of preference correlation.

The previous Pareto comparison only allows for a very high-level comparison of FS and CS. To better understand the conditions under which certain agents benefit from FS or CS, we compare the number of agents who are better off in either search discipline. Our numerical experiments show that all families are almost always better off in CS. We refer the reader to online Appendix H.1 for more details on families' statistics. For children, the combination of model parameters affects which approach appears more appealing. Figure 2 shows how many children are (strictly) better off (in terms of utilities) in CS and FS for different parameter combinations.

CS provides higher utility than FS for almost all children when agents are sufficiently impatient (e.g., $\delta = 0.8$) because CS allows agents to express interest in more potential match partners without risking wasted search efforts. Being interested in more agents increases the probability of getting matched at each time step, which is especially valuable to children when patience is low. On the other hand, FS incentivizes agents to focus on a smaller set of match candidates due to higher expected total search costs. When $\delta = 0.8$, families will, on average, be interested in 35.9 and 43.5




Figure 2 The ratio of children on average (strictly) better off in terms of utilities in either approach in the family-optimal equilibrium for different combinations of agents' patience and preference correlation level.

child types in FS and CS, respectively.¹⁰ Interestingly, more children benefit from FS than CS when agents are extremely patient and agents' preferences are almost completely aligned. Although unlikely to occur in practical child welfare settings, this explains why CSEs are less frequently Pareto improvements over FSEs under these conditions, as we previously saw in Figure 1.




Figure 3 The ratio of children on average (strictly) better off in terms of (expected discounted) match value in either approach in the family-optimal equilibrium for different combinations of agents' patience and the level of preference correlation.

Figure 3 shows that CS not only reduces wasted search efforts in many cases but also enables children to match with more preferred families. We calculate a child's match value as the child's utility, ignoring the expected search costs, which might be less relevant to a policymaker trying to improve child outcomes.

9. Empirical Evidence from a Field Implementation

To validate our model and understand the real-world implications of switching from an FS to a CS approach, we analyze children's outcomes for a multi-county region in Florida that started

¹⁰ The probability of matches occurring is another metric of interest to stakeholders. Our findings on match probabilities can be found in Section H.2.




Figure 4 Children eligible for adoption in Florida on July 1 of each year and whether they are adopted by June 30 of the following year (Florida Department of Children and Families 2019, 2024).

implementing a CS approach on July 1, 2018, by adopting a technology platform which we simply refer to as the *platform*. The region, one of 20 *circuits* in Florida, is administered on the state’s behalf by a non-profit organization — which we refer to as the *agency*. The agency’s previous FS approach relied on regular email announcements to the full pool of registered families to advertise children in need of adoption. Out of frustration with the difficulties in finding adoptive placements for children, the agency’s leadership decided to implement the technology platform and follow a CS discipline. In the new approach, caseworkers contact individual families listed on the platform whom they consider a good match for a child. The platform’s purpose in this approach is to provide an easy-to-navigate database of prospective families. Additionally, it offers family recruitment, evaluation, and recommendation services to support caseworkers.

Our analysis compares the agency’s performance to outcomes for children in Florida reported in the Adoption and Foster Care Analysis and Recruiting System (AFCARS) from the US Children’s Bureau, Administration on Children, Youth and Families (2023a,b) for FY2015 to FY2021, which is the most recently available data on all foster care cases and finalized adoptions. From 766,527 AFCARS foster care 6-month update records for Florida children during this period, we identified 10,286 children as legally free and clear for adoption with cases starting after October 1, 2014. While the platform’s usage is free to all agencies in Florida, the overwhelming majority of placements during the observed period were still searched via traditional methods, with only the considered agency utilizing CS as their primary method. While some other agencies occasionally employed the platform’s tools to help identify families for harder-to-place children, these constitute less than 4% of adoptions in the AFCARS dataset.

For the agency, we obtained case data from the platform about 279 children in need of adoptive placements who were listed by the agency before the AFCARS cut-off date of October 1, 2021. The platform provided its first matches around July 1, 2018. While the agency listed all children in

need of an adoptive placement on the platform, a small number of these children found placement through other channels, such as word of mouth within the agency, Florida’s online photolisting websites, and other contracted recruitment efforts for specific children. We include these placements in our analysis because serendipitous matches occur outside the primary (FS or CS) search channel in any system but cannot be identified in the AFCARS dataset. Details of how the data is assembled and pre-processed can be found in online Appendix I.1.

It should be noted that Florida experienced at least two dramatic shocks to its child welfare system in the years over which the platform was implemented. As shown in Figure 4, the statewide population of children legally free for adoption — which also includes children on a path to adoption by relatives and foster parents — increased by nearly 50% from 2015 to 2018 (Florida Department of Children and Families 2019, 2024). Quast et al. (2018) document a relationship between opioid prescriptions and child welfare system entries in Florida in the early 2010s that could partially explain this increase in children in need of adoption as the opioid crisis worsened. Despite an increase in children needing adoptive placements, the state’s reported metric of the percentage of children eligible for adoption on June 30 of some year and adopted by July 1 of the following year remained above 55% from 2017 until 2019. However, that metric dropped below 50% in the years after the coronavirus pandemic, as caseworker turnover, staffing shortages, and slower judicial processing times may have hindered adoptive searches. After reaching the lowest reported value of 46.7% for 2022-2023, the metric only returned to above 50% in the most recent report covering 2023-2024.

9.1. Statistical Approaches

Using the AFCARS and platform datasets, we assess the platform’s impact on children’s outcomes through two statistical approaches that take advantage of the datasets’ similar structure. First, we construct a benchmark for the set of agency children using the Florida AFCARS case data. Second, we statistically measure the treatment effect of the platform by appending the platform case data to the AFCARS case data and including a variable for being listed on the platform. This approach dampens the estimated platform effect and its statistical significance, as the 279 children of the platform are double-counted in the AFCARS dataset without the treatment variable. Thus, we expect real treatment effects to be larger than those estimated.

Table 1 presents child statistics for the AFCARS and platform data, highlighting some key differences: The platform’s population of children tends to be older, more male, and more likely to have a clinically diagnosed disability — all factors associated with greater difficulty in placing children. The difference in the age distribution between the AFCARS and platform datasets is statistically significant at the 0.05 level by both a Welch’s t-test and a Mann–Whitney U test.

Table 1 Summary Statistics for AFCARS and Platform Data

<i>Attribute</i>	<i>Covariate</i>	AFCARS (N=10,286)	Platform (N=279)
		<i>Mean (SD) or Percent</i>	<i>Mean (SD) or Percent</i>
Case Duration (years)		1.47 (1.17)	1.60 (0.91)
Adopted before End of Data Horizon		46%	59% [†]
Age at TPR (years)	<i>Age</i>	7.71 (5.16)	8.39 (4.73)
Sex			
Female	<i>Female</i>	48.9%	42.3%
Male		51.1%	57.7%
Race (may be multiple)			
American Indian or Alaskan Native		0.3%	0.0%
Asian		0.6%	0.0%
Black or African American	<i>Black</i>	36.1%	23.3%
Native Hawaiian/Other Pacific Islander		0.2%	0.0%
White		70.5%	65.2%
Other		N/A	11.5%
Hispanic or Latino Ethnicity	<i>Hispanic</i>	14.6%	3.6%
Clinical Disability Diagnosis	<i>Disability</i>	32.7%	37.6%

[†] Includes 144 adoptions via the platform and 21 children listed on the platform but adopted by non-relative/non-foster families found off-platform.

Using a chi-squared test, the difference is also significant at that level for the Female variable but not for the Disability variable. Fewer children on the platform are Black, which reflects regional variations in Florida’s population. Due to differences in how the datasets treat multi-racial children, we provide analysis in Section I.2.3 of the Appendix that shows how estimates of the platform’s performance improve if the Black variable refers only to children with the Black or African American variable exclusively selected as a race variable in the AFCARS dataset.

We also note that the AFCARS dataset’s inability to explicitly identify children in need of adoptive search resources results in a conservative assessment of the platform’s performance; i.e., it will underestimate the platform’s impact. Specifically, managers with experience in Florida child welfare agencies report that some children in the AFCARS dataset may go through the termination of parental rights (TPR) judicial process with an already-identified adoptive placement with a non-relative, such as a teacher, church member, or neighbor.

Both of our statistical approaches model children’s time to placement using a Cox proportional hazards model (Cox 1972). The Cox proportional hazards model includes a baseline hazard function that describes how the likelihood of a placement changes over time, as well as a parameter for each covariate that affects the baseline hazard. The following characteristics were available in both AFCARS and the platform data, allowing us to control for them in the hazards model (for discrete variables, one of the categories is omitted for statistical identification):

1. female (versus male, the omitted category);
2. Black or African-American (versus no such designation, the omitted category);
3. Hispanic or Latino ethnicity designation (versus no such designation, the omitted category);
4. clinical disability diagnosis (versus no diagnosis, the omitted category);

5. age in years upon termination of parental rights; and
6. federal fiscal year (e.g., October 1, 2014, to September 30, 2015, for FY2015) of the termination of the parental rights (TPR) order, with FY2015 as the omitted category.

Focusing on Model 1 — the Cox proportional hazards model without a platform effect that is used for benchmarking — let $t = 0$ denote the TPR date for some child i . The hazard of adoption finalization at time t is given by the semi-parametric Cox model

$$h(t | X_i) = h_0(t) \exp\left(\beta_1 \text{Female}_i + \beta_2 \text{Black}_i + \beta_3 \text{Hispanic}_i + \beta_4 \text{Disability}_i + \beta_5 \text{Age}_i + \beta_6 \text{Age}_i^2 + \sum_{y=2016}^{2021} \beta_y \mathbb{1}\{\text{FY}_i = y\}\right), \quad (13)$$

where $h_0(t)$ is a baseline hazard rate common to all children¹¹ and X_i collects the covariates listed above. The indicator $\mathbb{1}\{\text{FY}_i = y\}$ flags the fiscal year in which TPR occurred. Estimation relies on Cox’s partial likelihood, which is obtained using the hazard rates of all children. Maximizing the log-partial likelihood yields the coefficient estimates. An exponentiated coefficient $\exp(\beta_k) > 1$ indicates that covariate k is associated with a faster-than-baseline adoption rate, whereas $\exp(\beta_k) < 1$ signals a slower rate. We explain how the estimated model in Equation (13) changes when the platform effect is included for the second approach in Section 9.3.

9.2. Benchmark Against Statewide Outcomes

Using only the statewide AFCARS data, we construct a benchmark for the outcomes of children served by the platform. Model 1 of Table 2 shows how demographic factors affect the time to placement using the AFCARS dataset for Florida children, i.e., the maximum-likelihood estimation of the Cox model in Equation (13). Controlling for the fiscal year in which the search started, covariates associated with a faster placement are being female and having Hispanic ethnicity; factors associated with a slower time to or diminished likelihood of adoption are being older, having a disability, and being Black. Using this fitted model, we collected the necessary covariate data on each child served by the platform and predicted the likelihood of a finalized adoption at monthly intervals, as explained below. If a child is adopted, the benchmark adoption probability continues accumulating as if the child’s search continued until the earlier date of February 1, 2023, or the child’s 18th birthday.

By adding the predicted adoption probabilities of all children served by the platform, we can establish a benchmark against which to compare the actual adoptions of children on the platform. Let \mathcal{C} represent the set of children on the platform. For any child $i \in \mathcal{C}$ with attributes X_i , we

¹¹ The baseline hazard rate is estimated using the methodology of Breslow (1972).

Table 2 Cox Proportional Hazards Models of Time Until Adoption

	Model 1	Model 2
Female	1.080** (2.615)	1.089** (2.965)
Black	0.768*** (−8.379)	0.769*** (−8.441)
Hispanic	0.827*** (−4.313)	0.841*** (−3.963)
Age at TPR (years)	0.889*** (−9.563)	0.891*** (−9.507)
(Age at TPR) ²	0.999 (−0.791)	0.999 (−1.072)
Disability	0.881*** (−3.766)	0.866*** (−4.345)
TPR in FY2016	0.967 (−0.605)	0.954 (−0.844)
TPR in FY2017	1.052 (0.918)	1.031 (0.560)
TPR in FY2018	0.927 (−1.419)	0.888* (−2.238)
TPR in FY2019	0.753*** (−5.191)	0.750*** (−5.308)
TPR in FY2020	0.516*** (−10.847)	0.522*** (−10.827)
TPR in FY2021	0.493*** (−8.613)	0.503*** (−8.626)
Platform		1.272** (2.974)
Platform Cases	0	279
N	10,286	10,565
Concordance	0.680	
Log-likelihood ratio test	1969.148 on 12 d.f.	2038.169 on 13 d.f.

Note: exponentiated coefficients; z-statistics in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

can use the survival probability to predict the adoption probability function using the estimated hazard function $\hat{h}(\cdot|X_i)$ in Model 1

$$\hat{\pi}(X_i, t) := 1 - \exp\left(-\int_0^t \hat{h}(u|X_i) du\right) \quad (14)$$

for the likelihood that child i has a finalized adoption within t years.

A child's maximum possible search horizon from TPR until the child turns 18 is denoted by τ_i^e . We use τ_i^d to represent the time between the TPR date and the earliest possible adoption finalization date after registration on the platform. We calculate τ_i^d as the time between the case creation date and the TPR date, plus an additional three months to account for the legally required period in Florida during which a child must reside with the adoptive family before the adoption is finalized.¹²

¹² We provide an alternate analysis in which we assume $\tau_i^d = 0$, as well as an analysis excluding platform children whose placement was identified through other channels in online Appendix I.2.

To account for the time that a child might have already been eligible for adoption before being listed on the platform, we calculate a conditional survival probability for child i at any time $t > \tau_i^d$ since TPR as

$$\tilde{\pi}(X_i, t, \tau_i^e, \tau_i^d) := \frac{\hat{\pi}(X_i, \min\{t, \tau_i^e\}) - \hat{\pi}(X_i, \tau_i^d)}{1 - \hat{\pi}(X_i, \tau_i^d)}, \quad (15)$$

which provides the AFCARS benchmark $\hat{\mu}(t)$ for the number of expected matches by time t :

$$\hat{\mu}(t) := \sum_{i \in \mathcal{C}} \tilde{\pi}(X_i, t, \tau_i^e, \tau_i^d). \quad (16)$$

Figure 5 shows the adoptions achieved through the platform compared to this benchmark from the AFCARS proportional hazards model. The results show that the platform has outperformed the commonly used two-year and three-year search window benchmarks. For children listed on the platform before October 1, 2021, 138 adoptions were finalized within two years — 123 enabled by the platform and 15 through other channels — while the predicted number for children on the platform was only 117.2. At the three-year mark, this difference between adoptions achieved and the predicted number extends to an extra 31.6 adopted children, or a 24% increase over the benchmark. However, the platform’s performance fell short of the AFCARS benchmark one year after the search began, even though the platform not only made up for this gap but surpassed the benchmark at the two-year and three-year marks. The difference in performance may be attributed to a limitation of the AFCARS datasets mentioned above: some cases included in the benchmark might not have actually required a search for a family but rather had a potentially faster and easier non-relative adoptive placement. One subject-matter expert shared that it is common for foster families to network and recruit on behalf of children in their care who might need adoptive placements; these children might have an adoptive resource identified at TPR who is not a relative or foster parent and be included in our benchmark dataset.

Despite this disadvantage, the platform’s over-performance over longer time horizons may indicate that the technology helps caseworkers to be more persistent in their search efforts for hard-to-place children. When Avery (2000) investigated the longest-waiting children in New York, caseworkers were found to be pessimistic about the children’s chances at adoption. They also neglected to use the search tools available to them. However, in addition to the evidence of usage from the actual placements generated by the platform, a survey of the platform’s 73 active users in Florida suggests a positive assessment of caseworker satisfaction with the new CS-driven system. A total of 51 respondents — mostly caseworkers but also some supervisors and recruitment specialists — responded to the survey in 2023. About 30% of the respondents had been in their current role for at least five years, so many of the respondents were familiar with other search methods besides




Figure 5 Actual adoptions by the agency using the platform and other channels compared to Florida AFCARS benchmark model. Dashed lines represent the 95% confidence interval defined by the Poisson binomial distribution.

the platform. When asked how satisfied they were with the platform, approximately 33% were very satisfied, and 31% were somewhat satisfied. Only 5% were very dissatisfied, and only 9% were somewhat dissatisfied. Thus, this case study shows that providing leads of potential adoptive families may encourage caseworkers to persist in finding families.

9.3. Platform Effect in Cox Proportional Hazards Model

To provide additional evidence of the platform’s effectiveness compared to statewide outcomes, we re-estimated the Cox proportional hazards model on a combined dataset to verify that engagement with the online adoption matching platform speeds the journey from TPR to adoption finalization. Because AFCARS restrictions bar direct record linkage and obscure dates of birth (i.e., to the month of birth), a given child can appear in both treatment and control groups — an approach that biases estimates toward understating, rather than exaggerating, any positive or negative platform effect.

Since there is a significant difference between the TPR date for some children and the date on which they were added to the platform, we treat the platform variable as time-varying: the child’s timeline starts at TPR, with the platform variable’s value switching from zero to one at a point three months after the platform case is created. The three-month delay accounts for the minimum period Florida law requires before an adoption can be finalized after placement. Thus, any adoption within the three months after listing cannot be attributed to the platform’s search.

More precisely, when compared with the Cox proportional hazards model introduced in Equation (13), we add a binary covariate to account for whether the child has been on the platform or not. We make the term $OnPlatform_i(t)$ time dependent, and the binary variable $OnPlatform_i(t)$ is

turned on three months after the child is listed on the platform. Each case’s timeline still begins on the TPR date.¹³

We call this Model 2, with resulting coefficients found in Table 2. We see that control variables behave as expected. Girls experience higher adoption hazards than boys, while Black and Hispanic children tend to be adopted more slowly than children of other racial groups. A clinical disability diagnosis and a child’s increasing age are also linked to decreased adoption hazards, as well as TPR in pandemic years or immediately before.

This analysis finds that the platform is consistently associated with faster adoption. The estimated platform hazard ratio is 1.272, indicating a 27.2% increase in the hazard rates compared to not using the platform (i.e., faster adoptions on average). This ratio is significant at the 1% level, despite the biases against the platform built into the dataset construction mentioned above. Taken together, these results confirm that the agency has seen accelerated adoption and finalization rates compared to statewide outcomes after implementing the platform.

Conclusion

Treating the search for adoptive families as an operations challenge, we develop the first formal game-theoretic model for the child welfare system adoption process and introduce a novel search-and-matching framework to compare caseworker-driven and family-driven search paradigms. We characterize the Nash equilibria of these models, showing that agents adopt threshold strategies under a mild tie-breaking condition and that equilibria form a non-empty complete lattice. We find that caseworker-driven search better avoids wasted search efforts and, in most settings, improves outcomes. While caseworker-driven equilibria can Pareto dominate all family-driven equilibria, the converse does not hold. Although caseworker-driven equilibria do not always Pareto dominate those of family-driven search, they typically perform better for most agents, as confirmed through extensive numerical analysis. However, we also identify conditions — notably, highly patient agents with strongly correlated preferences — under which caseworker-driven search may harm children. Yet such conditions are rare in practice: families usually desire a timely placement after completing the adoption approval process, and preferences tend to be idiosyncratic. When families are sufficiently impatient, caseworker-driven search unambiguously benefits all children.

Empirically, we analyze adoption outcomes for 279 children in need of adoptive placements who were served by a Florida agency that shifted from family-driven to caseworker-driven search. Using Cox proportional hazards models that control for observable characteristics as a benchmark, we find a 24% increase in the likelihood of adoption finalization within three years. Similarly, a hazard

¹³ We provide an alternative analysis with non-time-varying platform variables, that is, where the timeline for platform children begins once they are listed on the platform, in online Appendix I.2.

model fitted on a combined dataset indicates 27% increases in the instantaneous adoption rate despite conservative bias in the data construction.

Our work provides a foundation for future empirical research into managing the operation of adoptive placement searches, especially how to allocate caseworkers' efforts across different search channels. While existing research in child welfare literature emphasizes the effectiveness of highly specialized recruiters, we show how technology can help frontline caseworkers achieve better outcomes for children. Furthermore, our benchmarking approach that utilizes AFCARS data provides a replicable tool for evaluating agency performance and child welfare system reforms. Additional study of caseworker behavior, family engagement, and children's outcomes under different search practices is essential to improving services for this vulnerable population.

References

- Abdulkadiroğlu A, Che YK, Yasuda Y (2015) Expanding “Choice” in School Choice. *American Economic Journal: Microeconomics* 7(1):1–42.
- Adachi H (2003) A search model of two-sided matching under nontransferable utility. *Journal of Economic Theory* 113(2):182–198.
- Akbarpour M, Combe J, He Y, Hiller V, Shimer R, Tercieux O (2020a) Unpaired kidney exchange: Overcoming double coincidence of wants without money. *Proceedings of the 21st ACM Conference on Economics and Computation*, 465–466.
- Akbarpour M, Li S, Gharan SO (2020b) Thickness and information in dynamic matching markets. *Journal of Political Economy* 128(3):783–815.
- Albrecht JW, Gautier PA, Vroman SB (2006) Equilibrium directed search with multiple applications. *Review of Economic Studies* 73(4):869–891, URL <http://dx.doi.org/10.1111/j.1467-937X.2006.00400.x>.
- Altinok A, MacDonald D (2023) Designing the menu of licenses for foster care. *Available at SSRN 4466506*.
- Andersson T, Ehlers L, Martinello A (2018) Dynamic Refugee Matching. Technical report, Lund University, Lund, Sweden.
- Arnosti N, Johari R, Kanoria Y (2021) Managing congestion in matching markets. *Manufacturing & Service Operations Management* 23(3):620–636.
- Arnosti N, Shi P (2020) Design of lotteries and wait-lists for affordable housing allocation. *Management Science* 66(6):2291–2307.
- Atakan AE (2006) Assortative Matching with Explicit Search Costs. *Econometrica* 74(3):667–680.
- Auster S, Gottardi P, Wolthoff RP (2025) Simultaneous search and adverse selection. *Review of Economic Studies* URL <http://dx.doi.org/10.1093/restud/rdaf014>, advance online publication.
- Avery RJ (2000) Perceptions and practice: Agency efforts for the hardest-to-place children. *Children and Youth Services Review* 22(6):399–420.
- Avery RJ, Butler J, Schmidt EB, Holtan BA (2009) AdoptUsKids national photolisting service: Characteristics of listed children and length of time to placement. *Children and Youth Services Review* 31(1):140–154.
- Baccara M, Collard-Wexler A, Felli L, Yariv L (2014) Child-Adoption Matching: Preferences for Gender and Race. *American Economic Journal: Applied Economics* 6(3):133–158.
- Baccara M, Lee S, Yariv L (2020) Optimal dynamic matching. *Theoretical Economics* 15(3):1221–1278.
- Bansak K, Ferwerda J, Hainmueller J, Dillon A, Hangartner D, Lawrence D, Weinstein J (2018) Improving refugee integration through data-driven algorithmic assignment. *Science* 359(6373):325–329.

- Berenguer G, Haskell WB, Li L (2023) Managing volunteers and paid workers in a nonprofit operation. *Management Science*, *Forthcoming* .
- Breslow NE (1972) Contribution to discussion of paper by dr cox. *Journal of the Royal Statistical Society, Series B* 34:216–217.
- Chade H (2001) Two-sided search and perfect segregation with fixed search costs. *Mathematical Social Sciences* 42(1):31–51.
- Chade H, Eeckhout J, Smith L (2017) Sorting through Search and Matching Models in Economics. *Journal of Economic Literature* 55(2):493–544, ISSN 0022-0515, URL <http://dx.doi.org/10.1257/jel.20150777>.
- Chade H, Smith L (2006) Simultaneous Search. *Econometrica* 74(5):1293–1307, ISSN 0012-9682, 1468-0262, URL <http://dx.doi.org/10.1111/j.1468-0262.2006.00705.x>.
- Cheremukhin A, Restrepo-Echavarria P, Tutino A (2020) Targeted search in matching markets. *Journal of Economic Theory* 185:104956.
- Children’s Bureau (2022) The AFCARS report: Preliminary FY2021 estimates. URL <https://www.acf.hhs.gov/sites/default/files/documents/cb/afcars-report-29.pdf>, accessed: 2024-02-28.
- Children’s Bureau, Administration on Children, Youth and Families (2023a) Afcars adoption file 2021 [Data set]. URL <https://doi.org/10.34681/psb7-a026>, accessed: 2024-02-28.
- Children’s Bureau, Administration on Children, Youth and Families (2023b) AFCARS Foster Care File, 6-month periods (FY2016A - 2022B). URL <https://doi.org/10.34681/yjxr-zz92>, accessed: 2024-02-28.
- Combe J, Dur U, Tercieux O, Terrier C, Ünver MU (2025) Market design for distributional objectives in (re) assignment: An application to improve the distribution of teachers in schools, working Paper, SSRN: 5112041.
- Combe J, Dur U, Tercieux O, Terrier C, Ünver MU (2022) Market design for distributional objectives in (re)assignment: An application to improve the distribution of teachers in schools. Technical report, Boston College, Department of Economics.
- Cox DR (1972) Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)* 34(2):187–202.
- Delacrétaz D, Kominers SD, Teytelboym A (2020) Matching Mechanisms for Refugee Resettlement. Technical report, University of Oxford, Oxford, UK.
- Dierks L, Slaugh V, Ünver MU (2024) Child welfare platform design to improve outcomes for children with disabilities. *Available at SSRN 4778791* .
- Eeckhout J (1999) Bilateral Search and Vertical Heterogeneity. *International Economic Review* 40(4):869–887.

- Feldman SW, Price KM, Ruppel J (2016) Not too late: Effects of a diligent recruitment program for hard to place youth. *Children and Youth Services Review* 65:26–31.
- Florida Department of Children and Families (2019) Adoption incentive annual report. Office of Child Welfare, URL <https://www.myflfamilies.com/sites/default/files/2023-02/2019%2520Adoption%2520Incentive%2520Report.pdf>, accessed: 2024-02-28.
- Florida Department of Children and Families (2022) A comprehensive, multi-year review of the revenues, expenditures, and financial position of all community-based care lead agencies with system of care analysis. Technical report, Florida Department of Children and Families, Tallahassee, FL, URL <https://www.myflfamilies.com/programs/childwelfare/dashboard/>, state Fiscal Years 2020–2021 and 2021–2022.
- Florida Department of Children and Families (2024) Adoption incentive annual report. Office of Child and Family Well-Being, URL <https://www.myflfamilies.com/sites/default/files/2024-12/2024%20Annual%20Adoption%20Incentive%20Report.pdf>, accessed: 2025-03-11.
- Fradkin A (2017) Search, Matching, and the Role of Digital Marketplace Design in Enabling Trade: Evidence from Airbnb. Technical report, Boston University Questrom School of Business.
- Gale D, Sotomayor M (1985) Ms. Machiavelli and the stable matching problem. *The American Mathematical Monthly* 92(4):261–268.
- Gypen L, Vanderfaeillie J, De Maeyer S, Belenger L, Van Holen F (2017) Outcomes of children who grew up in foster care: Systematic-review. *Children and Youth Services Review* 76:74–83.
- Halaburda H, Piskorski MJ, Yildirim P (2018) Competing by restricting choice: The case of matching platforms. *Management Science* 64(8):3574–3594, URL <http://dx.doi.org/10.1287/mnsc.2017.2797>.
- Hanna M, McRoy R (2011) Innovative Practice Approaches to Matching in Adoption. *Journal of Public Child Welfare* 5(1):45–66.
- Hitsch GJ, Hortacısu A, Ariely D (2010a) Matching and Sorting in Online Dating. *American Economic Review* 100(1):130–163.
- Hitsch GJ, Hortacısu A, Ariely D (2010b) What makes you click?—Mate preferences in online dating. *Quantitative Marketing and Economics* 8(4):393–427.
- Honka E, Chintagunta P (2017) Simultaneous or sequential? search strategies in the U.S. auto insurance industry. *Marketing Science* 36(1):21–42, URL <http://dx.doi.org/10.1287/mksc.2016.0995>.
- Immorlica N, Lucier B, Manshadi V, Wei A (2023) Designing approximately optimal search on matching platforms. *Management Science* 69(8):4609–4626, URL <http://dx.doi.org/10.1287/mnsc.2022.4601>.
- Kanoria Y, Saban D (2021) Facilitating the Search for Partners on Matching Platforms. *Management Science* .
- Kasy M, Teytelboym A (2020) Adaptive Combinatorial Allocation. *arXiv:2011.02330 [econ, stat]* ArXiv: 2011.02330.

- Kawasaki T, Wada R, Todo T, Yokoo M (2021) Mechanism design for housing markets over social networks. *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems*, 692–700.
- Kerimov S, Ashlagi I, Gurvich I (2023) On the optimality of greedy policies in dynamic matching. *Operations Research* .
- Kircher P (2009) Efficiency of simultaneous search. *Journal of Political Economy* 117(5):861–913, URL <http://dx.doi.org/10.1086/644791>.
- Knuth DE (1997) *Stable marriage and its relation to other combinatorial problems: An introduction to the mathematical analysis of algorithms*, volume 10 (American Mathematical Soc.).
- Kushel MB, Yen IH, Gee L, Courtney ME (2007) Homelessness and health care access after emancipation: Results from the midwest evaluation of adult functioning of former foster youth. *Archives of pediatrics & adolescent medicine* 161(10):986–993.
- Lauermann S, Nöldeke G (2014) Stable marriages and search frictions. *Journal of Economic Theory* 151:163–195.
- Lauermann S, Nöldeke G, Tröger T (2020) The Balance Condition in Search-and-Matching Models. *Econometrica* 88(2):595–618.
- Lee S, Niederle M (2015) Propose with a rose? Signaling in internet dating markets. *Experimental Economics* 18(4):731–755.
- Leshno JD (2022) Dynamic matching in overloaded waiting lists. *American Economic Review* 112(12):3876–3910.
- Lushin V, Katz CC, Julien-Chinn FJ, Lalayants M (2023) A burdened workforce: Exploring burnout, job satisfaction and turnover among child welfare caseworkers in the era of covid-19. *Children and Youth Services Review* 148:106910.
- Ma H, Fang F, Parkes DC (2020) Spatio-Temporal Pricing for Ridesharing Platforms. Technical report, Columbia Business School, 3022 Broadway, Uris Hall 423, New York, NY, 10027, USA.
- MacDonald DE (2019) Foster Care: A Dynamic Matching Approach. Technical report, Arizona State University, Phoenix, AZ, USA.
- Mennle T, Weiss M, Philipp B, Seuken S (2015) The Power of Local Manipulation Strategies in Assignment Mechanisms. *Proceedings of the 24th International Conference on Artificial Intelligence*, 82–89, IJCAI’15, ISBN 978-1-57735-738-4.
- Quast T, Storch EA, Yampolskaya S (2018) Opioid prescription rates and child removals: Evidence from florida. *Health Affairs* 37(1):134–139.
- Riley NS (2019) Adoptions powered by algorithms. *Wall Street Journal* URL <https://www.wsj.com/articles/adoptions-powered-by-algorithms-11546620390>, accessed: 2024-02-28.

- Robinson-Cortés A (2019) Who Gets Placed Where and Why? An Empirical Framework for Foster Care Placement. Technical report, California Institute of Technology, 315 Baxter Hall, Pasadena, CA 91125, USA.
- Roby JL, White H (2010) Adoption activities on the internet: A call for regulation. *Social Work* 55(3):203–212.
- Roth AE (1984) The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory. *Journal of Political Economy* 92(6):991–1016.
- Roth AE (1991) A Natural Experiment in the Organization of Entry-Level Labor Markets: Regional Markets for New Physicians and Surgeons in the United Kingdom. *The American Economic Review* 81(3):415–440.
- Shi P (2023) Optimal matchmaking strategy in two-sided marketplaces. *Management Science* 69(3):1323–1340.
- Shimer R, Smith L (2000) Assortative Matching and Search. *Econometrica* 68(2):343–369.
- Shimer R, Smith L (2001) Matching, Search, and Heterogeneity. *Advances in Macroeconomics* 1(1).
- Slaugh VW, Akan M, Kesten O, Ünver MU (2016) The Pennsylvania Adoption Exchange Improves Its Matching Process. *Interfaces* 46(2):133–153.
- Slaugh VW, Dierks L, Scheller-Wolf A, Trapp AC, Ünver MU (2025) Child welfare services in the united states: An operations research perspective. *Nonprofit Operations and Supply Chain Management: Theory and Practice* 321–347.
- Smith L (2006) The Marriage Model with Search Frictions. *Journal of Political Economy* 114(6):1124–1144, ISSN 0022-3808, 1537-534X, URL <http://dx.doi.org/10.1086/510440>.
- Sönmez T, Ünver MU, Yenmez MB (2020) Incentivized kidney exchange. *American Economic Review* 110(7):2198–2224.
- Spindler A (1970) Social and rehabilitation services: A challenge to operations research. *Operations Research* 18(6):1112–1124.
- Stigler GJ (1961) The economics of information. *Journal of Political Economy* 69(3):213–225, URL <http://dx.doi.org/10.1086/258464>.
- Tarski A (1955) A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics* 5(2):285–309, ISSN 0030-8730, 0030-8730, URL <http://dx.doi.org/10.2140/pjm.1955.5.285>.
- Triseliotis J (2002) Long-term foster care or adoption? The evidence examined. *Child & Family Social Work* 7(1):23–33.
- US Administration for Children and Families (2024) AFCARS Adoption and Foster Care Analysis and Reporting System Dashboard. URL https://tableau-public.acf.gov/views/afcars_dashboard_main_page/mainpage, accessed: 2025-06-30.

- Vandivere S, Malm KE, Zinn A, Allen TJ, Mcklindon A (2015) Experimental evaluation of a child-focused adoption recruitment program for children and youth in foster care. *Journal of Public Child Welfare* 9(2):174–194.
- Weitzman ML (1979) Optimal Search for the Best Alternative. *Econometrica* 47(3):641.
- Wright R, Kircher P, Julien B, Guerrieri V (2021) Directed search and competitive search equilibrium: A guided tour. *Journal of Economic Literature* 59(1):90–148.
- Yamatani H, Engel R, Spjeldnes S (2009) Child welfare worker caseload: What’s just right? *Social Work* 54(4):361–368.
- You B, Dierks L, Todo T, Li M, Yokoo M (2022) Strategy-proof house allocation with existing tenants over social networks. *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, 1446–1454.
- Ünver MU (2010) Dynamic Kidney Exchange. *Review of Economic Studies* 77(1):372–414, ISSN 00346527.

e-Companion

A. Home Studies

Before a family can be considered for adoption, they have to complete a home study. State regulations determine minimum requirements for home studies. For example, 2024 Florida Statutes Chapter 63 Section 092 states:

The preliminary home study must be made to determine the suitability of the intended adoptive parents and may be completed prior to the identification of a prospective adoptive child. The study must include, at a minimum, the following:

- *An interview with the intended adoptive parents.*
- *Records checks of the department's central abuse registry, which the department shall provide to the entity conducting the preliminary home study, and criminal records correspondence checks under s. 39.0138 through the Department of Law Enforcement on the intended adoptive parents.*
- *An assessment of the physical environment of the home.*
- *A determination of the financial security of the intended adoptive parents.*
- *Documentation of counseling and education of the intended adoptive parents on adoptive parenting, as determined by the entity conducting the preliminary home study. The training specified in s. 409.175(14) shall only be required for persons who adopt children from the department.*
- *Documentation that information on adoption and the adoption process has been provided to the intended adoptive parents.*
- *Documentation that information on support services available in the community has been provided to the intended adoptive parents.*
- *A copy of each signed acknowledgment of receipt of disclosure required by s. 63.085.*

However, most agencies perform more comprehensive studies than required by law to ensure a smoother process. In practice, home studies typically also include additional information on

- Family background, including education and health
- Relationships and social environment
- Parenting experiences
- A description of daily life routines
- Details about home and neighborhood
- Reasons for seeking an adoptive placement
- Recommendations by the caseworker about what types of children the family is suited for (e.g., especially regarding special needs, but also other factors such as gender or age)

B. Increased Uncertainty or Heterogeneity

Our analytical model assumes a relatively limited amount of variability and uncertainty: the value of every match is known a priori. The only uncertainty is whether the child is compatible with the considered family, with the suitability probability p being the same for all matches. This is not without loss of generality. In practice, compatibility is a spectrum, and values may be uncertain. This means that while caseworkers and families may have prior beliefs based on their own experiences and potentially supplemented by recommender systems, the actual value of a match is not known until the investigation is conducted. Unfortunately, explicitly modeling such higher degrees of uncertainty is intractable in a model where both market sides are strategic and more than one match may be investigated per time period.

This is because if values are uncertain, family utility is no longer monotonic in terms of how selective a child is. While children being less selective (i.e., considering families it has lower expected value for) still increases a family's utility if it causes the child to be interested in them, it often reduces a family's utility if the child was already interested in them (as the child has some probability of instead matching with a different family they didn't even consider before). This may lead to non-convergence of best responses and non-existence of equilibria in pure strategy, even if only either the suitability or market thickness is heterogeneous.

PROPOSITION 10. *If suitability probabilities are heterogeneous, i.e., if there exist different $p_{c,f}$ for different child/family pairs, or if families have heterogeneous probabilities λ_f to be present, then pure strategy equilibria do not always exist.*

Proof. Consider the following example in FS where each child/family pair has its own probability $p_{c,f}$ for being suitable. Let $C = \{c_1, c_2\}$, $F = \{f_1, f_2, f_3\}$, and let valuations and suitability probabilities be according to the following tables

$v_c(f)$	f_1	f_2	f_3
c_1	2	1	0.5
c_2	0.5	0	1

$v_f(c)$	f_1	f_2	f_3
c_1	20	0.11	2
c_2	5	0	2

$p_{c,f}$	f_1	f_2	f_3
c_1	0.1	0.99	1
c_2	0.5	0.5	0.5

Let $\lambda = 1$ (i.e., all family types are always present). Further, assume children are very impatient and have negligible search costs, i.e. $\delta_C = \kappa_C = 0$, while families are very patient $\delta_F = 1$ and have search cost $\kappa_F = 0.1$. Note that this implies that a) both children are so impatient, that they will be interested in all families and b) families will be so patient, that they are at most interested in a single-child type (unless both children types give almost the same expected utility, which cannot happen in pure strategy with these parameters). Note that setting these extreme values is not required for equilibrium non-existence and is only done for brevity.

In this setting, it is not rational for f_2 to be interested in c_2 , independent of the other agent's strategies. Further, if f_1 is interested in c_1 in equilibrium, f_2 cannot be interested in c_1 as doing so would result in negative utility, as c_1 prefers f_1 to f_2 : the probability that c_1 doesn't successfully match with f_1 times the probability that f_2 and c_1 are compatible times $v_{f_2}(c_1)$ equals $(1 - 0.1) \times 0.99 \times 0.11 = 0.09801 \leq 0.1 = \kappa_F$. However, if f_1 is not interested in c_1 , f_2 would obtain positive utility (i.e., $0.99 \times 0.11 = 0.1089 \geq 0.1$) and therefore would be interested.

Thus, we only have 4 possible equilibrium candidates, depending on in which child families f_1 and f_3 are interested.

1. If f_1 and f_3 are interested in c_1 in equilibrium, then f_1 will obtain utility $0.1 \times 20 - 0.1 = 1.9$. However, f_1 deviating to be (solely) interested in c_2 would yield utility $0.5 \times 5 - 0.1 = 2.4 > 1.9$, a contradiction.

2. If f_1 is interested in c_1 and f_3 is interested in c_2 in equilibrium, then f_3 will obtain utility $0.5 \times 2 - 0.1 = 0.9$. However, f_3 deviating to be (solely) interested in c_1 would yield utility $(1 - 0.1) \times 1 \times 2 - 0.1 = 1.7 > 0.9$, a contradiction.

3. If f_1 and f_3 are interested in c_2 in equilibrium, then f_1 will only be chosen if f_3 is incompatible, thus obtaining expected utility $0.5 \times 0.5 \times 5 - 0.1 = 1.15$. However, f_1 deviating to be (solely) interested in c_1 would yield utility $0.1 \times 20 - 0.1 = 1.9 > 1.15$, a contradiction.

4. If f_1 is interested in c_2 and f_3 is interested in c_1 in equilibrium, then f_2 is interested in c_1 . Since c_1 prefers f_2 and has a $p_{c_1, f_2} = 0.99$ likelihood of being compatible, this implies that f_3 only obtains utility $(1 - 0.99) \times 1 \times 2 - 0.1 = -0.08$, an immediate contradiction.

In conclusion, there can be no equilibrium in pure strategies, because f_1 always prefers the child type f_3 is not interested in, while f_3 always prefers the child type f_1 is interested in.

To see the nonexistence of equilibria for heterogeneous λ_f , consider the following example. Let $C = \{c_1, c_2\}$, $F = \{f_1, f_2, f_3\}$ and let valuations and λ_f be according to the following tables

$v_c(f)$	f_1	f_2	f_3	$v_f(c)$	f_1	f_2	f_3		f_1	f_2	f_3
c_1	2	1	0.5	c_1	4	0.11	2		0.1	0.99	1
c_2	0.5	0	1	c_2	5	0	1	λ_f			

Let $p = 1$ (i.e., all families are suitable). Further, assume children are impatient and have negligible search costs, i.e. $\delta_C = \kappa_C = 0$, while families are very patient $\delta_F = 1$ and have search cost $\kappa_F = 0.1$. Note that this implies that a) both children are so impatient, that they will be interested in all families and b) families will be so patient, that they are at most interested in a single-child type (unless both children types give almost the same expected utility, which cannot happen in pure strategy with these parameters). Note that setting these extreme values are not required for equilibrium non-existence and is only done for brevity.

In this setting, it is not rational for f_2 to be interested in c_2 , independent of the other agent's strategies. Further, if f_1 is mutually interested in c_1 in equilibrium, f_2 cannot be interested in c_1 as doing so would result in negative utility: due to c_1 preferring f_1 whenever present (10% of the time), the immediate expected value of showing interest is only $(1 - 0.1) \times 0.11 = 0.099$, less than the search cost of 0.1. However, if f_1 is not mutually interested in c_1 , f_2 is the child's first choice and would obtain positive utility (i.e., $0.11 \geq 0.1$). Therefore, f_2 would be interested.

Thus, we only have 4 possible equilibrium candidates, depending on which child families f_1 and f_3 are interested in.

1. If f_1 and f_3 are interested in c_1 in equilibrium, f_1 will obtain utility $4 - 0.1 = 3.9$. However, f_1 deviating to be (solely) interested in c_2 would yield utility $5 - 0.1 = 4.9 > 3.9$, a contradiction.

2. If f_1 is interested in c_1 and f_3 is interested in c_2 in equilibrium, f_3 will obtain utility $1 - 0.1 = 0.9$. However, f_3 deviating to be (solely) interested in c_1 , it would successfully match whenever f_1 is not present, yielding utility $(1 - 0.1) \times 2 - 0.1 = 1.7 > 0.9$, a contradiction.

3. If f_1 and f_3 are interested in c_2 in equilibrium, f_1 will obtain no utility, as c_2 will always match successfully with f_3 , a contradiction since f_1 can guarantee positive utility by being interested in c_1 .

4. If f_1 is interested in c_2 and f_3 is mutually interested in c_1 in equilibrium, f_2 is interested in c_1 . Since c_1 prefers f_2 over f_3 , f_3 only obtains utility $(1 - 0.99) \times 1 \times 2 - 0.1 = -0.08$, an immediate contradiction.

In conclusion, there can be no equilibrium in pure strategies, because f_1 always prefers the child type f_3 is not interested in, while f_3 always prefers the child type f_1 is interested in.

□

As the same effect can be recreated by uncertainty about match values, we immediately obtain the following.

COROLLARY 1. *If values are uncertain, pure strategy equilibria do not always exist.*

While equilibria still exist in mixed strategy (i.e., where agents play randomized strategies), it is intractable to fully characterize or calculate mixed strategy equilibria in such a complex system.

While we, therefore, restrict our formal analysis to a low level of uncertainty, it is important to note that uncertainty typically favors CS: In CS, the caseworker can decide whether to investigate another family based on the expected utility gain while already knowing the true value of previously investigated matches. Additionally, investigating a family does not prevent the caseworker from going back and matching with previously investigated families. This allows caseworkers more flexibility than FS, where all investigation decisions must be made simultaneously, increasing the chance of “wasted” investigations.

However, our insight that neither approach always dominates the other still holds. Similar to before, the lower cost of being mutually interested in CS changes strategic considerations on both sides of the market and can lead to more matches in FS (e.g., if families are so patient that they are only interested in a small set of very attractive children in CS).

C. Additional Lemmas and Propositions

C.1. Proposition 11

Here, we show that processing families in decreasing order of $v_c(f)$ is optimal for children in CS.

PROPOSITION 11. *In CS, c 's utility is maximized if the caseworker processes families in decreasing order of $v_c(f)$.*

Proof. Assume child c is active at the current time step. Note that families without interest in c do not affect c 's utility in any way. For the remaining families, c faces a Pandora's box problem (Weitzman 1979) where c receives a payoff of $v_c(f)$ with probability p and a payoff of 0 with probability $1 - p$ when the box corresponding to family f is opened. Notice that the reservation value of the box corresponding to family f is higher than for family f' if and only if $v_c(f) > v_c(f')$.

C.2. Proposition 12

Strategy s_f is a simple threshold strategy with threshold $z \in \mathbb{R}$ for family f if $s_f = \mathbb{1}[v_f(c) \geq z]$, $\forall c \in C$. Here, we show that there exist instances where families cannot best respond with a simple threshold strategy.

PROPOSITION 12. *In FS, there exists an instance with a family f and other agents' strategies s_{-f} , such that no simple threshold strategy s_f is a best response for f .*

Proof. Consider the following example: Let $C = \{c_1, c_2\}$, $F = \{f_1, f_2\}$ and let valuations be according to the following tables for some $\epsilon > 0$.

$v_c(f)$	f_1	f_2
c_1	1	$1 - \epsilon$
c_2	1	$1 - \epsilon$

$v_f(c)$	f_1	f_2
c_1	1	1
c_2	$1 - \epsilon$	$1 - \epsilon$

Suppose strategy profile s is such that $s_c(f) = s_f(c) = 1$ for all $c \in C$, $f \in F$. If ϵ is small enough and $(1 - \lambda p)p < \kappa_F \leq p$, then it is optimal for f_2 to only be interested in c_2 , even though f_2 strictly prefers c_1 .

D. Remaining Proofs

D.1. Proof of Proposition 1

In FS, one way to express c 's utility is as follows:

$$u_c^{FS}(s) = \left[1 - \lambda p \sum_{f \in M_c(s)} \beta_{cf}(s) \right] \delta_C u_c^{FS}(s) + \lambda \sum_{f \in M_c(s)} \left[\beta_{cf}(s) p v_c(f) - \kappa_C \right]. \quad (17)$$

The probability of a match forming between c and f at the current time step is $\lambda p \beta_{cf}(s)$ if there is mutual interest, in which case c obtains a value of $v_c(f)$ and leaves the process. For any active family that showed interest, c incurs search costs κ_C . If c remains unmatched, then c receives $\delta_C u_c^{FS}(s)$. By pulling $\delta_C u_c^{FS}(s)$ out of the sums, we get

$$u_c^{FS}(s) = \delta_C u_c^{FS}(s) + \lambda \sum_{f \in M_c(s)} \left(\beta_{cf}(s) p (v_c(f) - \delta_C u_c^{FS}(s)) - \kappa_C \right). \quad (18)$$

The proof for families' utilities in FS and agents' utilities in CS is analogous and omitted.

D.2. Proof of Proposition 2

First of all, notice that whether agent i is interested in some agent j or not does not affect agent i 's utility if j is not interested in i . Further, for an arbitrary family f in FS, $\beta_{cf}(s)$ does not depend on s_f . By slightly modifying Equation (6), we can see that when f plays a best response in s it must hold that

$$u_f^{FS*}(s_{-f}) = \delta_F u_f^{FS*}(s_{-f}) + \frac{1}{n} \sum_{c \in M_f(s)} \left(\beta_{cf}(s) p (v_c(c) - \delta_F u_f^{FS*}(s_{-f})) - \kappa_F \right). \quad (19)$$

By Equation (19) it must hold for all $c \in C$ that

$$s_f(c) = \mathbb{1}[\beta_{cf}(s) p (v_c(f) - \delta_F u_f^{FS*}(s_{-f})) \geq \kappa_F] \quad (20)$$

when there is mutual interest between c and f , as s_f would otherwise not be a best response. That is, because all children c contribute non-negatively to f 's utility if and only if $\beta_{cf}(s) p (v_c(f) - \delta_F u_f^{FS*}(s_{-f})) \geq \kappa_F$. By our tie-breaking assumption, the claim of the proposition follows for families in FS. The proof for families in CS is analogous and thus omitted.

In the remainder of the proof, we show that the statement holds for children in FS. The proof for CS is again analogous and therefore omitted. Let c be an arbitrary child. For a best response s_c to s_{-c} we have that

$$u_c^{FS*}(s_{-c}) = \delta_C u_c^{FS*}(s_{-c}) + \lambda \sum_{f \in F} s_c(f) s_f(c) \left(\beta_{cf}(s) p (v_c(f) - \delta_C u_c^{FS*}(s_{-c})) - \kappa_C \right). \quad (21)$$

As for families, it must be the case that

$$s_c(f) = \mathbb{1}[\beta_{cf}(s) p (v_c(f) - \delta_C u_c^{FS*}(s_{-c})) \geq \kappa_C] \quad (22)$$

because for all $f, f' \in F$

$$\beta_{cf}(s) \geq \beta_{cf'}(s) \iff v_c(f) \geq v_c(f'), \quad (23)$$

and otherwise s_c would not be a best response to s_{-c} . Again, by our tie-breaking assumption, the claim of the proposition follows for children in FS.

D.3. Proof of Proposition 3

For FS, define a mapping $T^{FS} : Y \rightarrow Y$ as follows: $T^{FS} = (T_i^{FS})_{i \in A}$, where

$$T_c^{FS}(y) = \delta_C y_c + \lambda \sum_{f \in F} \mathbb{1}[\beta_{cf}(y)p(v_f(c) - \delta_F y_f) \geq \kappa_F] \left(\beta_{cf}(y)p(v_c(f) - \delta_C y_c) - \kappa_C \right)^+ \quad (24)$$

for all $c \in C$ and

$$T_f^{FS}(y) = \delta_F y_f + \frac{1}{n} \sum_{c \in C} \mathbb{1}[\beta_{cf}(y)p(v_c(f) - \delta_C y_c) \geq \kappa_C] \left(\beta_{cf}(y)p(v_f(c) - \delta_F y_f) - \kappa_F \right)^+. \quad (25)$$

for all $f \in F$. Note that any fixed point of T^{FS} (i.e., any y with $T^{FS}(y) = y$) is an equilibrium threshold profile in FS. We now show that T^{FS} is monotonically increasing according to \leq_C . Let $c \in C$, $y, y' \in Y$, and $y \leq_C y'$. We have

$$T_c^{FS}(y) = \delta_C y_c + \lambda \sum_{f \in F} \mathbb{1}[\beta_{cf}(y)p(v_f(c) - \delta_F y_f) \geq \kappa_F] \left(\beta_{cf}(y)p(v_c(f) - \delta_C y_c) - \kappa_C \right)^+ \quad (26)$$

$$\leq \delta_C y'_c + \lambda \sum_{f \in F} \mathbb{1}[\beta_{cf}(y')p(v_f(c) - \delta_F y'_f) \geq \kappa_F] \left(\beta_{cf}(y')p(v_c(f) - \delta_C y'_c) - \kappa_C \right)^+ \quad (27)$$

$$= T_c^{FS}(y'). \quad (28)$$

The inequality holds for the following reason: Suppose a family f is interested in c under $s(y)$ but not under $s(y')$. Since f is weakly less selective in $s(y')$, it must be the case that there exists another family f' with $v_c(f') > v_c(f)$ that is not mutually interested in c under $s(y)$ but under $s(y')$. Note that for any family that loses interest in c under $s(y')$, there must exist such a unique family that replaces it and is preferred by c .

Since each child c is weakly more selective under $s(y')$ than $s(y)$, we have for all $f \in F$

$$T_f^{FS}(y) = \delta_F y_f + \frac{1}{n} \sum_{c \in C} \mathbb{1}[\beta_{cf}(y)p(v_c(f) - \delta_C y_c) \geq \kappa_C] \left(\beta_{cf}(y)p(v_f(c) - \delta_F y_f) - \kappa_F \right)^+ \quad (29)$$

$$\geq \delta_F y'_f + \frac{1}{n} \sum_{c \in C} \mathbb{1}[\beta_{cf}(y')p(v_c(f) - \delta_C y'_c) \geq \kappa_C] \left(\beta_{cf}(y')p(v_f(c) - \delta_F y'_f) - \kappa_F \right)^+ \quad (30)$$

$$= T_f^{FS}(y'). \quad (31)$$

Note that T^{FS} maps elements from Y to Y and (Y, \leq_C) is a complete lattice. By Tarski's fixed point theorem, the claim follows for FS.

For CS, we define a mapping $T^{CS} : Y \rightarrow Y$ as follows: $T^{CS} = (T_i^{CS})_{i \in A}$, where

$$T_c^{CS}(y) = \delta_C y_c + \lambda \sum_{f \in F} \mathbb{1}[p(v_f(c) - \delta_F y_f) \geq \kappa_F] \beta_{cf}(y) \left(p(v_c(f) - \delta_C y_c) - \kappa_C \right)^+ \quad (32)$$

for all $c \in C$ and

$$T_f^{CS}(y) = \delta_F y_f + \frac{1}{n} \sum_{c \in C} \mathbb{1}[p(v_c(f) - \delta_C y_c) \geq \kappa_C] \beta_{cf}(y) \left(p(v_f(c) - \delta_F y_f) - \kappa_F \right)^+. \quad (33)$$

for all $f \in F$. Note that any fixed point of T^{CS} is an equilibrium threshold profile in CS. We now show that T^{CS} is monotonically increasing according to \leq_C . Let $c \in C$, $y, y' \in Y$, and $y \leq_C y'$. It holds that

$$T_c^{CS}(y) = \delta_C y_c + \lambda \sum_{f \in F} \mathbb{1}[p(v_f(c) - \delta_F y_f) \geq \kappa_F] \beta_{cf}(y) \left(p(v_c(f) - \delta_C y_c) - \kappa_C \right)^+ \quad (34)$$

$$\leq \delta_C y'_c + \lambda \sum_{f \in F} \mathbb{1}[p(v_f(c) - \delta_F y'_f) \geq \kappa_F] \beta_{cf}(y') \left(p(v_c(f) - \delta_C y'_c) - \kappa_C \right)^+ \quad (35)$$

$$= T_c^{CS}(y'), \quad (36)$$

because each family is weakly less selective under $s(y')$ than $s(y)$.

Similarly, since each child is weakly more selective under $s(y')$ than $s(y)$, we have for all $f \in F$

$$T_f^{CS}(y) = \delta_F y_f + \frac{1}{n} \sum_{c \in C} \mathbb{1}[p(v_c(f) - \delta_C y_c) \geq \kappa_C] \beta_{cf}(y) \left(p(v_f(c) - \delta_F y_f) - \kappa_F \right)^+ \quad (37)$$

$$\geq \delta_F y'_f + \frac{1}{n} \sum_{c \in C} \mathbb{1}[p(v_c(f) - \delta_C y_c) \geq \kappa_C] \beta_{cf}(y') \left(p(v_f(c) - \delta_F y'_f) - \kappa_F \right)^+ \quad (38)$$

$$= T_f^{CS}(y'). \quad (39)$$

Note that T^{CS} maps elements from Y to Y and (Y, \leq_C) is a complete lattice. By Tarski's fixed point theorem, the claim for CS follows.

D.4. Proof of Lemma 1

Proof. If $(c, f) \in M(s^{FS})$ and $(c, f) \notin M(s^{CS})$, then it holds that $\beta_{cf}(s^{FS})p(v_c(f) - \delta_C u_c(s^{FS})) \geq \kappa_C$ and $\beta_{cf}(s^{FS})p(v_f(c) - \delta_F u_f(s^{FS})) \geq \kappa_F$. Further, either $p(v_c(f) - \delta_C u_c(s^{CS})) < \kappa_C$ or $p(v_f(c) - \delta_F u_f(s^{CS})) < \kappa_F$. If the former is true we have

$$\delta_C u_c(s^{FS}) \leq v_c(f) - \frac{\kappa_C}{\beta_{cf}(s^{FS})p} \leq v_c(f) - \frac{\kappa_C}{p} < \delta_C u_c(s^{CS}), \quad (40)$$

since $0 < \beta_c(f, s^{FS}) \leq 1$. Otherwise, we get

$$\delta_F u_f(s^{FS}) \leq v_f(c) - \frac{\kappa_F}{\beta_{cf}(s^{FS})p} \leq v_f(c) - \frac{\kappa_F}{p} < \delta_F u_f(s^{CS}). \quad (41)$$

D.5. Proof of Lemma 2.

Proof. If c responds to s_{-c}^{CS} with s_c^{FS} , c is mutually interested in the same families as under s^{FS} since $M_c(s^{FS}) \subseteq M_c(s^{CS})$. Further, c 's expected costs are weakly lower in CS compared to FS. Hence, there exists a strategy for c in CS where c 's utility is weakly higher than $u_c(s^{FS})$. The fact that c plays a best response in s^{CS} completes the proof.

D.6. Proof of Theorem 1

Proof. It is easy to see why all CSEs can be Pareto improvements over all FSEs: Consider the following example: Let $C = \{c\}$, $F = \{f_1, f_2\}$ and let valuations be according to the following tables for some $\epsilon > 0$.

$v_c(f)$	f_1	f_2	$v_f(c)$	f_1	f_2
c_1	1	$1 - \epsilon$	c_1	1	1

Assume that $p \geq \max\{\kappa_C, \kappa_F\}$, and ϵ and δ_C are sufficiently small. If $p(1 - \lambda p) \geq \kappa_F$ then the matching correspondences of the unique FSE and the unique CSE are identical, and f_2 only incurs search costs in CS if the match between f_1 and c_1 is not suitable. Thus, an agent can be made strictly better off in CS compared to FS without making anyone else worse off by saving wasted search efforts. If $p(1 - \lambda p) < \kappa_F$, then in the unique FSE there is only mutual interest between c and f_1 while in the unique CSE, all agents have again mutual interest in each other. Furthermore, the CSE is a Pareto improvement over the FSE.

Let $s^{FS} \in S^{FS}$ and $s^{CS} \in S^{CS}$. We now show that if there exists an agent $i \in A$, such that $u_i(s^{FS}) > u_i(s^{CS})$, then there exists another agent $j \in A$ with $u_j(s^{FS}) < u_j(s^{CS})$. Suppose $i \in C$. If $u_i(s^{FS}) > u_i(s^{CS})$, we get by Lemma 2 that $M_c(s^{FS}) \not\subseteq M_c(s^{CS})$. Then, by Lemma 1, the claim immediately follows. Now suppose $i \in F$. Further, for the sake of contradiction, assume $M_c(s^{FS}) \subseteq M_c(s^{CS})$ for all $c \in C$, as otherwise by Lemma 1 the claim would immediately follow. Therefore, f 's increase in $u_f(s^{FS})$ can only come from the fact that there exists a child $c \in M_f(s^{FS})$ and another family $f' \in F$ with $v_c(f') > v_f(f)$ that is not interested in c under s^{FS} but under s^{CS} . Suppose c responds to s_{-c}^{CS} with s_c^{FS} in CS. By Lemma 2, this would imply $u_c(s^{FS}) \leq u_c((s_c^{FS}, s_{-c}^{CS}))$. However, note that c 's utility would strictly increase if c would be interested in f' instead of f since $v_c(f') > v_c(f)$. Hence, c does not play a best response in s^{CS} , a contradiction.

D.7. Proof of Proposition 4

Proof. Consider the following example: Let $C = \{c_1, c_2\}$, $F = \{f_1, f_2\}$ and let valuations be according to the following tables.

$v_c(f)$	f_1	f_2
c_1	L	ϵ
c_2	ϵ	L

$v_f(c)$	f_1	f_2
c_1	ϵ	L
c_2	L	ϵ

If agents are patient enough (i.e., δ_C and δ_F are close enough to 1) and search costs are small enough, there exist only two equilibria, independent of search technology. In the child-optimal equilibrium, s^{co} , children are only interested in their preferred choice (i.e., the agent type for which they have a match value of L) and in the family-optimal equilibrium s^{fo} families are only interested in their preferred choice. Note that if each child is mutually interested in at most one family in strategy profile s , then $u_i^{CS}(s) = u_i^{FS}(s)$ for all $i \in A$. As $p, \delta_F, \delta_C \rightarrow 1$, children's utilities will be weakly less than ϵ under s^{fo} while being positive and converging to L under s^{co} . Similarly, families' utilities will be weakly less than ϵ under s^{co} while being positive and converging to L under s^{fo} .

D.8. Proof of Proposition 5

Proof. Consider the following example: Let $C = \{c_1, c_2\}$, $F = \{f_1, f_2\}$ and let valuations be according to the following tables for some $\epsilon > 0$.

$v_c(f)$	f_1	f_2
c_1	1	$1 - \epsilon$
c_2	1	$1 - \epsilon$

$v_f(c)$	f_1	f_2
c_1	1	1
c_2	0	κ_F/p

If $\epsilon > 0$ is small enough and $p(1 - \lambda p) < \kappa_F$, then in the unique FSE s^{FS} , only c_1 and f_1 will be mutually interested in each other and c_2 and f_2 . Family f_2 will not be interested in c_1 in FS, independent of how patient f_2 is. That is, because the probability of actually matching with c_1 while facing competition from c_2 is too small, yet f_2 would have to incur search costs every time c_1 is active. However, if δ_F is sufficiently large, then the unique CSE matching correspondence is $M(s^{CS}) = \{(c_1, f_1), (c_1, f_2)\}$. In CS, c_2 , the “low-type” child, will remain unmatched.

D.9. Proof of Proposition 6

Proof. Consider the following example: Let $C = \{c\}$, $F = \{f_1, f_2, f_3\}$ and let valuations be according to the following tables for some $\epsilon > 0$.

$v_c(f)$	f_1	f_2	f_3
c	1	$1 - \epsilon$	$1 - 2\epsilon$

$v_f(c)$	f_1	f_2	f_3
c	1	$\kappa_F/p + \epsilon$	1

Choose $\delta_C = \delta_F = 0$. If both ϵ and κ_C are sufficiently small and further $\kappa_F < p$ we get the following: In the unique FSE s^{FS} , child c will be mutually interested in both f_1 and f_3 . In the unique CSE, all agents will have mutual interest in each other, and therefore f_3 's utility strictly decreases.

D.10. Proof of Proposition 7

Proof. Suppose all instance parameters except for δ_F are fixed. Let \mathcal{M} be the set of all child-family pairs (c, f) for which there exists $\delta_F \in [0, 1]$, such that c and f are mutually interested in each other under some FSE s^{FS} . Further, let v_{\min} denote the smallest value a family in \mathcal{M} has for a mutually interested child, i.e., $v_{\min} = \min_{(c, f) \in \mathcal{M}} v_f(c)$. For δ_F , such that

$$\delta_F \leq \frac{v_{\min} - \kappa_F/p}{\bar{v}}, \quad (42)$$

it follows that $p(v_f(c) - \delta_F \bar{v}) \geq \kappa_F$ for all $(c, f) \in \mathcal{M}$. Since \bar{v} is an upper bound for agents' utilities, for any such δ_F and any $(c, f) \in \mathcal{M}$, being interested in c is (weakly) advantageous for f under CS, independent of all other strategies. As in the proof of Lemma 2, this implies that c must be weakly better off in CS compared to FS.

D.11. Proof of Proposition 8

Proof. Consider the following example: Let $C = \{c_1, c_2\}$, $F = \{f_1, f_2\}$ and let valuations be according to the following tables for some $\epsilon > 0$.

$v_c(f)$	f_1	f_2	$v_f(c)$	f_1	f_2
c	1	$1 - \epsilon$	c	1	1

Choose p and κ_F , such that $p(1 - \lambda'p) < \kappa_F \leq p(1 - \lambda p)$. If κ_C , δ_C , and δ_F are sufficiently small, child c will be mutually interested in f_1 and f_2 in the unique FSE for λ . But in the unique FSE for λ' , only c and f_1 will be mutually interested in each other. However, if the difference between λ and λ' is very small, c 's utility can be smaller in the case where the market thickness indicator is λ , as the increased probability of f_1 being active might not compensate for the loss of f_2 's interest.

D.12. Proof of Proposition 9

Let $u = (u_i(s^{co-CS, \lambda}))_{i \in A}$ denote the vector of agents' utilities under $s^{co-CS, \lambda}$ and let $Y' = [u_c, \bar{v}]^n \times [0, u_f]^m \subseteq [0, \bar{v}]^{n+m} = Y$. We now define a mapping $T^\lambda : Y \rightarrow Y$ as follows: $T^\lambda = (T_i^\lambda)_{i \in A}$, where

$$T_c^\lambda(y) = \delta_C y_c + \lambda \sum_{f \in F} \mathbb{1}[p(v_f(c) - \delta_F y_f) \geq \kappa_F] \beta_{cf}(y) \left(p(v_c(f) - \delta_C y_c) - \kappa_C \right)^+ \quad (43)$$

for all $c \in C$ and

$$T_f^\lambda(y) = \delta_F y_f + \frac{1}{n} \sum_{c \in C} \mathbb{1}[p(v_c(f) - \delta_C y_c) \geq \kappa_C] \beta_{cf}(y) \left(p(v_f(c) - \delta_F y_f) - \kappa_F \right)^+. \quad (44)$$

As we have seen in a previous proof, T^λ is \leq_C -monotone on Y . We now show that $T^{\lambda'}$ maps from Y' to Y' , which by Tarski's fixed-point theorem yields the result. Since T^λ is \leq_C -monotone and u is the \leq_C -minimal element of Y' , it is sufficient to show that $T^{\lambda'}(u) \in Y'$.

Because u is an equilibrium threshold profile, we have that $T^\lambda(u) = u \in Y'$. Further, it can easily be verified that $T^\lambda(y) \leq_C T^{\lambda'}(y)$ for all $y \in Y$. Hence, $T^{\lambda'}(u) \in Y'$, which completes the proof.

E. Algorithm to Compute FS-TSs from Threshold Profiles

Algorithm 1 can be used to compute $s^{FS}(y)$ for a given threshold profile $y \in \mathbb{R}^{n+m}$. Since $\beta_{cf}(s)$ only depends on families $f' \in F$ with $v_c(f') > v_c(f)$ and for each child families are processed in decreasing order of $v_c(f)$, the final strategy profile satisfies the equations for FS from Proposition 1.

Algorithm 1: Thresholds to strategy profile

Input: $y \in \mathbb{R}^{n+m}$

Output: Strategy profile $s \in S$

$s_c(f) := 0$ and $s_f(c) := 0$ for all $c \in C$, $f \in F$

for $c \in C$ **do**

$U := F$

while $U \neq \emptyset$ **do**

$f := \arg \max_{f' \in U} v_c(f')$

if $\beta_{cf}(s)p(v_f(c) - \delta_F y_f) \geq \kappa_F$ **then**
 $s_f(c) := 1$

end

if $\beta_{cf}(s)p(v_c(f) - \delta_C y_c) \geq \kappa_C$ **then**
 $s_c(f) := 1$

end

$U := U \setminus \{f\}$

end

end

F. Limit Results

Here, we provide a collection of limit results that all illustrate how the differences between FS and CS disappear as certain parameters take on extreme values.

F.1. Negligible Search Costs

The next proposition shows that the games induced by CS and FS become identical as search costs become negligible.

PROPOSITION 13. *As $\kappa_C \rightarrow 0$ and $\kappa_F \rightarrow 0$, $|u_i^{FS}(s) - u_i^{CS}(s)| \rightarrow 0$ for all $s \in S$, $i \in A$.*

Proof. Assume that $\kappa_C = \kappa_F =: \kappa$ and let $s \in S$. For each family f , define $T_{s,f}^{FS} : Y \rightarrow \mathbb{R}$ and $T_{s,f}^{CS} : Y \rightarrow \mathbb{R}$ as follows:

$$T_{s,f}^{FS}(y) = \delta_F y_f + \frac{1}{n} \sum_{c \in C} s_c(f) s_f(c) \left(\beta_{cf}(s) p(v_f(c) - \delta_F y_f) - \kappa \right) \quad (45)$$

and

$$T_{s,f}^{CS}(y) = \delta_F y_f + \frac{1}{n} \sum_{c \in C} s_c(f) s_f(c) \beta_{cf}(s) \left(p(v_f(c) - \delta_F y_f) - \kappa \right). \quad (46)$$

Similarly, for all $c \in C$, let

$$T_{s,c}^{FS}(y) = \delta_C y_c + \lambda \sum_{f \in F} s_c(f) s_f(c) \left(\beta_{cf}(s) p(v_c(f) - \delta_C y_c) - \kappa \right) \quad (47)$$

and

$$T_{s,c}^{CS}(y) = \delta_C y_c + \lambda \sum_{f \in F} s_c(f) s_f(c) \beta_{cf}(s) \left(p(v_c(f) - \delta_C y_c) - \kappa \right). \quad (48)$$

Notice that the unique fixpoints $y^{FS}, y^{CS} \in Y$ of $T_s^{FS}(y)$ and $T_s^{CS}(y)$ define agents' utilities under s in FS and CS, respectively. For all $y \in Y$ and $f \in F$ it holds that

$$\lim_{\kappa \rightarrow 0} \left(\delta_F y_f + \frac{1}{n} \sum_{c \in C} s_c(f) s_f(c) \left(\beta_{cf}(s) p(v_f(c) - \delta_F y_f) - \kappa \right) \right) \quad (49)$$

$$= \delta_F y_f + \frac{1}{n} \sum_{c \in C} s_c(f) s_f(c) \left(\beta_{cf}(s) p(v_f(c) - \delta_F y_f) \right) \quad (50)$$

$$= \lim_{\kappa \rightarrow 0} \left(\delta_F y_f + \frac{1}{n} \sum_{c \in C} s_c(f) s_f(c) \beta_{cf}(s) \left(p(v_f(c) - \delta_F y_f) - \kappa \right) \right). \quad (51)$$

Similarly, for all $y \in Y$ and $c \in C$ we have that

$$\lim_{\kappa \rightarrow 0} \left(\delta_C y_c + \lambda \sum_{f \in F} s_c(f) s_f(c) \left(\beta_{cf}(s) p(v_c(f) - \delta_C y_c) - \kappa \right) \right) \quad (52)$$

$$= \delta_C y_c + \lambda \sum_{f \in F} s_c(f) s_f(c) \left(\beta_{cf}(s) p(v_c(f) - \delta_C y_c) \right) \quad (53)$$

$$= \lim_{\kappa \rightarrow 0} \left(\delta_C y_c + \lambda \sum_{f \in F} s_c(f) s_f(c) \beta_{cf}(s) \left(p(v_c(f) - \delta_C y_c) - \kappa \right) \right), \quad (54)$$

and hence $\lim_{\kappa \rightarrow 0} |u_i^{FS}(s) - u_i^{CS}(s)| = 0$ for all $i \in A$.

FS and CS do not necessarily become identical if only one side has negligible search costs. In both cases—i.e., if only $\kappa_C \rightarrow 0$ or only $\kappa_F \rightarrow 0$ —we can create instances where the sets of equilibria differ from each other in the two approaches.

F.2. High Match Success Probability and Market Thickness

Match success probability and market thickness indicator are strongly connected in our model, such that we cannot make any insightful statements about the limit behavior if only one of them approaches 1. However, if it is certain that a family of each type will be present at each time step and that each family would be a suitable match, we observe once more that FS and CS become equivalent in a slightly different way.

PROPOSITION 14. As $\lambda p \rightarrow 1$, $|u_i^{FS}(s^{FS}(y)) - u_i^{CS}(s^{CS}(y))| \rightarrow 0$ for all $y \in Y$, $i \in A$.

Proof. For each family f , define $T_f^{FS} : Y \rightarrow \mathbb{R}$ and $T_f^{CS} : Y \rightarrow \mathbb{R}$ as follows:

$$T_f^{FS}(y) = \delta_F y_f + \frac{1}{n} \sum_{c \in C} \mathbb{1}[\beta_{cf}^{FS}(y)p(v_c(f) - \delta_C y_c) \geq \kappa_C] \left(\beta_{cf}^{FS}(y)p(v_f(c) - \delta_F y_f) - \kappa_F \right)^+ \quad (55)$$

and

$$T_f^{CS}(y) = \delta_F y_f + \frac{1}{n} \sum_{c \in C} \mathbb{1}[p(v_c(f) - \delta_C y_c) \geq \kappa_C] \beta_{cf}^{CS}(y) \left(p(v_f(c) - \delta_F y_f) - \kappa_F \right)^+. \quad (56)$$

Similarly, for all $c \in C$, let

$$T_c^{FS}(y) = \delta_C y_c + \lambda \sum_{f \in F} \mathbb{1}[\beta_{cf}^{FS}(y)p(v_f(c) - \delta_F y_f) \geq \kappa_F] \left(\beta_{cf}^{FS}(y)p(v_c(f) - \delta_C y_c) - \kappa_C \right)^+ \quad (57)$$

and

$$T_c^{CS}(y) = \delta_C y_c + \lambda \sum_{f \in F} \mathbb{1}[p(v_f(c) - \delta_F y_f) \geq \kappa_F] \beta_{cf}^{CS}(y) \left(p(v_c(f) - \delta_C y_c) - \kappa_C \right)^+. \quad (58)$$

Notice that the unique fixpoints $y^{FS}, y^{CS} \in Y$ of $T^{FS}(y)$ and $T^{CS}(y)$ define agents' utilities under an FS-TS and CS-TS with threshold y in FS and CS, respectively. For all $y \in Y$ and $f \in F$, it holds that

$$\lim_{\lambda p \rightarrow 1} \left(\delta_F y_f + \frac{1}{n} \sum_{c \in C} \mathbb{1}[\beta_{cf}^{FS}(y)p(v_c(f) - \delta_C y_c) \geq \kappa_C] \left(\beta_{cf}^{FS}(y)p(v_f(c) - \delta_F y_f) - \kappa_F \right)^+ \right) \quad (59)$$

$$= \delta_F y_f + \frac{1}{n} \sum_{c \in C} \mathbb{1} \left[f = \operatorname{argmax}_{f' \in F: v_{f'}(c) - \delta_F y_{f'} > \kappa_F \wedge v_c(f') - \delta_C y_c > \kappa_C} v_c(f') \right] \left(v_f(c) - \delta_F y_f - \kappa_F \right) \quad (60)$$

$$= \lim_{\lambda p \rightarrow 1} \left(\delta_F y_f + \frac{1}{n} \sum_{c \in C} \mathbb{1}[\beta_{cf}^{CS}(y)p(v_c(f) - \delta_C y_c) \geq \kappa_C] \beta_{cf}^{CS}(y) \left(p(v_f(c) - \delta_F y_f) - \kappa_F \right)^+ \right). \quad (61)$$

Similarly, for all $y \in Y$ and $c \in C$ we have that

$$\lim_{\lambda p \rightarrow 1} \left(\delta_C y_c + \lambda \sum_{f \in F} \mathbb{1}[\beta_{cf}^{FS}(y)p(v_f(c) - \delta_F y_f) \geq \kappa_F] \left(\beta_{cf}^{FS}(y)p(v_c(f) - \delta_C y_c) - \kappa_C \right)^+ \right) \quad (62)$$

$$= \delta_C y_c + \lambda \sum_{f \in F} \mathbb{1} \left[f = \operatorname{argmax}_{f' \in F: v_{f'}(c) - \delta_F y_{f'} > \kappa_F \wedge v_c(f') - \delta_C y_c > \kappa_C} v_c(f') \right] \left(v_c(f) - \delta_C y_c - \kappa_C \right) \quad (63)$$

$$= \lim_{\lambda p \rightarrow 1} \left(\delta_C y_c + \lambda \sum_{f \in F} \mathbb{1}[p(v_f(c) - \delta_F y_f) \geq \kappa_F] \beta_{cf}^{CS}(y) \left(p(v_c(f) - \delta_C y_c) - \kappa_C \right)^+ \right). \quad (64)$$

and hence $\lim_{\lambda p \rightarrow 1} |u_i(s_i^{FS}(y)) - u_i(s_i^{CS}(y))| = 0$. This concludes the proof for families. For children, the proof is analogous and omitted.

In order to see why Proposition 14 holds, notice that the probability of child c matching with his first choice from $M_c(s)$ goes to 1 as $\lambda p \rightarrow 1$. Thus, the contribution of all other families from $M_c(s)$ goes to zero. As only these first choices contribute to utilities, the difference between FS and CS again disappears.

F.3. Low Market Thickness

As the supply on the family side becomes very small, the differences between CS and FS disappear in equilibrium.

PROPOSITION 15. *As $\lambda \rightarrow 0$, $|u_i^{FS}(s) - u_i^{CS}(s)| \rightarrow 0$ for all $s \in S$, $i \in A$.*

Proof. Notice that the unique fixpoints $y^{FS}, y^{CS} \in Y$ of $T_s^{FS}(y)$ and $T_s^{CS}(y)$ (recall definitions from the proof of Proposition 13 define agents' utilities under s in FS and CS, respectively. For all $y \in Y$ and $f \in F$ it holds that

$$\lim_{\lambda \rightarrow 0} \left(\delta_F y_f + \frac{1}{n} \sum_{c \in C} s_c(f) s_f(c) \left(\beta_{cf}(s) p(v_f(c) - \delta_F y_f) - \kappa_F \right) \right) \quad (65)$$

$$= \delta_F y_f + \frac{1}{n} \sum_{c \in C} s_c(f) s_f(c) \left(p(v_f(c) - \delta_F y_f) \right) \quad (66)$$

$$= \lim_{\lambda \rightarrow 0} \left(\delta_F y_f + \frac{1}{n} \sum_{c \in C} s_c(f) s_f(c) \beta_{cf}(s) \left(p(v_f(c) - \delta_F y_f) - \kappa_F \right) \right), \quad (67)$$

because $\beta_{cf}(s) \rightarrow 1$ as $\lambda \rightarrow 0$. Similarly, for all $y \in Y$ and $c \in C$ we have that

$$\lim_{\lambda \rightarrow 0} \left(\delta_C y_c + \lambda \sum_{f \in F} s_c(f) s_f(c) \left(\beta_{cf}(s) p(v_c(f) - \delta_C y_c) - \kappa_C \right) \right) \quad (68)$$

$$= \delta_C y_c \quad (69)$$

$$= \lim_{\lambda \rightarrow 0} \left(\delta_C y_c + \lambda \sum_{f \in F} s_c(f) s_f(c) \beta_{cf}(s) \left(p(v_c(f) - \delta_C y_c) - \kappa_C \right) \right), \quad (70)$$

and hence $\lim_{\lambda \rightarrow 0} |u_i^{FS}(s) - u_i^{CS}(s)| = 0$ for all $i \in A$.

The reason why the statement holds is that for very small λ , families do not have to worry about competition in FS, as it is likely that there is no other family active at any given time step.

F.4. Patient Agents

Here, we first need to revisit classical matching markets. The matching market *induced* by $(v, \delta_C, \delta_F, \kappa_C, \kappa_F, p, \lambda)$ is a tuple (C, F, \succ) , where $f \succ_c f'$ if and only if $v_c(f) > v_c(f')$ and $c \succ_f c'$ if and only if $v_f(c) > v_f(c')$. Further, $f \succ_c c$ if and only if $p v_c(f) \geq \kappa_C$ and $c \succ_f f$ if and only if $p v_f(c) \geq \kappa_F$. For the remainder of this section, we assume that $\delta_C = \delta_F =: \delta$.

PROPOSITION 16. *There exists $\bar{\delta} \in [0, 1)$, such that for all $\delta \in [\bar{\delta}, 1)$ the set of equilibrium matchings correspondences are identical in FS and CS and coincide with the set of stable matchings in the induced marriage market.*

For small p and large search costs, the statement above is trivial: If search costs become too large or the probability of a match being suitable becomes too small, no agent will have an incentive to

be interested in any potential match candidate in either FS or CS. Hence, no matches will form in any equilibrium.

We only prove the statement for iii) here, as i) and ii) are trivial. For this proof, we first need to revisit classical matching markets. The matching market *induced* by $(v, \delta_C, \delta_F, \kappa_C, \kappa_F, p, \lambda)$ is a tuple (C, F, \succ) , where $f \succ_c f'$ if and only if $v_c(f) > v_c(f')$ and $c \succ_f c'$ if and only if $v_f(c) > v_f(c')$. Further, $f \succ_c c$ if and only if $pv_c(f) \geq \kappa_C$ and $c \succ_f f$ if and only if $pv_f(c) \geq \kappa_F$.

Let (C, F, \succ) be a matching market where agents have strict preferences.

DEFINITION 5. A pair of functions $g = (g_C, g_F)$ is called a *pre-matching* if $g_C : C \rightarrow A$ and $g_F : F \rightarrow A$, such that if $g_C(c) \neq c$ then $g_C(c) \in F$ and if $g_F(f) \neq f$ then $g_F(f) \in C$.

We say that a pre-matching g *induces* a matching w if the function $w : A \rightarrow A$ defined by $w(i) = g(i)$ is a matching. Consider the following set of equations:

$$g_C(c) = \max_{\succ_c} (\{f \in F \mid c \succeq_f g_F(f)\} \cup \{c\}), \quad c \in C, \quad (71)$$

$$g_F(f) = \max_{\succ_f} (\{c \in C \mid f \succeq_c g_C(c)\} \cup \{f\}), \quad f \in F, \quad (72)$$

where the maxima are taken with respect to agents' preferences.

LEMMA 3 (**Adachi (2003)**). *If a matching w is stable, then the pre-matching g defined by w solves the above equations. If a pre-matching g solves the above equations, then g induces a stable matching w .*

With the above lemma, we can now prove Lemma 4. Assume that $\delta_C = \delta_F =: \delta$.

LEMMA 4. *There exists $\bar{\delta} \in [0, 1)$, such that for all $\delta \in [\bar{\delta}, 1)$ the set of equilibrium matching correspondences are identical in FS and CS and coincide with the set of stable matchings in the induced marriage market.*

Proof. Let $s \in S^{FS}$ and $c \in C$. c 's utility under s is the unique value $u_c^{FS}(s)$ that satisfies

$$u_c(s) = \delta u_c(s) + \lambda \sum_{f \in F} s_f(c) \left(\beta_{cf}(s) p(v_f(c) - \delta u_c(s)) - \kappa_C \right)^+ \quad (73)$$

$$\iff (1 - \delta)u_c(s) = \lambda \sum_{f \in F} s_f(c) \left(\beta_{cf}(s) p(v_f(c) - \delta u_c(s)) - \kappa_C \right)^+. \quad (74)$$

Therefore, as $\delta \rightarrow 1$ we get

$$\frac{1}{n} \sum_{f \in F} s_f(c) \left(\beta_{cf}(s) p(v_f(c) - \delta u_c^{FS}(s)) - \kappa_C \right)^+ \rightarrow 0. \quad (75)$$

For the sake of contradiction, assume that $|M_c(s)| \geq 2$. Let $f^* = \arg \max_{f \in F: s_{f^*}(c)=1} v_c(f)$ and $f \in M_c(s) \setminus \{f^*\}$. Because of Equation (75), it must hold that

$$\beta_{cf}(s) \left(v_c(f) - \delta u_c(s) - \kappa_C/p \right)^+ \rightarrow 0. \quad (76)$$

However, since $v_c(f^*) > v_c(f)$ and $\beta_{cf^*}(s) > \beta_{cf}(s)$ we get that

$$\beta_{cf^*}(s) \left(v_c(f^*) - \delta u_c(s) - \kappa_C/p \right)^+ \rightarrow \epsilon \quad (77)$$

for some $\epsilon > 0$, a contradiction. Hence, $M_c(s) = \{f^*\}$ if $pv_c(f^*) \geq \kappa_C$ or $M_c(s) = \emptyset$ otherwise. Notice that this is equivalent to the expression of Equation (71). Now that we have established that children will be mutually interested in at most one family, the same can similarly be shown for families, which completes the proof for FS. The proof for CS is analogous and thus omitted.

G. Family Preferences

Using data on 1,364 families across the state of Florida who completed a multi-step registration and approval process with the platform examined in Section 9, we provide high-level insights on preference alignment to help understand which parameter regions discussed in Section 8 best represent reality. It should be noted that, as a snapshot of the population registered on the platform, this may not necessarily be representative of all families looking to adopt in Florida. As these preferences are based on families' answers to a registration questionnaire, they are not binding on how families respond when they are solicited for their interest in specific children, nor do they indicate the family's suitability to care for such a child.




Figure 6 Acceptable child ages for 1,364 families in Florida active on the platform.

First, as shown in Figure 6, age preferences vary significantly. While a mean lower age limit of 1.8 combined with a standard deviation of 3.0 indicates a high willingness among families to adopt younger children, a mean upper age limit of 10.5 with a standard deviation of 4.8 indicates a reluctance to accept older children. The age that families find most acceptable is five years old, which is within the minimum and maximum age range for 75% of families.

Concerning the child's gender, almost 70% of families chose "no preference," with 12.7% preferring boys and 15.4% preferring girls. Although not pronounced, this difference in gender preference

presents a challenge due to the higher number of boys that become available for adoption (see Table 1 in Section 9) and partially explains why male children have worse adoption hazard rate coefficients. We note that — while outside the scope of what our model considers — choosing “no preference” may also indicate an interest in sibling groups, and the mean maximum number of children that a family expresses interest in is 2.2, with a standard deviation of 1.2.

Ethnicity and race offer additional attributes over which families express heterogeneous preferences. While 70% of families expressed an openness to a child of any ethnicity, 21% expressed an interest in white children, 16% in biracial children, and between 7% and 14% in each of four other racial categories. Half of families are interested in children of Hispanic or Latino ethnicity, which is indicated separately from race.




Figure 7 Preferences over child attributes for 1,364 families in Florida active on the platform.

Family preferences over seven additional attributes shown in Figure 7 further diminish the plausibility of highly aligned preferences among families on some dimensions. For example, while nearly all families are open to a child impacted by neglect or substance abuse, only 41% of families would consider or definitely be interested in a child with a criminal record.

H. Numerical Evaluation: Supplementary Material

H.1. Families




Figure 8 The ratio of families that are on average (strictly) better off in either approach in the family-optimal equilibrium for different combinations of agents’ patience and the level of preference correlation.

From Figure 8, we can see that the majority of families achieves a higher utility under CS, and almost no families achieve a higher utility under FS.

H.2. Match Probabilities

Figure 9 shows that, on average, children are more likely to get matched in CS than FS at any given time step.




Figure 9 Children's average match probability (averaged over all children and over all instances) in the family-optimal equilibrium for different combinations of agents' patience and the level of preference correlation.

I. Empirical Analysis: Supplementary Material

I.1. Data

Our main analysis relies on multiple data sets: the AFCARS Foster Care 6-month File (Children's Bureau, Administration on Children, Youth and Families 2023b), the AFCARS Adoption File (Children's Bureau, Administration on Children, Youth and Families 2023a), and case history data from the platform. From 766,527 AFCARS foster care 6-month update records for Florida children, we identified 10,286 children as legally free and clear for adoption with cases starting after October 1, 2014, which is the closest federal fiscal year cut-off for the data in Figure 4. Regrettably, and despite significant efforts, the agency implementing the platform could not extract information from Florida's statewide case management system and assemble its own comparable case history dataset. However, we provide validation in online Appendix I.3 that uses circuit-level data in state reports to verify that the platform had roughly average performance for all adoptions compared to other circuits in the state. Filtering children based on case goals and the relationships with adoptive families proved to be the biggest obstacles to using the agency's data. However, the agency was helpful in manually tracking down the outcomes of children who left the platform without an adoptive placement. In Figure 10, we provide an outline of how we combined different data sources to create the dataset used and how it is used in various analyses.

For our analysis, children's timelines begin with the termination of parental rights (TPR) order and end with the last status update, which could be the adoption finalization date. For AFCARS

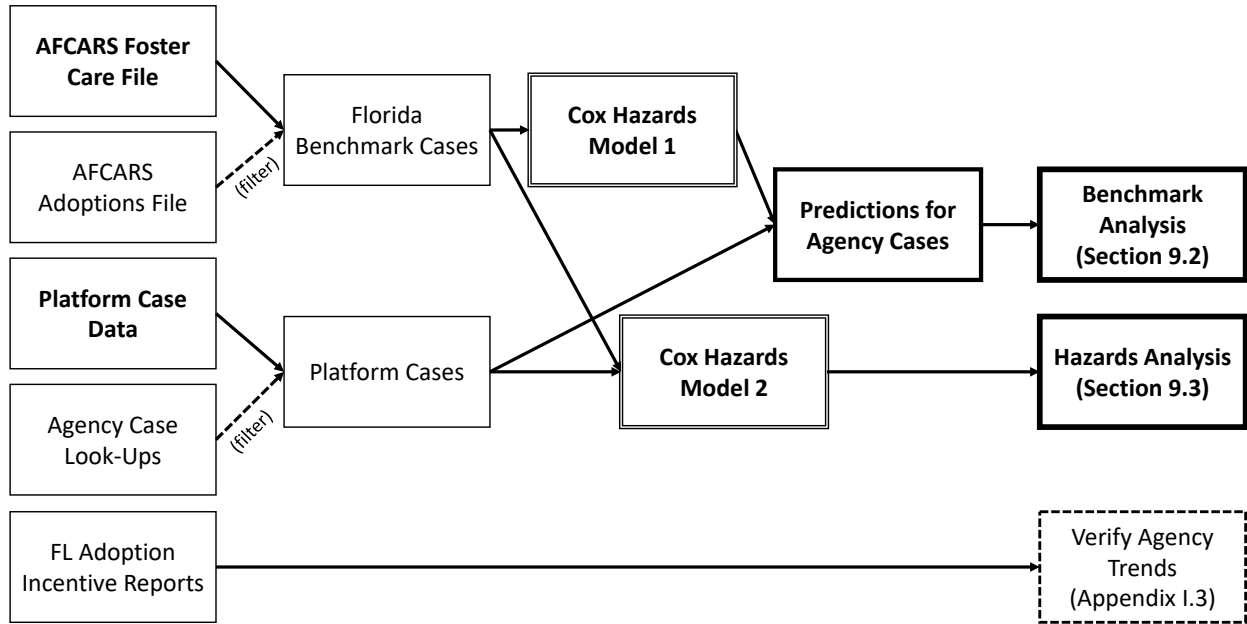


Figure 10 Data sources used in the analyses of Section 9 and online Appendix I.3.

cases to qualify for this analysis, children had to have a case goal of adoption in their final record or previously had a case goal of adoption with (a) a resulting non-relative adoption, (b) a final case goal of “emancipation”, or (c) a discharge reason of “emancipation”. We excluded children listed in the Adoption File as being adopted by a relative, step-parent, or foster parent. Because the most recent AFCARS Adoptions File only covers adoptions through September 30, 2021, we excluded case data past that date as we would not be able to tell if adoptions resulted from relative, step-parent, or foster placements. We also excluded any cases with a duration less than 120 days — a 30-day appeal window after TPR before a search can start, plus a legal minimum of 90 days between placement with the adoptive family and adoption finalization — or over 18 years.

We note that the AFCARS datasets’ inability to explicitly identify children in need of adoptive resources results in a conservative benchmark on the platform’s performance; i.e., it will make the platform appear less helpful than it actually was. Specifically, some children included in the AFCARS dataset may quickly find placements with non-relatives, such as teachers, church members, or neighbors. Such children have short times until adoption, but would not have been listed on the platform because they had an identified adoptive placement at TPR. While we filter out cases with implausibly short durations of 120 days or fewer, some of these cases with durations exceeding 120 days may still be included in the AFCARS benchmark model, as we cannot identify them separately. Child welfare professionals have also mentioned the possibility that users of the AFCARS reporting system may inadvertently classify adoptions by foster parents as non-relative adoptions. These potential limitations show the importance of more explicitly identifying children




Figure 11 Actual adoptions by the agency using the platform and other channels compared to Florida AFCARS benchmark model assuming the platform case creation date as the start of the hazards model time horizon.

for whom an active search is being conducted in child welfare data sources to aid in evaluating performance.

We received case data from the platform about 279 children in need of adoptive placements who were listed by the agency before the AFCARS cut-off date of October 1, 2021. The platform provided its first matches around July 1, 2018. Because the platform allows caseworkers to register children and indicate the time since TPR using time intervals, the TPR date is estimated using the average of the interval endpoints and subtracting that value from the platform case creation date. Of the children listed on the platform, 165 had finalized adoptions by February 1, 2023. Of the finalized adoptions, the adoptive parents for 144 children were found through the platform, and 21 were actively listed on the platform when a match was found through some other channel, such as serendipitous encounters in the agency’s office. We excluded an additional 66 children who were listed at one time on the platform but achieved permanency through adoptive placement with relatives or foster care parents or were reunified. Our data only includes activity and case updates through February 1, 2023; children adopted after that date are not counted as adopted in our dataset.

I.2. Robustness Checks

We now provide three additional model variations for our empirical analysis as a robustness check.

I.2.1. Analysis without Conditional Adoption Probabilities or Time-varying Platform Effects We observe similar results from this comparison to the AFCARS benchmark, which is shown in Figure 11. For children listed on the platform before October 1, 2021, the predicted number of finalized adoptions within two years was 121.4. However, 138 adoptions were finalized

within two years. Within three years of listing, this difference extends to an extra 27.3 adopted children, or a 20% increase over the benchmark.

Similarly, we present an alternate model to Model 2, where we assume children’s timelines start when the respective adoptive search starts: at TPR for the AFCARS children and at case creation for the platform children. This means we can treat the platform indicator as static and employ a standard, non-time-varying Cox hazards model. Formally, we use a constant binary covariate over time, $OnPlatform_i$, for each child i . This covariate is 1 for any child who has been listed on the platform at any point in time before adoption and time $t = 0$ for this child is when the child’s case is created on the platform.

Table 3 Cox Proportional Hazards Model of Time Until Adoption:
Model 1 from the Main Text and New Model with Unconditional Adoption Probabilities and Time-constant

	Platform Effect	
	Model 1	Model 3
Female	1.080** (2.615)	1.089** (2.955)
Black	0.768*** (−8.379)	0.764*** (−8.659)
Hispanic	0.827*** (−4.313)	0.830*** (−4.243)
Age at TPR (years)	0.889*** (−9.563)	0.892*** (−9.459)
(Age at TPR) ²	0.999 (−0.791)	0.999 (−1.077)
Disability	0.881*** (−3.766)	0.870*** (−4.198)
TPR in FY2016	0.967 (−0.605)	0.967 (−0.602)
TPR in FY2017	1.052 (0.918)	1.040 (0.714)
TPR in FY2018	0.927 (−1.419)	0.911 (−1.762)
TPR in FY2019	0.753*** (−5.191)	0.763*** (−5.010)
TPR in FY2020	0.516*** (−10.847)	0.527*** (−10.648)
TPR in FY2021	0.493*** (−8.613)	0.511*** (−8.429)
Platform		1.175* (2.005)
Platform Cases	0	279
N	10,286	10,565
Concordance	0.680	0.680
Log-likelihood ratio test	1969.148 on 12 d.f.	2023.878 on 13 d.f.

Note: exponentiated coefficients; z-statistics in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

We call this Model 3, with resulting coefficients found in Table 3. Control variables still behave as expected, with Girls experiencing higher adoption hazards than boys, while Black and Hispanic




Figure 12 Actual adoptions by the agency using the platform and other channels compared to Florida AFCARS benchmark model, excluding children adopted via channels other than the platform.

children tend to be adopted more slowly than children of other racial groups. A clinical disability diagnosis and a child’s increasing age are also still linked to decreased adoption hazards, as well as TPR in pandemic years or immediately before. The platform is also still consistently associated with faster adoption, with an estimated platform hazard ratio of 1.175 and statistical significance at the 5% level.

I.2.2. Excluding Off-platform Adoptions We also present results in Figure 12 that exclude the children who were adopted by families found through channels other than the platform. With these 21 adoptions removed, the platform continues to outperform the state benchmark. The number of platform adoptions by the two-year and three-year marks is 12% and 15%, respectively, over the statewide benchmark for adoptions. The number of adoptions achieved exceeds the confidence interval from the Poisson binomial distribution for times until adoption of 23 months and longer.

I.2.3. A More Exclusive AFCARS Race Variable. We also consider an alternate version characterization of the *Black* variable, as the approach in Section 9 takes a conservative approach due to differences in how the AFCARS and platform data treat multi-racial children. The AFCARS dataset allows children to have multiple race categories indicated, while the platform only allows one race to be selected and also has a category labeled “other” that presumably is used for some multi-racial children.

In this alternate approach, we treat the Black variable as referring to children identified *exclusively* as Black; i.e., only the Black race category has value 1 in the AFCARS data. Table 4 shows the updated hazards models. For all three models, the coefficient of the Black variable decreases, which indicates that multi-racial children experience a higher likelihood of adoption. Because the number of children designated as Black in the platform’s dataset does not change — i.e., unlike

Table 4 Cox Proportional Hazards Models of Time Until Adoption with Exclusively Black AFCARS Race

Variable	Model 1 ^b	Model 2 ^b	Model 3 ^b
Female	1.078** (2.555)	1.087** (2.902)	1.087** (2.893)
Black	0.755*** (−8.277)	0.752*** (−8.495)	0.751*** (−8.549)
Hispanic	0.830*** (−4.242)	0.842*** (−3.927)	0.833*** (−4.167)
Age at TPR (years)	0.889*** (−9.558)	0.891*** (−9.513)	0.892*** (−9.454)
(Age at TPR) ²	0.999 (−0.731)	0.999 (−1.002)	0.999 (−1.017)
Disability	0.882*** (−3.707)	0.868*** (−4.282)	0.872*** (−4.142)
TPR in FY2016	0.971 (−0.522)	0.959 (−0.753)	0.971 (−0.519)
TPR in FY2017	1.062 (1.088)	1.041 (0.730)	1.050 (0.883)
TPR in FY2018	0.931 (−1.350)	0.892* (−2.164)	0.914 (−1.694)
TPR in FY2019	0.757*** (−5.105)	0.754*** (−5.216)	0.766*** (−4.923)
TPR in FY2020	0.516*** (−10.850)	0.522*** (−10.821)	0.527*** (−10.648)
TPR in FY2021	0.495*** (−8.561)	0.505*** (−8.575)	0.513*** (−8.375)
Platform (Time-varying in Model 2 ^b)		1.298** (3.231)	1.200* (2.274)
Platform Cases	0	279	279
N	10,286	10,565	10,565
Concordance	0.680		0.680
Log-likelihood ratio test	1968.562 on 12 d.f.	2040.332 on 13 d.f.	2023.160 on 13 d.f.

Note: *exponentiated coefficients; z-statistics in parentheses.*

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

the analysis in the main paper, they are assumed to be exclusively Black due to the presence of the “other” category — the platform’s hazards coefficient increases slightly.

Figure 13 shows the corresponding benchmark comparison. Narrowing the Black definition lowers the predicted three-year adoptions for the platform cohort from 131.4 to 129.7, widening the observed-versus-expected gap and slightly strengthening the evidence that the caseworker-driven search platform accelerates placements.

I.3. Validation from State Reports

To validate our empirical analysis — especially to understand how the circuit that implemented the platform compares to statewide averages before and after implementation — we use the “Adoption Incentive” annual reports published by the Florida Department of Children and Families (2019, 2024). The analysis using AFCARS data implicitly assumes that the agency is representative of statewide patterns; if the agency already outperformed statewide averages — controlling for the




Figure 13 Actual adoptions by the agency using the platform and other channels compared to Florida AFCARS benchmark model when the Black variable refers children in the AFCARS data who only have a Black racial designation.

population of children — before implementing the platform, the analysis would overestimate the effect of the platform.

The Adoption Incentive reports provide annual statistics for how the 20 circuits in Florida perform on various measures. We note that Florida provides data on 19 entities; while most are individual circuits, some circuits are split or combined. Of the provided statistics, adoption success is best measured by the “number of children who were eligible for adoption on 7/1 who were adopted by 6/30,” which is displayed in Figure 4 in the main paper. In this case, eligibility refers to children for whom a termination of parental rights order has been granted. We refer to this metric as the *adoption clearance rate*. Using the case timelines of the children on the platform, approximately 40% of the children eligible for adoption every year on July 1 belonged to the set of 279 children who required search services through the platform. The remaining children likely already had a path to adoption identified through a foster parent or relative and would be expected to have a faster path to adoption.

Table 5 Mean annual adoption clearance rate for the circuit that implemented a caseworker-driven search platform compared to state averages. The clearance rate refers to the percentage of children eligible for adoption on July 1 of a year whose adoption is finalized by June 30 of the following year.

Comparison Period	Before	After
	7/1/2014-6/30/2018	7/1/2019-6/30/2024
Platform-Implementing Circuit Annual Mean % Adopted	57%	58%
Statewide Annual Mean % Adopted	55%	51%
Circuit-Statewide Ratio Mean	103%	115%
Mean Rank (of 19 circuits)	9.25	5.60

In Table 5, we compare the circuit’s average adoption case clearance rate performance against the statewide average for the four years before implementing the platform and the five years after implementation. Because the monthly case creation peaked after July 1, 2018, as the platform’s usage gradually ramped up into fall 2018, we disregard the annual report for July 1, 2018, to June 30, 2019, as a transition period. Thus, we compare the mean across annual statistics from July 1, 2014, until June 30, 2018, against July 1, 2019, until June 30, 2024. In the four years before implementation, the agency’s performance was only 3% higher than the statewide average, corresponding to an average ranking among all circuits of 9.25 out of 19. Thus, we expect the statewide AFCARS case data used for benchmarking to accurately reflect the agency’s performance without the platform.

Considering the averages over the five years since implementation, the circuit has seen a slight increase in its own performance and outperformed statewide averages. We note that the statewide average for the percentage of eligible children adopted decreased compared to before 2018, which could reflect higher acuity in children’s needs or increased difficulties in casework and judicial processes from the COVID-19 pandemic. Considering the five years of data after implementation, the mean ratio for the circuit’s case clearance rate compared to the statewide rate increased to 14.6%. While it is difficult to directly link the percentage of eligible children adopted over a one-year time frame to the outcomes explored in Section 9, this data indicates circuit has improved its performance in relation to the state as a whole and lends credence to the value of our previous analysis using benchmarks from statewide AFCARS data.