

The Pennsylvania Adoption Exchange Improves Its Matching Process

Vincent W. Slauch, Mustafa Akan, Onur Kesten
Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA 15213
vslauch@cmu.edu, akan@cmu.edu, okesten@andrew.cmu.edu

M. Utku Ünver
Department of Economics, Boston College, Chestnut Hill, MA, 02467
unver@bc.edu

The Pennsylvania Adoption Exchange (PAE) helps case workers representing children in state custody by recommending prospective families for adoption. We describe PAE's operational challenges using case worker surveys. We analyze child outcomes through a regression analysis of data on child outcomes over multiple years and improve the implementation of a family match ranking algorithm through a spreadsheet tool. Using a discrete-event simulation of PAE, we justify the value of a statewide adoption network and demonstrate the importance of the family preference information quality on the percentage of children who successfully find adoptive placements. Finally, we detail a series of simple improvements implemented by PAE to increase the adoptive placement rate through collecting more valuable information and aligning incentives for families to provide useful preference information.

Key words: community OR; public service; decision support; market design

Child adoption has become increasingly important in the United States and worldwide. In 2000, the total number of adopted children in the United States was 1.6 million, making up 2.5 percent of all children under 18 (cf. U.S. Census Bureau (2003)). Of these, 87% were U.S.-born and adopted through domestic channels and 40% were adopted through public child welfare agencies. The number of adoptions continues to rise: In 2008, approximately 136,000 children were adopted in the United States, which is a 6% increase over 2000 and a 15% increase over 1990 (cf. Flango and Flango (1995)). In this paper, we focus on the process related to the adoption of children in state custody who tend to be older, with special needs and other characteristics that make them more difficult to place in adoptive homes.

Despite the prospective parents' desire to adopt and children's need for permanent families, the adoption process for children in state custody suffers from inefficiencies due to overburdened case workers, challenges in coordinating various institutions involved, and

a mismatch in the types of children desired by families and those available for adoption. Currently, approximately 560,000 children and teens are in the U.S. foster care system, twice the number in foster care in 1987. On average, a child waits 3.65 years to be adopted. During this process children might change homes several times, an experience which causes long-term psycho-social difficulties (Rushton and Minnis 2002).

One of the most important roles that governmental Children and Youth Services (CYS) agencies play is to find adoptive families for children who are in state custody. For children who are older and have special needs, the task of finding a family is very challenging and agencies may not always succeed. According to Howard and Brazin (2011), over 25,000 children are emancipated each year from the child welfare system in the United States. Numerous studies have shown that children who “age out” of foster care without finding a permanent family suffer from high levels of unemployment, homelessness, early parenthood, and incarceration. Approximately 40% of teenagers who age out of foster care become homeless within 18 months (U.S. General Accounting Office (1999)). Legislation such as the federal Fostering Connections to Success and Increasing Adoptions Act of 2008 has made finding a family for every child an explicit priority of state governments.

The Pennsylvania Adoption Exchange (PAE) was established in 1979 by the state’s Department of Public Welfare to support county and nonprofit agencies as they attempt to find adoptive families for children who are difficult to place. In addition to promoting children through a website and hosting in-person matching events, PAE maintains a database with detailed attributes about children and preferences of families. PAE is mandated by the state to make match recommendations for families and children. We collaborated with managers at the Pennsylvania Statewide Adoption and Permanency Network (SWAN), a non-profit organization with the responsibility of administering PAE for the state of Pennsylvania, to implement changes in the matching recommendation process.

PAE’s match recommendation function has two main goals. First, it helps overcome geographical and institutional barriers in the adoption search process. Pennsylvania has 67 counties that are supported by 82 non-profit organizations, therefore a statewide perspective on the population of families and children can help identify prospective matches that an individual might miss. Second, the match recommendation system is supposed to save overburdened social workers the effort of sorting through extensive data—about 100

attributes for each child’s special needs and demographic characteristics with corresponding information about families’ preferences. Furthermore, PAE managers believe that case workers sometimes have too high expectations — i.e., they are waiting for the “perfect” family — so the process has a secondary goal of promoting a decision-making structure that helps evaluate the trade-offs of different prospective families.

In this paper, we investigate PAE to scientifically define the mission of an adoption exchange and help improve its matching process. Our project contributes an interesting and societally important new public policy and non-profit application of market design. Similar to other examples of applied research on matching markets, we focus on simple changes to PAE that address its most significant challenges, and we identify key elements of the child adoption market to investigate in future research. Our project has resulted in the collection of additional information, process changes, and a spreadsheet matching tool that is currently being used by PAE staff to recommend families. While the idea of a computerized matching tool for ranking families is not novel, we believe that we are the first to link its effectiveness to an increased rate of successful adoption through a discrete-event simulation of the adoption network. We also are the first to notice that the matching process may distort incentives for families in revealing their preferences truthfully and propose simple remedies.

The remainder of our paper is organized as follows. First, we provide context for the problem of a match recommendation system as it relates to research on the design of matching markets. We then characterize the challenges that PAE faces and assess the current system through case worker surveys and a regression analysis using child outcome data from 2005 to 2013. We describe how PAE recommends prospective families for children by comparing child needs to family preferences on a set of approximately 100 attributes using a spreadsheet tool. Based on this understanding of PAE’s role in the matching process, we analyze the value of the network and the information available to the network through a discrete-event simulation of PAE’s operation. We discuss our recommendations for the information that PAE collects, the decision rules for match recommendations, and the interaction with system participants. Finally, we conclude with a discussion of improvements that were implemented and possible future improvements for adoption and other domains.

Design of Adoption Markets

Despite the growing prominence and urgency of finding families for children in state custody, little is known about how to operationalize child adoption; i.e., how to define, analyze, and improve the family finding process. More generally, adoption has thus far received little attention from the social sciences. Fisher (2003) provides an account of how adoption has been overlooked by sociologists and other social scientists. In economics, only a limited number of studies deal with documenting and estimating the preferences of potential adoptive parents (see Baccara, Collard-Wexler, Felli, and Yariv (2013) and the references therein). The fundamental supply-and-demand imbalance for children of different demographic characteristics was first described using an economics framework by Landes and Posner (1978). The context of and reaction to this paper from the social sciences, which labeled it as the “baby selling paper,” is documented in Goodwin (2010). An empirical study by Baccara, Collard-Wexler, Felli, and Yariv (2013) identifies biases in preferences of prospective adoptive families for infant child adoption. They show that an adopted child’s desirability to prospective families depends heavily on the child’s age, gender, and race, with some of the greatest disparities accounted for by the child’s race. From social work literature, Hanna and McRoy (2011) offer an overview of the goals of matching for adoption. They also document tools to assess the quality of a family-child match that other governmental and nonprofit organizations have developed concurrent to PAE’s efforts.

We view the PAE as a two-sided matching market and rely on market design literature to frame our approach to the problem while using operations research and economics techniques to improve the current recommendation practice. In doing so, we focus on the economic aspects of resource allocation decisions that PAE faces surrounding the selection of match leads to be offered to system participants. As described in Roth (2010), “The challenge [of market design] is to understand complex markets well enough to fix them when they’re broken, and implement new markets and market-like mechanisms when needed.” PAE exists in a complex market in which the participants are families and difficult-to-place children, as well as their case workers, and our project attempts to help PAE overcome challenges by redesigning its matching process.

Early market design work focused on understanding and improving centralized clearing-houses that operate in the absence of prices and face institutional and ethical constraints. The seminal work of Gale and Shapley (1962) introduced the formal two-sided matching

framework. This theory was subsequently advanced and adapted to important applications such as the design of National Residency Matching Program in the US for matching medical school graduates to internships, residencies, and fellowships at hospitals (see Roth and Sotomayor (1990) for a detailed account of this approach and Roth and Peranson (1999) for redesign of the centralized matching algorithm). This approach was also retooled to be used in resource allocation problems such as the assignment of students to public schools (see Abdulkadiroğlu and Sönmez (2003)). On the other hand, the market design for kidney exchange clearinghouses have benefitted from algorithmic developments in resource allocation problems (see, e.g., Roth, Sönmez, and Ünver (2004, 2005), and for a survey on the matter see Sönmez and Ünver (2013)).

Another recent strand of literature in matching market design has focused on introducing new ideas to improve the functioning of a centralized or decentralized matching market rather than designing new clearinghouses to conduct the matching. The role of PAE can be compared to that of online dating sites such as eHarmony and advanced job matching websites such as Matchpoint Careers, which not only serve as two-sided platforms where interested parties meet, but also as advanced recommendation tools (cf. Zax (2013)). PAE's main function of decreasing frictions in the child and family search are similar to those of dating and job matching web sites. In particular, like the dating and job matching web sites, PAE can only recommend matches; it has little power to enforce them. Families, children, and their case workers all have complex preferences, and significant uncertainties exist in the attraction between individual families and children. Our contribution in the current paper is similar to some of the recent approaches where the goal is to make low-cost high-return recommendations for improvement rather than introducing radically new matching systems and algorithms. For example a recent market design paper by Coles et al. (2010) reports the introduction of a signaling tool for the academic job market for new economics PhDs. When PhD students apply for an academic position through the job search database of econjobmarket.org — the main academic job market application database for economics PhDs — they can now signal to, at most, two positions about their interest. Another market design paper, Lee, Niederle, Kim, and Kim (2011) reports an experiment to measure the effects of the use of a similar signaling device as mentioned above for online dating. Similarly, more recent literature on market design in live-donor organ exchange has focused on introducing new additional tools to increase the number of

transplants rather than new matching algorithms, such as the introduction of non-directed donor chains (see Roth et al. (2006), and Rees et al. (2009)).

In social sciences, the inner workings of online dating and job search markets have been documented in more detail relative to adoptions. For example, see Lee (2007) and Hitsch, Hortacısu, and Ariely (2010) for empirical assessment of sorting in online dating, and Niederle and Roth (2003) for empirical assessment of the scope of job market. A common theme of these papers is that a centralized platform; i.e., frictionless matching leads to better sorting and higher scope with respect to traditional decentralized search with friction.

PAE can even be compared to a new company called OrganJet, which uses private jets to transport patients in order to overcome the inefficiencies of regionally isolated organ donation networks (Ata, Skaro, and Tayur (2012)). A related study by Arikan et al. (2012) shows that broader sharing of the bottom 15% of kidneys (in terms of quality) from deceased donors leads to significantly increased procurement rates for those organs.

While previous strands of literature from operations and market design have had a remarkable impact on public policy with regard to a number of real-world applications, some in similar settings such as job market recommendations, we are not aware of any prior research on child adoption from a market design perspective. Our paper is the first as far as we know.

Child Adoptions in Pennsylvania

The statewide network's primary goal is to help find permanent families for children in state custody. Children who fail to be placed upon initial attempts at the county level are provided with extra services at the state level, including match recommendations from PAE. Between 2007 and 2012, an average of 239 children per year registered to received match recommendations from PAE. The Office of Children, Youth and Families (OCYF) of Pennsylvania's Department of Public Welfare mandates that children without an identified adoptive family must be registered with PAE within ninety days of termination of parental rights (TPR).

PAE managers believe that the best scenario for children is to be placed with a permanent family, even for a child who is at the threshold of reaching the age of majority. For these older children, a stable family is viewed as playing an important role as the child

enters adulthood, and statistics indicate that children who do not receive a family are at great risk of homelessness and incarceration. For children who are placed, the time during which the child is a legal orphan should be minimized. Furthermore, the suitability of the family for the child is also important. In particular, some families are better prepared than others to handle children with certain needs, whether medical, behavioral, or psychological. Based on their expectations and on how they view success in raising a child, some families are better suited than others for adopting children with troubled pasts.

Investigations into whether a family and child could be matched require a substantial time investment by case workers representing the children, social workers representing the families, and even the children and families, themselves. Overburdened case workers have limited time to review and process lengthy family profiles, and the human element of the matching process often necessitates face-to-face meetings before a decision can be made whether or not to pursue an adoption. Second, PAE stores limited information for use in understanding families' preferences, and discrepancies between families' stated and revealed preferences can cause difficulties for selecting the best families. Prior to TPR, a CYS worker seeks a suitable family for the child. If there is no clear kinship adoption possibility or potential family within the agency's local network, the CYS worker may contact SWAN and request match suggestions from PAE. The case worker must register the child with PAE in the ninety days after TPR if the worker has not identified a resource for a child through filing an intent to adopt. Through working with PAE regional coordinators, the CYS worker then receives between five and ten families to consider and pursue. After identifying interested families, the worker will then arrange an interview and consult with a committee comprised of social worker, legal, and other professionals to attempt to choose a family with which to place the child. After a series of family/child visits of increasing duration and decreasing supervision, if the match is deemed suitable, the child is placed with the family. According to Pennsylvania law, an adoption can then be finalized after six months.

Survey of Case Workers

Building on interviews with PAE managers and case workers, a PAE manager performed a survey of all case workers in spring 2011 to gain a broader picture of feelings about PAE across the state. Survey recipients included both county case workers and social workers

known as “child-specific recruiters” who work for private non-profit agencies and serve as an additional resource to county case workers in finding families for hard-to-place children. The manager sent the survey to the email addresses of the county case worker and/or the child-specific recruiter for all children active at that time with PAE. Seventy-seven completed responses were received — forty-three (56%) from public case workers and thirty-four (44%) from private social workers. The survey was divided into three sections: the first set of questions was general, the second set tried to elicit opinions about the current system, and the third aimed to assess attitudes regarding possible changes. In this section, we summarize the findings of the survey (see the Supplemental Files for the entire survey and further details about the survey results).

Respondents were first asked to indicate the helpfulness of various avenues of finding families: the centralized match recommendation system (known as “electronic match recommendations”), inquiries from a public website (www.adoptpakids.org) that PAE uses to promote children seeking adoptive families, decentralized lists of prospective families who registered with non-profit organizations (“affiliates’ lists of waiting families”) with which the worker had some connection, and in-person matching events sponsored by SWAN or held locally. Least helpful was PAE’s centralized matching system: 65% of respondents said that this never or rarely served as the initial source of prospective families for children who are successfully placed.

The survey’s second part gauged case workers’ satisfaction with the centralized system. The design of this part of the survey was assisted by the information gathered through family registration forms (CY 131) and child registration forms (CY 130). The first set of questions in this section addressed the Resource Family Registry, which is the database of prospective families and children seeking adoptions. For example, no respondents strongly agreed and only 32% agreed with the statement that “PAE does a good job of recommending the most suitable families via electronic matches from the Resource Family Registry for each child.”

In the last part of the survey, respondents testified to the difficulty of making placement decisions and case worker bias. More respondents agreed or strongly agreed (53%) than disagreed or strongly disagreed (37%) with the statement that they “know of case workers who struggle to make placement decisions for children because of emotional attachments to those children.” More interestingly, even more respondents agreed (65%) than disagreed

(23%) with the statement that they “know of case workers whose personal preferences lead to negative perceptions towards some families.”

Survey responses demonstrate both the ineffectiveness of the current matching system and the case workers’ mistrust of match recommendations from PAE. However, case workers expressed more positive views about possible helpfulness of the underlying registration data, with over 60% of respondents agreeing or strongly agreeing that family data is “helpful” for screening and child data is “accurate.” Taken together, the survey results indicate the potential value of a statewide matching network, and help to justify our efforts to improve PAE’s ability to help case workers find families for children in state custody.

Analysis of Child Outcomes

We reviewed registration and outcome information about children served by PAE to better understand adoption trends in Pennsylvania and the varying levels of difficulty in trying to find adoptive placements. PAE managers overcame significant challenges related to the decentralized nature of the adoption process in Pennsylvania to prepare this data set for our use. We are the first to analyze the relationship between child outcomes and child attributes upon registration in Pennsylvania, and the results of the analysis have provided insights about what children might require additional adoption resources and information to share with case workers as part of training on best practices.

Between 2005 and 2013, PAE assisted in the family-finding process for 1,853 children seeking adoptive families. This set of children was a subset of children in state custody with the goal of adoption; only when the matching process encounters difficulties at the county level does the search process shift to the state level. The mean age of a child upon registration with PAE was 9.41 years, and the median age was 9.63 years. Boys comprised 57.8% of all PAE registrants. Of these 1,853 children, outcomes were known for 1,514 children, as 283 were still active upon the creation of the data set in May 2013, and outcomes were missing for 56 children. Otherwise, child outcomes are known and grouped into categories, each of which has a value between 0 and 1 that is given by PAE managers. The most desirable outcome, a finalized adoption, has a value of 1. Emancipation, which can be referred to as “aging out” of the system, is the least desirable outcome and has a value of 0. Other positive outcomes include permanent guardianship arrangement (0.8) and living with a relative (0.7), among other scenarios. An outcome of “hold” with a foster

care arrangement is considered a neutral outcome and has value of 0.5. Other negative outcomes include placement in a residential facility (0.2) and a goal change so that the child's case worker is no longer seeking an adoptive placement for the child (0.1).

Of the children for whom outcome data are known, the county case worker succeeded in finding a finalized adoptive placement for 41.4% (627 of 1,514) of children. Another 19.1% of children have lesser positive outcomes with values of 0.7 or 0.8. Negative outcomes with values less than or equal to 0.2 are experienced by 26.2% of children, with 12.4% of children aging out of the system. The remaining portion (13.3%) of children have neutral outcomes. Using the values given by PAE managers, the expected outcome value for a child in the data set is 0.64.

Regression Model

We developed a linear regression model and a logistic regression model to analyze the relationship between children's attributes when they were registered with PAE and their outcomes. Eighty-eight factors were created from registration data, and the outcome was used as the dependent variable. Specifically, the outcome value was used for the linear model and a binary variable with positive outcomes having value one and neutral and negative outcomes having value zero for the logit model. The square of the child's age upon registration was included to represent the increasing importance of age for older children. The date of registration was expressed in the fractional number of years after January 1, 2005. Gender was represented by a binary variable, as were whether the child had a designation of being African-American and/or Hispanic. Another binary variable represented whether the child had more than one race designation. Eighteen binary variables represented the items under the sections labeled "Educational Status" and "Special Needs" on the CY 130 forms, and an additional variable counted the number of binary variables with positive responses. The last category of variables was 61 questions in the CY 130 section labeled "Characteristics of Child," with a binary variable for each question representing a "Yes" answer. A variable that counted the number of positive responses in the section was also used.

Starting with these 88 variables, we performed a backward stepwise procedure using the Akaike information criterion for both the linear and logit models to select which variables to include in the model. Linear and logit regressions were performed on the union of the variables from these two models (Table 1). To roughly assess model performance for the

linear model, rounding the predicted outcome to the nearest outcome value and classifying it as positive, negative, or neutral, the model classifies 60.0% into the correct one of three categories. Only 8.1% of children had a positive or negative outcome that was errantly predicted as the opposite outcome. All other prediction errors involved the neutral outcome. When binary child outcomes are considered for the logistic models, 75% of outcomes are correctly predicted. For children with negative/neutral outcomes, 59% were correctly predicted, and 85% were correctly predicted for children with positive outcomes.

The child's age upon registration plays a significant role in the regression model and confirms that older children become increasingly difficult to place. For the linear model, the predicted likelihood of success decreases by 0.034 per year at age eight and by 0.087 per year at age 16. The accumulative decrease in the expected outcome compared to a newborn child is then 0.054 for an eight year-old child and 0.535 for a 16 year-old child. Although PAE managers anticipated this general trend, they found the quantification of this relationship to be very helpful as they instruct case workers around the state on best practices. Specifically, they can encourage case workers to register children with the PAE as soon as possible in the adoption process and perform a search through PAE in parallel with family reunification or other placement efforts.

Mental retardation, which had a coefficient of -0.109 for the linear model, was the most negative of the significant special needs factors from the "Child's Statuses" section of the CY 130 form for both models. Two special needs factors — having siblings and attending school in a general education setting — had statistically significant positive coefficients for both models. For the "Characteristics of Child" attributes, the result that surprised PAE managers was a linear regression coefficient of -0.118 for children who use foul or bad language, which was also significantly negative for the logit model. PAE managers found this information valuable to share with case workers as part of training on how to identify challenges to a successful placement. Other factors that were significantly negative with 95% confidence for both models were a difficulty accepting and obeying rules, a desired contact with siblings, contact with birth parents, and contact with the former foster family. The child's gender was revealed not to have a significant effect in either model, but outcome value decreases of 0.034 and 0.051 (for the linear model) were expected for children with African-American and Hispanic designations, respectively, although the logit model did not determine designation as African-American to have a significant effect.

Table 1 We choose 28 factors from the available 88 factors to model child outcomes using ordinary least squares and logistic regression methods. Age upon registration, which is a negative factor for children six years of age and older, was the most important factor for predicting outcomes.

	<i>Dependent variable:</i>							
	Outcome Value				Outcome (Binary)			
	<i>Ordinary Least Squares</i>				<i>Logistic</i>			
						Freq.	Import.	
Constant	0.794	***	(0.046)	1.516	***	(0.372)		
Age upon Registration (years)	0.020	**	(0.009)	0.102		(0.075)		High
(Age upon Registration) ²	-0.003	***	(0.0005)	-0.017	***	(0.004)		High
Registration Year (after 2005)	-0.009	**	(0.004)	-0.059	*	(0.031)		
Male	0.019		(0.017)	0.100		(0.128)	57.1%	High
African-American	-0.034	**	(0.017)	-0.198		(0.132)	42.5%	High
Hispanic	-0.051	**	(0.024)	-0.303	*	(0.179)	14.1%	High
SPECIAL NEEDS								
Mental Retardation Diagnosis	-0.109	***	(0.031)	-0.562	**	(0.230)	9.0%	High
Multiple Placement History	-0.035	*	(0.018)	-0.189		(0.137)	45.6%	Medium
Drug Exposed Infant	-0.020		(0.026)	-0.100		(0.202)	11.6%	Medium
Emotional Disability	-0.019		(0.022)	-0.071		(0.162)	20.2%	Medium
General Education	0.064	***	(0.019)	0.353	**	(0.146)	37.1%	
Siblings	0.085	***	(0.019)	0.465	***	(0.143)	47.3%	High
CHILD CHARACTERISTICS								
<i>Blind</i>	-0.164	*	(0.085)	-0.899		(0.611)	1.0%	Medium
<i>Uses Foul or Bad Language</i>	-0.118	***	(0.027)	-0.613	***	(0.194)	15.0%	Medium
<i>History of Running Away</i>	-0.086	**	(0.043)	-0.443		(0.321)	4.2%	High
<i>Desires Contact with Siblings</i>	-0.079	***	(0.020)	-0.443	***	(0.152)	59.4%	Low
<i>In Contact with Former Foster Family</i>	-0.064	***	(0.022)	-0.353	**	(0.162)	18.8%	Low
<i>Rejects Father Figures</i>	-0.061	**	(0.031)	-0.345		(0.230)	8.5%	Low
<i>Difficulty Accepting and Obeying Rules</i>	-0.061	***	(0.022)	-0.337	**	(0.160)	36.9%	Low
<i>In Contact with Birth Parents</i>	-0.058	***	(0.020)	-0.327	**	(0.153)	26.0%	Low
Num. of Characteristics Present	0.007	***	(0.003)	0.034	*	(0.020)		
<i>Parent(s) with Criminal Record</i>	0.017		(0.018)	0.087		(0.138)	51.6%	Low
<i>Difficulty Relating to Others</i>	0.018		(0.022)	0.101		(0.168)	31.0%	Low
<i>Speech Problems</i>	0.024		(0.024)	0.176		(0.191)	18.4%	Low
<i>Previous Adoption or Disruption</i>	0.038	*	(0.021)	0.220		(0.155)	24.1%	Low
<i>Strong Ties to Foster Family</i>	0.041	**	(0.018)	0.226	*	(0.134)	54.2%	Low
<i>Vision Problems</i>	0.042	*	(0.023)	0.224		(0.175)	17.1%	Low
<i>High Achiever</i>	0.054	**	(0.025)	0.283		(0.190)	13.1%	Low
Observations			1,514			1,514		
R ²			0.345					
Adjusted R ²			0.333					
Akaike Inf. Crit.			1,697					

Note: The values in parentheses indicate the standard deviation. Italicized variable names refer to “Characteristics of Child” questions on CY 130 form. Also, *p<0.1; **p<0.05; ***p<0.01. “Freq.” refers to an attribute’s prevalence among the observation. “Import.” refers to the existing default weight given to the attribute before our analysis.

To understand how the regression results compare to managerial intuition about the different factors’ relative importance, we compare them to PAE managers’ existing classification the factors. As part of a previous attempt to create a family ranking tool that had encountered difficulties, managers divided the factors from the CY 130 form based on their

perceived importance into groups with 15 factors as “high,” 18 factors as “medium,” and 41 factors as “low.” Of the ten significant factors with the most negative OLS coefficient, managers had assessed three of the factors (Hispanic, mental retardation diagnosis, and history of running away) as of high importance and two factors (blind and uses foul or bad language) as of medium importance. They classified the remaining five factors as of low importance. These five factors, which include two behavioral traits and three related to a child’s social connections, merit closer attention and a higher weight in the matching process to help identify families more suited to a child’s needs. As a result of this analysis, they ultimately decided to reclassify almost all significant ($p < 0.1$) factors with a negative OLS coefficient as of high importance in the match assessment tool, which we will further discuss later in the paper. The lone exception to this rule was multiple placement history, which remained at medium priority due to the low magnitude of its coefficient. PAE managers particularly appreciated the suggestion to increase the importance of a factors related to a child’s social connections, as they identified those connections as a frequent obstacles to adoption.

PAE managers identified four other characteristics, which represent some of the most severe behavioral needs on the CY 130 form (e.g., abusing animals), as of high importance that did not appear as significant factors. However, increased managerial attention to finding families related to these children’s needs might have led to their exclusion from the final model for predicting child outcomes, and we do not make a recommendation about whether to reduce the emphasis on these child characteristics.

Match Assessment Tool

PAE is only one of several governmental and non-profit institutions that have developed tools to assess a possible match between a child and family. Hanna and McRoy (2011) describe the practice of matching in adoption as a means of finding families that have the right capabilities for handling a child’s special needs and identifying gaps in a family’s capabilities. They emphasize the need for standardization and data collection, point to match assessment tools as an important part of the family finding process, and review seven different tools used in practice by public and private agencies. Although some tools use more attributes (up to 277) than PAE, the design of PAE’s existing matching tool surpasses all seven tools in terms of nuance in how attributes are weighted and its ambitions for helping to find families within a statewide network.

Table 2 PAE managers use data on 76 attributes to recommend families for children. Weights displayed represent weights used in the existing algorithm.

Category	Number of Attributes (Scoring Weight)			Child Attribute Values (Family Preference Values)
	High (100)	Medium (10)	Low (1)	
Child Demographics				
Age	1	0	0	Current age (Max/min age)
Race/Ethnicity	6	0	0	Applicable/ <i>not applicable</i> (Preferred/not preferred)
Gender	1	0	0	Male/female (Male/female/either)
Child Status				
Educational Status	0	1	0	Applicable/ <i>not applicable</i> (Approved/not approved)
Special Needs	6	7	0	Applicable/ <i>not applicable</i> (Approved/not approved)
Characteristics of Child				
Health	0	3	7	Yes/ <i>no/unknown</i> (Acceptable/will consider/unacceptable)
Education	0	1	7	Yes/ <i>no/unknown</i> (Acceptable/will consider/unacceptable)
Characteristics and Behaviors	5	5	11	Yes/ <i>no/unknown</i> (Acceptable/will consider/unacceptable)
Connections and History	0	0	14	Yes/ <i>no/unknown</i> (Acceptable/will consider/unacceptable)
Contact with Birth Family	0	0	1	Yes/ <i>no/unknown</i> (Acceptable/will consider/unacceptable)

Note: Items in italics indicate child attribute values for which the attribute does not count as part of the total matching score.

According to its intended design — which we formally express in Algorithm 1 — PAE’s match assessment tool computes a family’s score between 0% and 100% for a child based on 78 pairs of child attribute values and family preferences from CY 130 and CY 131 registration forms (Table 2). PAE managers had assigned a number of possible points to each of these pairs, which is 100 for items of high importance, 10 for items of medium importance, and 1 for items of low importance. When a special need or other attribute is not applicable or of an unknown state for a child, no points are either eligible or awarded to each family for that attribute. We note that PAE also does not give preference to families who say that a special need is undesirable when matching a child who does not have that special need. This practice could create incentives for families to hide special needs that they can accommodate if the families want to be considered for children without those special needs. For attributes that are applicable for a child, a positive family response — a child’s age within the family’s range, a matching gender, and answers of preferred,

Algorithm 1 Design of existing algorithm for computing scores for each family for a given child.

Inputs:

- 1) Attribute value c_k for each attribute $k = 1, 2, \dots, K$ corresponding to Table 2 for a single child
- 2) Preference value f_{jk} for each family $j = 1, 2, \dots, F$ and each attribute $k = 1, 2, \dots, K$
- 3) Weight $w_k \geq 0$ for each attribute $k = 1, 2, \dots, K$

Output: Set of family matching scores $\{y_1, y_2, \dots, y_F\}$

```

for  $j = 1$  to  $F$  do
  Points possible  $x_j^{TOT} \leftarrow 0$ 
  Points earned  $x_j \leftarrow 0$ 
  for  $k = 1$  to  $K$  do
    if  $c_k \notin \{\text{"not applicable"}, \text{"unknown"}, \text{"no"}\}$  then
       $x_j^{TOT} \leftarrow x_j^{TOT} + w_k$ 
      if  $f_{jk} = \text{"will consider"}$  then
         $x_j \leftarrow x_j + w_k/2$ 
      else
        if  $f_{jk}$  is compatible with  $c_k$  (see Table 2) then
           $x_j \leftarrow x_j + w_k$ 
        end if
      end if
    end if
  end for
   $y_j \leftarrow x_j/x_j^{TOT}$ 
end for
return  $\{y_1, y_2, \dots, y_F\}$ 

```

approved, or acceptable — receives all possible points for the item. An answer of “will consider” receives 50% of the possible points. Otherwise, the family’s answer receives no points. For a specific child, the family’s score is simply the sum of points received for their answers divided by the sum of possible points for the child.

The appendix provides an example match score for a child and family. Based on the child’s presenting needs, a total of 425 points are possible for any family matched to the child, including 100 for age, 100 for race, 100 for gender, 60 for special needs, and 65 for characteristics (not shown). The family receives 100 points because the child is within the family’s stated minimum and maximum ages, 100 points because the family’s preference for a white child matches the child’s race, and 100 points as the family stated that either gender was acceptable. The family receives 40 out of 60 possible points for the child’s special needs, losing points as the family is not approved to adopt a child with “drug exposed infant” and “abuse history” designations. For child characteristics, the family is allotted 44 of 65 possible points for a total of 384 out of 425 possible points, which translates to a score of 90.35%.

In recent years, PAE has struggled to make match recommendations that help case workers to find and assess families. Specifically, coordinating algorithm design with an

information technology contractor and managing data collected over time across Pennsylvania’s 67 counties proved difficult for PAE, and child case workers received unhelpful or illogical match recommendations as a result of a flawed implementation of Algorithm 1. The automated match recommendations were even abandoned for a period of over two years during which PAE regional coordinators manually searched through CY 130 and CY 131 forms to provide match suggestions. In the case worker survey, no respondents strongly agreed and only 30% of case workers agreed (and 27% disagreed or strongly disagreed) with the statement that they were “satisfied with the electronic match recommendations” that they received.

Recognizing the shortcomings of the current rules for choosing matches, we worked with the SWAN managers to redesign the matching tool. This resulted in a spreadsheet-based algorithm that use PAE data about families and children to select matches. For ease of implementation, we focused on policies that had the same form as the PAE’s match assessment tool. In particular, we considered policies that are based on a point system and ranked families according to some compatibility criteria. Rather than making specific assumptions about the relative importance of each criteria, our method offers the PAE managers the flexibility to select his desired weights and any other geographical constraints (see Figure 7). To find select prospective families for a child, the spreadsheet tool computes a ranked list of families for a child using CY 130 and CY 131 information stored in tables elsewhere in the matching tool.

The matching algorithm that we implemented differs in two aspects from those that have been applied in prominent, centralized two-sided matching applications (e.g., medical match, school choice) enumerated earlier in the paper. First, the algorithms studied in that literature, such as Gale and Shapley’s celebrated deferred acceptance algorithm, produce a batch of final matches that are concurrently implemented. By contrast, our algorithm generates a list of mere recommendations, which may be implemented in conjunction with the judgment of the professionals. In this sense, our approach is closer to that adopted by the literature on semi-decentralized matching platforms such as those for online dating and job matching, discussed earlier. The literature on the latter however almost exclusively focused on estimating participants preferences as opposed to increasing match quality, which we pursue here. The second, more subtle difference is that in centralized matching,

participants are required to rank-list the available options e.g., students providing preference rankings over various schools, whereas here such ranking information is impossible to directly elicit because of limitations such as market size and informational asymmetry. One main function of the above algorithm can be seen as constructing such preferences from given pieces of information in the CY 130 and CY 131 forms and using them as the basis of recommendations.

At first, the tool replicated the matching rules and was used in conjunction with offline PAE data. Through this exercise, the Excel tool was recognized to be valuable not only as a brainstorming tool but as a working prototype to execute customized matches with improved rules. In particular, custom weights can be specified for each attribute to help determine which attributes are most important for selecting a family for a child. The matching tool prototype has also allowed PAE regional coordinators to gain intuition about the matching rules and begin to think about what matching rules help produce the best matches for children.

Before our proposed matching tool was implemented, we tested it offline via a spreadsheet prototype which replicated the matching rules that were used in conjunction with PAE data. Through this exercise, the spreadsheet tool was recognized to be valuable not only as a validation tool but as a working prototype for PAE regional coordinators to work with us to design simple and easy-to-implement improvements (see Process Improvements and Results section for specific details). We have supported several development iterations of the matching tool to add features and functionality. PAE started to make match recommendations based on the matching tool in November 2013 after almost a three year hiatus.

Simulation of the Pennsylvania Adoption Exchange

To examine the impact of a simple matching tool's effective use on the network's overall adoption rate, we represent PAE's matching process as a discrete-event simulation. We show the value of a statewide pool of families compared to a decentralized search process, and analyze how the ability of PAE to predict a match's success increases the number of matches. Specifically, we model how different levels of information about child attributes and family preferences affect the number of matches and the number of attempts before a successful match. We rely on the results of the regression analysis from the previous

section to identify the most important child attributes for the simulation and additionally introduce relevant family registration data to model family preferences.

As an alternative to conventional techniques such as clinical trials, which would require many years to evaluate, discrete-event simulation has long been used to estimate the effects of policy changes, especially organ allocation policies (cf. Ata et al. (2012)). Similarly, our discrete-event simulation model of PAE's operation is modular and based on input parameters estimated with real data. Some of the simulation studies on organ allocation use a finite-horizon model, which is necessary since the data is highly time-dependent and reaching steady-state is very unlikely unless there is an alternative therapy for transplant patients. However, we assume stationary parameters in our simulation model, which is justified in the context of child adoption as the population characteristics of children and families in the system do not change dramatically over time.

The adoption network is divided into regions that constitute separate adoption networks defined by geographical and/or institutional barriers. Children may only be adopted by families that reside within the same region. Decreasing the number of regions to increase the size of each region provides each separate network with more prospective families and more children to match. We model the state of the PAE before our project as 20 separate regions due to the ineffectiveness of the central matchmaker. In that case, each region may correspond to a large county or a coalition of smaller counties in Pennsylvania. Because county case workers do face geographical limitations in matching, we do not expect a perfect centralized matchmaker to operate as a single region. Instead, managerial insights are primarily motivated by two cases: doubling the region size — i.e., dividing the state into 10 regions instead of 20 regions — and a system with four regions corresponding to PAE regional coordinators who provide match recommendations to county case workers.

Children

Children are defined by their age, number of special needs, and region in which they reside. A younger child is generally preferred to an older child, and a child with fewer special needs is generally preferable to a child with more special needs. The age attribute corresponds to the child's age upon registration with PAE. Using available data for the 1,853 children who have been registered with PAE, we fit the data using the input analyzer tool of @Risk and compared alternatives using a q-q chart. A beta distribution with shape parameters $\alpha = 5.7736$ and $\beta = 4.8877$ and scaled to be within the interval $[-5.4648, 22.738]$ was

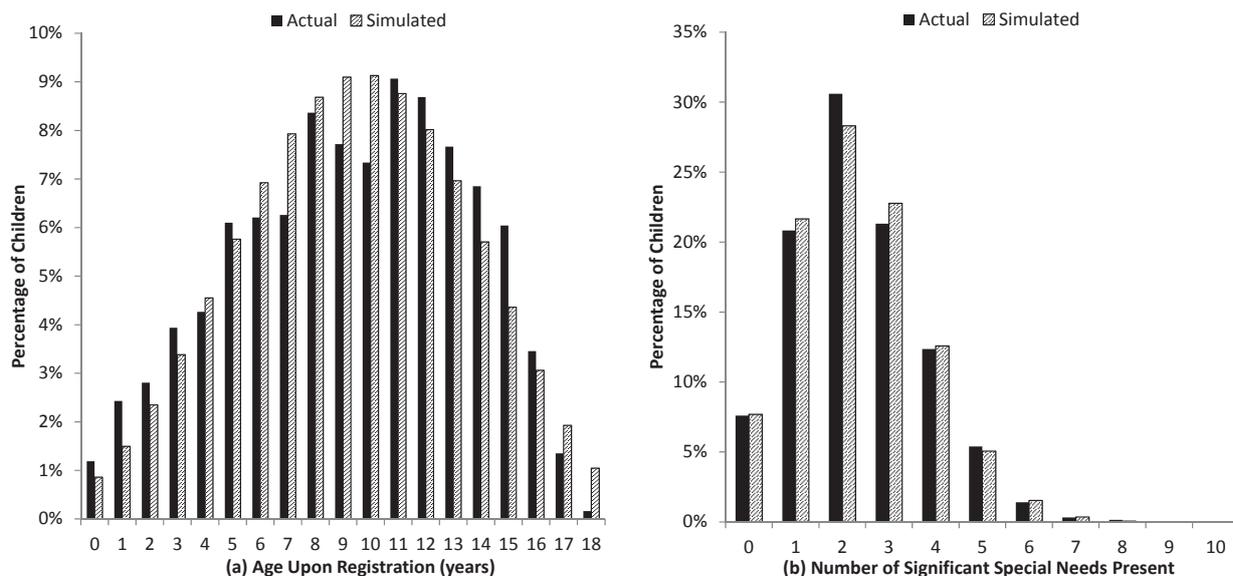


Figure 1 Based on age upon registration data for 1,853 children, we simulate the child's age using a beta distribution and use a binomial random variable to simulate the number of significant special needs present for each child.

determined to produce the best fit (Figure 1a). In the simulation, any age values outside the interval $[0.0, 19.0]$ were discarded and resampled.

The special needs attribute corresponds to a count of the presenting attributes out of the ten child special needs attributes that had a significant negative coefficient with a value less than -0.05 in the OLS regression analysis. This cutoff is arbitrary and used only in the simulation analysis to designate attributes of high importance. As with age, we fit the data using @Risk's input analyzer, and modeled the number of special needs present as a binomial random variable with parameters $n = 14$ and $p = 0.16741$ (Figure 1b). Any special needs values greater than 10 were discarded and resampled.

The registration age and number of significant negative special needs are positively correlated with a correlation coefficient of 0.230. Therefore, we used a Normal-To-Anything (NORTA) process with two base vectors that have a correlation coefficient of 0.239, which was obtained via a simulation approach. The standard multivariate normal vectors that follow a NORTA distribution are transformed to the age and special needs distributions using the method as described by Biller and Ghosh (2006).

The value of the child's region attribute is a random variable uniformly distributed over all regions. Children arrive in the system as a Poisson process with a rate of 239 per year,

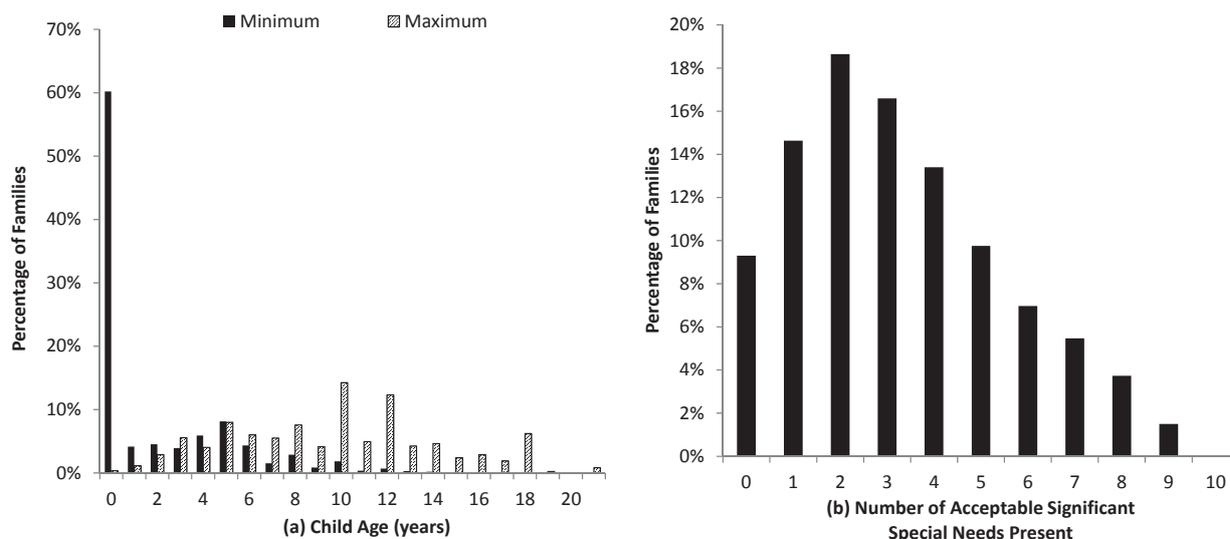


Figure 2 To simulate family preferences, we sample actual preferences from 2,194 registered families.

which is the average number of children who were fully registered with PAE annually between 2007 and 2012.

Families

Families are defined by their region of residence and their preferences for an adoptive child’s age and maximum number of special needs, as well as the relative weight of the age preference compared to the special needs preference. They arrive to the matching system as a Poisson process with rate 282 per year, which is the mean number of families to fully register with PAE as approved adoption resources each year between 2007 and 2012. Managers estimate that there are in the order of 1,000 prospective families available at any point in time, which implies that the expected time in system is 3.55 years by Little’s Law. Thus, we model the family’s time in the system as an exponential random variable with a mean of 3.55 years, as the distribution of families’ time in system is not tracked by PAE. We note that higher number of families compared to the number of children in the system creates a disparity in the distribution of children available to adopt and the distribution of family preferences, which is reflected in the PAE system through children who “age out” of the system without an adoptive placement.

We model the families’ behavior as myopic, accepting the first child that they are offered for which their utility of a match with the child is positive. The model behind the family acceptance decision is discussed in the appendix. Once a family accepts a child, the family

departs from the system. The values for a child's minimum age, maximum age, and number of acceptable special needs are sampled together from the data on 2,194 families (Figure 2). As with children, families are uniformly distributed over the regions.

Matching Process

We represent the matching process as a series of events that take place upon the arrival of a child in the system, which is viewed as the driver of the matching process by PAE managers. Matches are offered sequentially to families within the child's region. In each system, families are sorted according to criteria that correspond to PAE's operation with different levels of information. The highest-ranking family is selected and offered the child as a match. If a family accepts a match (i.e., its utility for the match is positive), both the family and child depart from the system. Otherwise, the family remains in the system and a match is attempted for the child with a total of ten families in sequence. For simplicity, we model the matching process as an instantaneous event, although in practice time elapses between sequential matching attempts. If no match is found for the child, the child departs from the system. A flowchart representing this process is provided in the appendix.

We present three methods for ranking families to investigate the value of information. The three methods and their interpretations are as follows:

1. *Critical Attribute (CA)* represents a system in which case workers can search for families based on either the age or special needs attribute due to constraints on their search time and effort. We represent the matching process before our collaboration with PAE as following the CA policy.
2. *Unknown Weight (UW)* represents a simple version of a centralized matching system that is limited in its ability to properly incorporate family preference information. Age and special needs attributes are given equal weight in this model.
3. *Full Information (FI)* represents an improved version of PAE's centralized matching system with families sorted based on a known age preference, special needs preference, and preference weighting term.

Simulation Results

We begin our analysis by comparing the simulated mean percentage of children matched over the five-year horizon based on the attraction variability and number of regions for CA, UW, and FI decision rules. The adoption rate increases with the amount of information

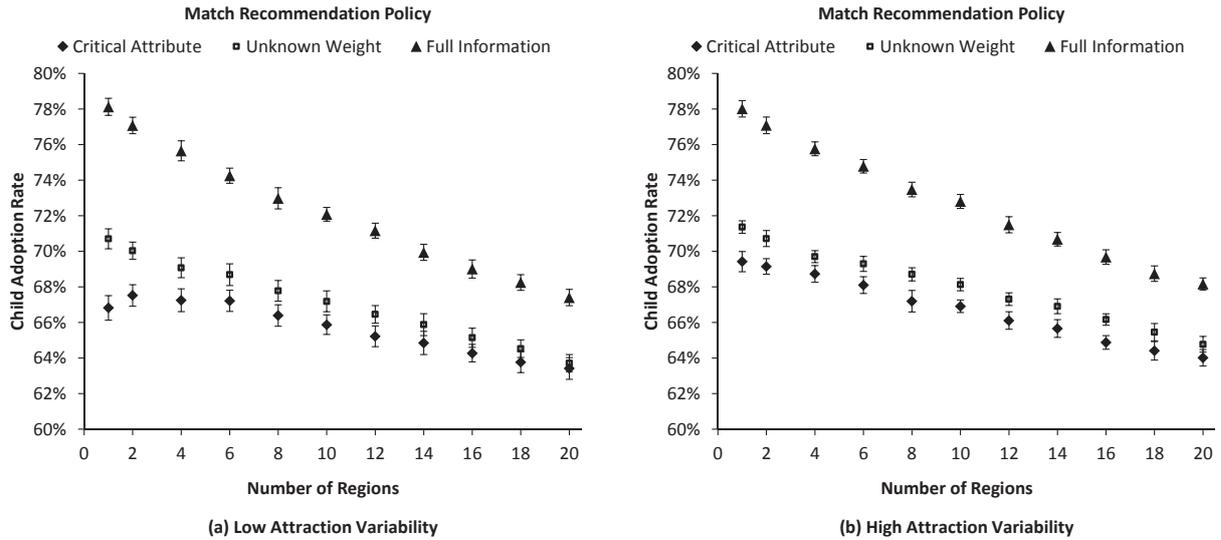


Figure 3 The child adoption rate increases with the quality of information available for matching and decreases with the number of regions (i.e., the segmentation of the network). Bands represent 95% confidence intervals.

about the families' preferences utilized in the match recommendation process (Figure 3); i.e., CA exhibits a lower adoption rate than UW, which in turn has a rate lower than that of FI. The UW policy only slightly improves upon the CA policy with a maximum increase of 3.9 percentage points. However, the FI policy improves on the CA policy by between 4.0 and 11.3 percentage points in the child adoption rate.

The mean adoption rate always either increases as the number of regions decreases or shows a statistically insignificant decrease. A completely centralized system (i.e., one region) results in an adoption rate that is between 3.4 and 10.7 percentage points higher — depending on the attraction variability and the recommendation rule — than the completely decentralized case with 20 regions. This validates the role of a statewide network. With a larger pool of families, it is more likely for a family to exist that seeks the type of child being matched or can accommodate the child's special needs.

In addition to an increase in the adoption rate for the UW and FI policies compared to the CA policy, the better use of information also corresponds to a decrease in the mean number of match attempts until a successful adoption for children who are successfully adopted (Figure 4). We study this metric as a proxy for two important secondary measures of success for the adoption network: the workload for case workers and time in system for the child. Fewer attempts until success means less work for overburdened case workers

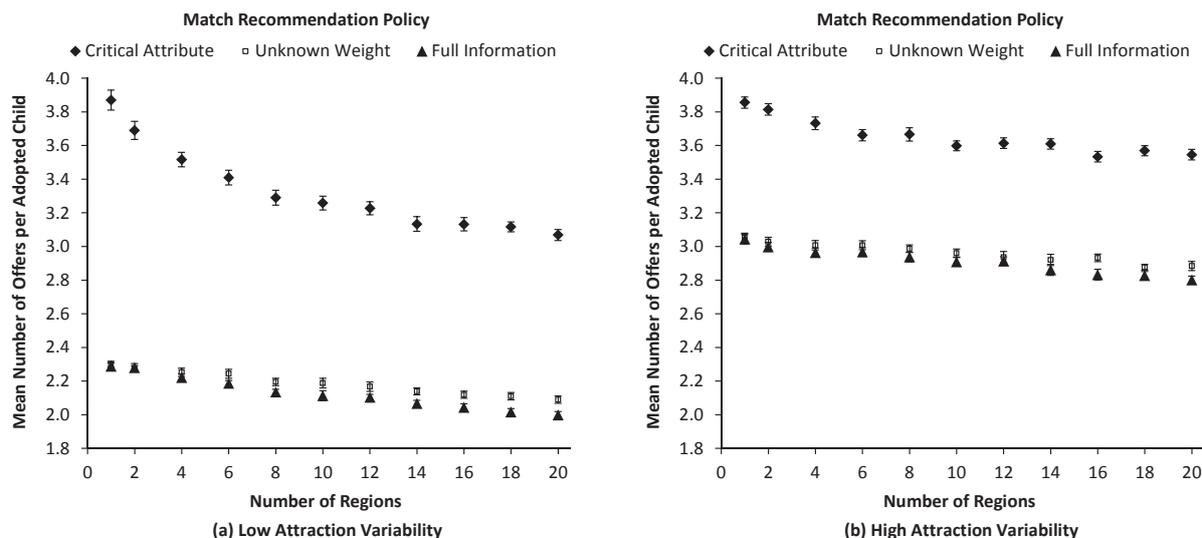


Figure 4 Improving the information available for matching reduces the average number of attempts before a successful adoption. Bands represent 95% confidence intervals.

and a shorter time in state custody for a child. Depending on the number of regions, the UW and FI policies cause the mean number of match attempts until success to decrease between 32% and 41% when attraction variability is low and between 17% and 21% when attraction variability is high.

We further investigate the effect of the attraction variability, which represents the unpredictability of attraction between an individual family and child. When the attraction variability is high, match success is inherently more difficult to predict, which results in an increase in the mean number of attempts per successful adoption of up to 0.81 (Figure 4). Comparing the change in the mean adoption rate as the number of regions decreases, the difference between the low and high attraction variability cases (in relation to the calibration point) is almost always less than 1%. This indicates that the underlying match unpredictability has relatively little impact on the mean adoption rate compared to the matching rule and number of regions. The only exceptions are for the CA policy with 1 or 2 regions when high attraction variability results in an adoption rate that is 1.0-2.0 percentage points higher than the adoption rate when the attraction variability is low. In these cases, the lower match success predictive power of the CA policy for ranking families, the high attraction variability, and the small number of regions means that it is relatively more likely for families who are offered children later in the sequence of ten offers to be stronger candidates.

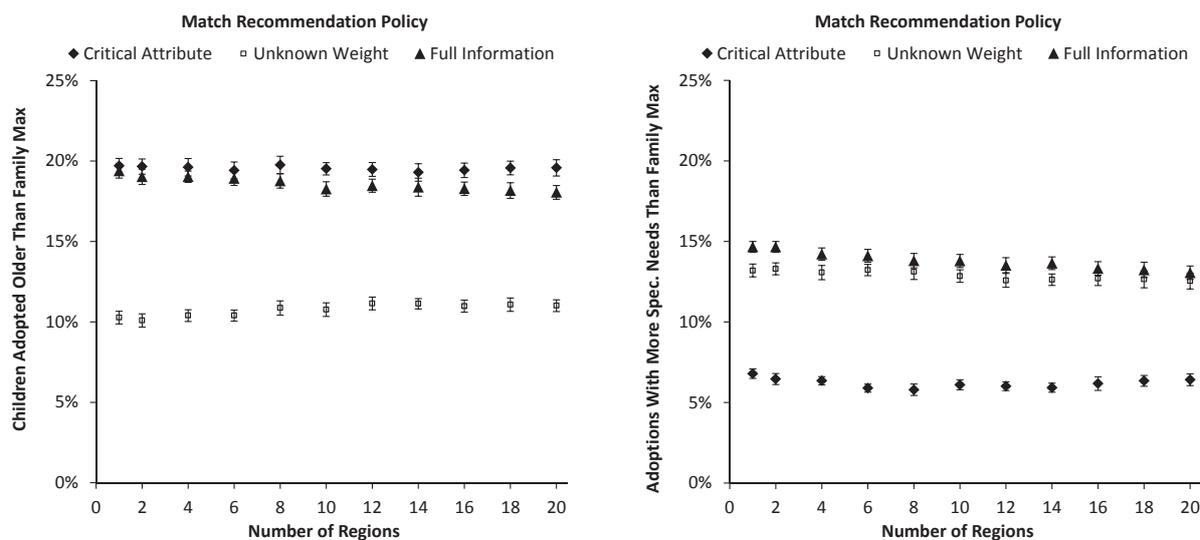


Figure 5 Families' preferences are violated more frequently in adoptions for the Critical Attribute and Full Information policies than the Unknown Weight policy for a system with low attraction variability. Bands represent 95% confidence intervals.

We additionally consider metrics that provide insights into the quality of matches for children and families. Specifically, we record the percentage of families who adopt a child with attribute values larger than the family's stated maximums (Figure 5). We use these metrics to represent the quality of matches for children, especially in terms of the special needs attribute as a higher stated tolerance for special needs represents the increased ability of a family to accommodate a child's needs. We also record the average nominal family utility for adoptions, which we use to represent the quality of matches from the perspective of families (Figure 6).

We first examine the UW policy, which assumes that families give equal weight to the age and special needs components of the utility function, and the case of low attraction variability. Under the UW policy, the percentage of families who adopt children with attributes above stated maximums is between 10.1% and 11.1% for the age attribute and between 12.5% and 13.3% for the special needs attribute, depending on the number of regions and attraction variability. Compared to the UW policy, the CA policy results in matches for which more families adopt children higher than their stated maximums more frequently for the age attribute (19.3% to 19.8%) and less frequently for the special needs attribute (5.8% to 6.8%). Given that the penalty for exceeding a child's age in a utility function is relatively low for children who are younger than teenagers, the difference

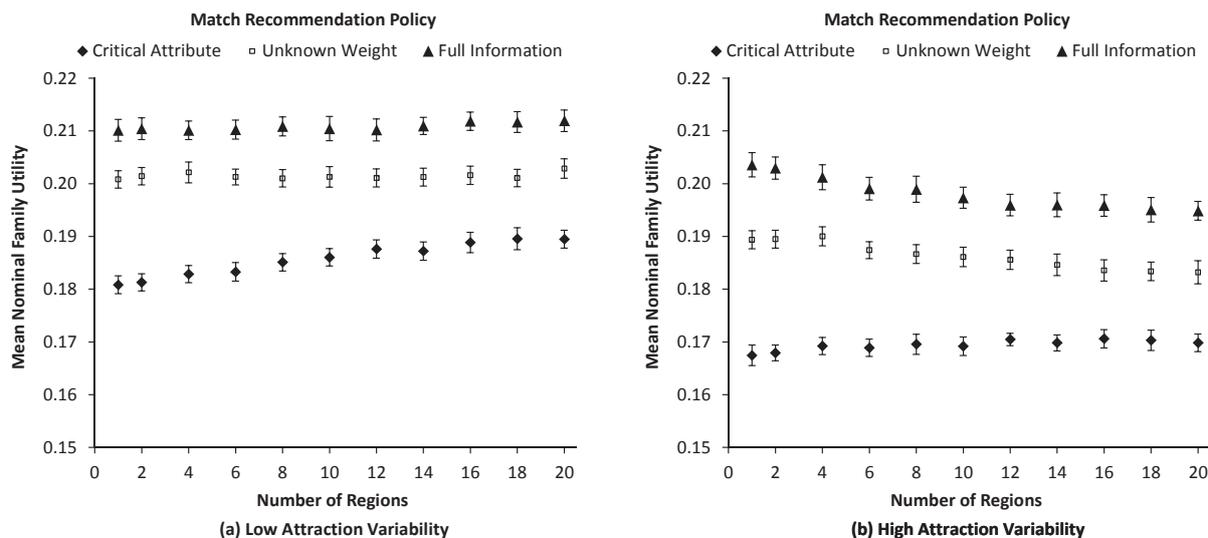


Figure 6 The mean family nominal utility of successful matches increases with the quality of matching information. Bands represent 95% confidence intervals.

between the UW and CA policies can be explained by the practice of the CA policy to ignore the age attribute entirely if it is nominally less important than the special needs attribute.

For the FI policy, the percentage of families who adopt children with attribute values greater than their stated maximums is between 18.0% and 19.4% for age, which is almost as high as the CA policy. For the special needs attribute, it is between 13.1% and 14.7%, which is higher than the UW policy. This helps to explain the success of the FI policy with respect to the mean adoption rate; it uses knowledge of the weighting term α to violate families' preferences when relatively insignificant to propose plausible matches that have a greater chance of success. This explanation is supported by Figure 6, which shows that the mean family utility increases from CA to UW and from UW to FI. We omit a corresponding chart for high attraction variability that shows a similar relationship between the three policies, although in all cases the percentage of families who adopt children with attribute values higher than their stated maximums is higher due to increased attraction variability.

The managerial insight regarding match quality is that improving the matching rate corresponds to more matches that violate families' stated age and special needs preferences when either of those preferences may be relatively less important to families. This highlights the importance of better understanding families' preferences as expressed through registration data and actual match rejection decisions. For example, allowing families to

state a preference of “willing to accept training” for individual child special needs questions in registration data could allow PAE to better estimate how families weight questions. Similarly, PAE should expect to offer training more frequently to families on how to handle specific child special needs to maintain match quality for children.

These results justify the value of a statewide adoption network and show that the quality of information about family preferences is critical to its success. Furthermore, we have shown that better information improves secondary metrics of system performance that can be interpreted as reducing case workers’ workloads and time in system for children. If the adoption network can even just double the size of its regions so that the number of regions is 10 instead of 20 while improving how it elicits family preferences for matching (i.e., follow the FI policy), the number of successfully adopted children can increase by approximately 21 children per year.

Process Improvements and Results

To achieve the possible improvement in the adoption rate demonstrated in the simulation, we worked with PAE to improve the information it collects and its family ranking tool for matching. In this section, we discuss these changes, as well as the incentives that affect how participants reveal their preferences. A spreadsheet matching tool in use by PAE regional coordinators provides the most tangible evidence of improvements from our collaboration. While using a computerized tool for matching is not unique to our project, we are the first to connect the quality of preference information to the overall network adoption rate and show how underlying incentives for how families reveal their preferences can diminish the usefulness of recommendations. We also suggest improvements to the family ranking algorithm that are novel to the practice of matching in child adoptions.

Registration Information

Through our interviews with child case workers and discussion of the simulation results, PAE managers came to recognize new potential for gathering information during the registration process as a driver of the overall network adoption rate. Specifically, our research collaboration has led to the collection of additional information to use in making match recommendations. Although PAE managers believed revising the CY 130 and CY 131 forms to be an arduous process, especially since the forms had been recently revised, they are beginning to collect child and family information through an online survey format with a

new set of questions. Their intuition about data that would be most valuable for predicting matches informed their selection of these new questions. In particular, these new questions focus on the child's positive attributes, such as interests or hobbies, that might predict attraction between children and families. Other questions focus on family attributes that could be compared to child or child case worker preferences for families without certain pets or other children in certain age ranges. The questions have received approval from the state for use by PAE and are in the process of being implemented.

Furthermore, as a result of our project, PAE managers have begun tracking the results of match recommendations for future analysis. While the set of information used for match recommendations is currently based on managerial intuition, we have encouraged PAE to maintain data about match attempts and their results in order to scientifically evaluate the power of different questions to predict matches. Econometric analysis of child attributes, family preferences, and results of match attempts would allow PAE managers to better estimate the probability of success of a child-family match and assess which questions are more or less important in predicting a match.

Spreadsheet Matching Tool

After several design and feedback iterations, PAE regional coordinators have begun to use our matching tool prototype to suggest families to county-level case workers. The matching tool prototype has also allowed them to gain insights into the matching rules and to begin to think about what matching rules help produce the best matches for children. We have supported the addition of features, such as geographical preferences for families, to improve the tool's value to PAE managers.

Compared to other matching tools discussed in Hanna and McRoy (2011), we use the same underlying framework of linearly weighted questions to score a family's suitability for a child, and add three simple innovations. First, the user can directly specify the weights for each attribute to help determine which attributes are most important for selecting a family for a child. PAE managers and regional coordinators observed that children are "labelled" by certain special needs (e.g., fire starter, animal abuser) when the underlying behavior that prompted the label was viewed as innocuous. They felt having the ability to adjust matching tools weights based on knowledge of the severity of a child's special needs could produce better matches. This feature also allowed the easy change of default weights for factors identified as important in the regression analysis. Second, the user can

state geographical preferences for the family's county of residence, which can be important in assessing the feasibility of a match if continuing community or familial relationships is important for children. Finally, as noted in Hanna and McRoy (2011), social workers who assist families can use the tool to identify shortcomings in the family's capabilities for a child and prepare appropriate support mechanisms. To this end, we have added score summaries by category so that users can more easily identify a family's strengths and challenges, as well as output reports for use by the matching committee that show how the child and family compare for each attribute.

Information Incentives

We also examined how families and child case workers interact with PAE to understand intentional or unintentional behavior patterns that may reduce the matching process's effectiveness. Through conversations with PAE managers and preliminary runs of the matching system, we noticed that the system is vulnerable to strategic manipulation by families in how they complete their CY 131 forms. Because rejecting a child is very easy for families — a telephone hotline is available to review details of and accept/reject an available child for which the family is recommended — families have an incentive to overstate their willingness to accept children with special needs. This allows a family to gain additional information and be considered for more children. Furthermore, this behavior does not necessarily result from conscious manipulation of the PAE system; different families may be inherently more or less strict in how they interpret the difference between an "acceptable" and "will consider" response to a specific special need. The current system gives families the incentive to err on the side of choosing "accept," which makes differentiating between families more difficult for PAE.

We recommended a process change and an algorithmic feature to overcome the challenge of families' overstating their tolerance for children with special needs. First, we recommended that matching occur in small batches so that PAE regional coordinators — who use the matching tool — can observe if families are chosen too frequently and further investigate the appropriateness of those families as recommended matches. PAE managers initially welcomed this suggestion, and decided that monthly matching meetings would work well with the adoption framework. They also augmented it based on their experiences with a rule that PAE coordinators should wait 30 days between successive recommendations of the same family. Second, the family-child match score was adjusted for three

criteria — race, age, and gender — to reward families whose preferences more closely fit the child’s attributes. For example, a family who indicates a preference of male or female receives a higher score on the gender attribute than a family who indicates a preference of “either” if the child is of the preferred gender.

Using a matchmaking experiment, we show that the rewarding of narrow preferences more effectively spreads the recommendations over the pool of families. For each active child, we calculated the top five matches (plus ties) from the list of active families using a scoresheet with and without rewards for narrow preferences over age, race, and gender. Without rewards for narrow preferences, we noticed that only 7.7% of families received at least one match, although we should note that this number would increase in practice when geographical preferences are applied. With rewards for narrow preferences, the number of families who received at least one match increased by 41%. We proposed further rewarding narrow preferences based on the number of child special needs that the family states are acceptable, based on the acceptable number of all special needs or the ten special needs identified as significant in the regression analysis.

Another phenomenon that PAE managers felt needed to be addressed is placement decision making by child case workers. SWAN managers and even case workers themselves indicate that some case workers also struggle with emotional attachment to children to an extent that dispassionately making a placement decision according to the child’s best interests can be very difficult. In other words, some case workers’ emotional attachment can cause them to hold out for the “perfect” family, when another family likely to be suitable for the child is available. PAE managers have found the spreadsheet matching tool to be valuable as a mechanism to enforce the conceptualization of trade-offs. In conversations between case workers and PAE regional coordinators who use the matching tool, observing that no family is likely to be a 100% match can lead to discussions about the strengths and weaknesses of a family-child match.

Conclusions

In our collaboration with PAE described in this paper, we have characterized the operation of an adoption exchange that recommends prospective families for children in state custody. A survey of child case workers indicated that PAE faces significant hurdles to aiding child case workers in the matching process. To focus on the goal of reducing the number of

children who age out of the foster care system, we identified predictors of a child's outcome that informed best practices for case workers throughout Pennsylvania and the relative importance of factors used in a match assessment tool. We next used a simulation of PAE to justify the value of a statewide network and show how the adoption rate increases with the quality of information about family preferences. Finally, we worked with PAE to collect additional valuable information, intelligently rank families using a spreadsheet matching tool, and implement process changes to encourage families to give narrow preferences. We believe that these improvements increase PAE's value to case workers in the course of finding a family, and will increase the percentage of children who find permanent placements. Furthermore, we have worked with PAE to begin collecting additional data about the matching process to scientifically compare families' stated preferences to their actual decisions. This will enable future work that would analyze the matching weights and relative value of the registration questions.

The challenge of making match recommendations in a two-sided matching market are by no means confined to PAE. Many other states, multi-state partnerships, and non-profit organizations operate similar exchanges with recommendation processes. It is, then, only natural that the results of this project can be utilized by national and regional adoption exchanges that function similarly to PAE. Furthermore, the procedure for selecting foster care families as children enter the foster care system also bears similarities. Foster care systems can benefit from this research, as they face similar challenges but with additional time pressures. Given the importance of analytics in social work, the proposed empirical and computational models and solutions can have a huge impact on how children in state custody are adopted. It is a daunting yet extremely promising endeavor, as we firmly believe that this research will inspire both the operations management literature and the economics literature.

Acknowledgments

The authors wish to thank Jane Johnston, Pennsylvania Adoption Exchange Division Manager, and the

managers and staff of the Pennsylvania Statewide Adoption and Permanency Network for their extensive

support of this project.

Appendix

Matching Tool Spreadsheet

Match Scoring Tool for Ranking Families					
Child ID (Fictitious)		Family ID (Fictitious)		Points	
				Earned	Possible
				384	425
Score	90.35%				
Weight	Child Info	Family Pref	Points	Pts Possible	
DEMOGRAPHIC INFORMATION					
100	Age	13		100	100
	Low Age		10		
	High Age		14		
100	Race/Ethnicity				
	African American	NA	Not Preferred	0	0
	Hispanic	NA	Not Preferred	0	0
	White	Applicable	Preferred	100	100
	American Indian/Alaskan Native	NA	Not Preferred	0	0
	Asian	NA	Not Preferred	0	0
	Native Hawaiian/Other Pacific Islander	NA	Not Preferred	0	0
100	Child Gender	Female	Either	100	100
SPECIAL NEEDS INFORMATION					
10	Drug Exposed Infant	Applicable	Not Approved	0	10
10	Emotional Disability	Applicable	Approved	10	10
100	HIV	NA	Not Approved	0	0
10	MH Diagnosis	NA	Approved	0	0
100	MR Diagnosis	NA	Not Approved	0	0
10	Multiple Placement History	Applicable	Approved	10	10
100	Physical Disability	NA	Approved	0	0
10	Runaway History	NA	Approved	0	0
100	Sexual Abuse History	NA	Approved	0	0
100	Siblings	NA	Approved	0	0
10	Special Education Student	Applicable	Approved	10	10
100	Special Medical Care	NA	Approved	0	0
10	Abuse History	Applicable	Not Approved	0	10
10	Neglect History	Applicable	Approved	10	10

Figure 7 PAE regional coordinators use a spreadsheet with customizable attribute weights that automatically computes scores for all families for a given child. (“NA” refers to an attribute that is not applicable for a child.)

Weight	CHILD CHARACTERISTICS	Child Info	Family Pref	Points	Pts Possible
	1. Does child have significant health issues?	No	Acceptable		
1	2. Does child have allergies or asthma? (may require treatment)	Yes	Acceptable	1	1
10	3. Is child hyperactive? (may require treatment)	Yes	Acceptable	10	10
1	4. Does child have speech problems? (may require treatment)	No	Acceptable		
1	5. Does child have hearing problems? (may require treatment)	No	Acceptable		
1	6. Is child legally deaf?	No	Will Consider		
1	7. Does child have vision problems? (may require treatment)	No	Will Consider		
10	8. Is child legally blind?	No	Unacceptable		
1	9. Does child have dental problems? (may require treatment)	No	Acceptable		
1	10. Does child have orthopedic problems (special shoes, braces, etc)	No	Acceptable		
10	11. Does child have seizures?	No	Will Consider		
1	13. Is child a high achiever in school?	No	Acceptable		
1	14. Does child achieve at grade level in regular classes?	Yes	Acceptable	1	1
1	15. Does child achieve below grade level in regular classes?	No	Acceptable		
1	16. Is child in special education classes?	No	Acceptable		
1	17. Does child have a learning disability?	No	Acceptable		
1	18. Does child need classes for the emotionally or behaviorally handicapped?	Yes	Acceptable	1	1
1	19. Does child need tutoring in one or more subjects?	No	Acceptable		
10	20. Does child have serious behavior problems in school?	Yes	Will Consider	5	10
1	21. Is child generally quiet and shy?	No	Acceptable		
1	22. Is child generally outgoing and noisy?	Yes	Acceptable	1	1
1	23. Does child have emotional issues that requires therapy?	Yes	Acceptable	1	1
1	24. Does child tend to reject father figures?	No	Will Consider		
1	25. Does child tend to reject mother figures?	No	Will Consider		
1	26. Does child have difficulty relating to others and relating to other children?	Yes	Acceptable	1	1
1	27. Does child frequently wet the bed?	No	Acceptable		
1	28. Does child frequently soil him/herself?	No	Will Consider		
100	29. Does child masturbate frequently or openly?	No	Unacceptable		
1	30. Does child have poor social skills?	Yes	Acceptable	1	1
10	31. Does child have problem with lying?	Yes	Will Consider	5	10
10	32. Does child have problem with stealing?	No	Will Consider		
10	33. Does child frequently start physical fights with other children?	Yes	Will Consider	5	10
100	34. Does child abuse animals?	No	Unacceptable		
10	35. Is child destructive with clothing, toys, etc.?	Yes	Will Consider	5	10
10	36. Does child use foul or bad language?	No	Acceptable		
1	37. Does child have frequent temper tantrums?	Yes	Acceptable	1	1
1	38. Does child have difficulty accepting and obeying rules?	Yes	Acceptable	1	1
100	39. Does child exhibit inappropriate sexual behavior?	No	Will Consider		
100	40. Does child have a history of running away?	No	Will Consider		
100	41. Does child have history of playing with matches, setting fires?	No	Unacceptable		
1	42. Does child have strong ties to birth family?	Yes	Acceptable	1	1
1	43. Does child have strong ties to foster family?	No	Acceptable		
1	44. Is continued contact with siblings desirable?	No	Acceptable		
1	45. Does child have a previous adoption disruption?	No	Acceptable		
1	46. Was child sexually abused?	No	Will Consider		
1	48. Was child exposed to promiscuous sexual behavior?	No	Will Consider		
1	49. Was child conceived by rape?	No	Will Consider		
1	50. Was child conceived as a result of prostitution?	No	Unacceptable		
1	51. Are one or both parents addicted to alcohol?	Yes	Will Consider	0.5	1
1	52. Are one or both parents dependent on substances other than alcohol?	Yes	Acceptable	1	1
1	53. Do one or both parents have a criminal record?	Yes	Acceptable	1	1
1	54. Are one or both parents mentally retarded?	No	Unacceptable		
1	55. Do one or both parents have a mental illness?	Yes	Will Consider	0.5	1
1	56. Does agency lack information about one or both parents?	No	Acceptable		
1	57. Is child in contact with birth parents?	Yes	Acceptable	1	1
	58. Is child in contact with siblings?	No			
	59. Is child in contact with extended birth family?	No			
	60. Is child in contact with former foster family?	No			

Figure 8 The spreadsheet tool also includes a section for “Child Characteristics” information.

Matching Model and Simulation Details

Children are defined by a type $c = \{a, s, r\}$, which reflects that child's desirability on two attributes — age $a \in [0, 19]$ years and number of significantly negative special needs $s \in \{0, \dots, 10\}$ — and a residence region $r \in \{1, \dots, R\}$, where R is the total number of regions.

Families are defined by their type $f = \{a_{MIN}, a_{MAX}, s', r, \alpha\}$, which represents their range of acceptable ages $[a_{MIN}, a_{MAX}]$ with $a_{MIN}, a_{MAX} \in \{0, \dots, 19\}$ and $a_{MIN} < a_{MAX}$, their tolerance for child special needs $s' \in \{0, \dots, 10\}$, a weight $\alpha \in [0, 1]$ to express the relative importance of age and special needs, and a region attribute $r \in \{1, \dots, R\}$. We also define a utility function for the purpose of indicating whether a family will accept an offered child. A child age component and a child special needs component comprise the utility function, and their relative weight is dictated by the weighting term α . With only limited information about family preferences, we use a uniform distribution for the weighting term, as justified for preference modeling with limited information by Kennan (2006).

We define a family's utility for a match with a child of type c as

$$u(c; f) := \alpha (u^{AGE}(a, a_{MAX})) + (1 - \alpha) (u^{SN}(s, s')) + \epsilon, \quad (1)$$

where ϵ is an error term that represents the randomness of a child's attractiveness to a family. We let ϵ be an independent random variable that follows a normal distribution with mean μ and standard deviation σ . The term σ represents families' variability in their attractiveness for individual children. Without data to connect families' stated preferences to their acceptance decisions, we consider cases of $\sigma = 0.1$ and $\sigma = 0.2$, to which we refer as low attraction variability and high attraction variability, respectively. Given σ , μ then becomes a tuning parameter for the simulation. The value of ϵ is only revealed when a match is attempted between a family and a child.

Lacking data to directly estimate families' preferences, we instead rely on the analysis of factors related to child outcomes from the previous section to create a model for family preferences. Specifically, we use the resulting coefficients from a linear regression model based on the child's managerially weighted outcome as a response variable and factors of age (linear), age (quadratic), and the number of significant negative special needs. The resulting model is

$$Outcome(c) = 0.8356 + 0.0426a - 0.0045a^2 - 0.0476s, \quad (2)$$

for which the intercept and all three coefficients are significant at a 99.9% confidence level. We use these coefficients from this model to estimate the age and special needs components of the utility function. Given a child's age a and a family's preferred minimum age a_{MIN} and maximum age a_{MAX} , we define the age component of a family's utility for a child as

$$u^{AGE}(a, a_{MIN}, a_{MAX}) := \begin{cases} 0.0426(a - a_{MAX}) - 0.0045(a^2 - a_{MAX}^2) & \text{if } a \geq a_{MIN} \\ 0 & \text{if } a < a_{MIN} \end{cases},$$

which represents the difference in the effect of age upon outcome between the child's age upon registration and the family's maximum preferred age. For a family that prefers a child between 0 and 12 years old, the age component of the utility function is 0.235 for a child of age 4, 0.190 for age 8, and 0 for age 12. For older

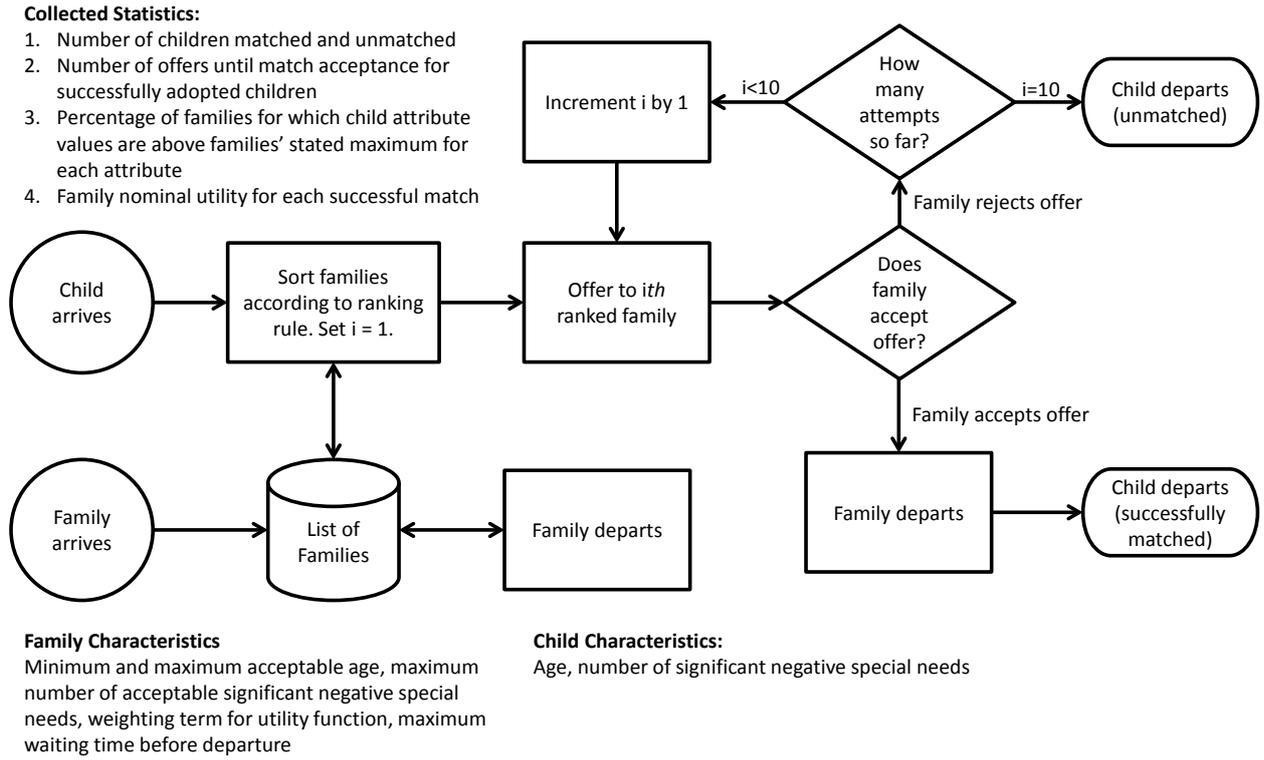


Figure 9 When a child becomes available, we rank prospective families and sequentially make up to ten match attempts. A child is successfully adopted if at least one family accepts the child. Otherwise, the child is not adopted.

children, the value is -0.149 for a child of age 14, -0.334 for age 16, and -0.554 for age 18. We note that, due to the quadratic term, the utility component is not strictly decreasing in age for very young children, but we ignore this effect just for simulation purposes as children younger than 3 only represent about 6% of the population. In general, these children are not difficult to place and are not the focus of PAE. Similarly, the special needs component of the family's utility for a child is defined as

$$u^{SN}(s, s') := -0.0476(s - s'),$$

which represents the difference in the effect of the number of significant negative special needs and the family's number of corresponding acceptable special needs.

With this model, we can more precisely define the three family ranking methods:

1. *Critical Attribute (CA)*: If $0.0426a - 0.0045a^2 < -0.0476s$, then families are sorted according to $u^{AGE}(a, a_{MIN}, a_{MAX})$. Otherwise, families are sorted according to $u^{SN}(s, s')$.
2. *Unknown Weight (UW)*: Families are sorted according to their nominal utilities, which disregards the error term, for the child with the two attributes equally weighted (i.e., $\alpha = 0.5$) to represent α unknown.
3. *Full Information (FI)*: Families' types are known to the matchmaker, which given a child of type c can rank the families according to their nominal utility $\alpha(u^{AGE}(a, a_{MAX})) + (1 - \alpha)(u^{SN}(s, s'))$.

The simulation is initialized by starting with a pool of 1,000 randomly generated families. The replication length is five years and is preceded by a one-year warm-up period. The family attraction variability tuning parameter is set to $\mu = -0.1875$ for low attraction variability $\sigma = 0.1$ and $\mu = -0.315$ for high attraction variability $\sigma = 0.2$, which corresponds to a 64% success rate for the Critical Attribute rule and matches the expected quality-adjusted outcome value for children in the PAE system between 2005 and 2013. We calibrated the simulation using the CA decision rule to represent the process by which county case workers manually searched through families' records focusing on their suitability for a small subset of child attributes, which most accurately describes PAE's functioning before changes were implemented as part of our collaboration. For each scenario — defined by a matching policy and number of regions — we used 25 replications so that we are 95% confident that the resulting mean match rate is within 1% of the true mean match rate.

The simulation was implemented in Java and relied upon the Java Simulation Library described in Rossetti (2008) for simulation functions.

References

- Abdulkadiroğlu, A., T. Sönmez. 2003. School choice: A mechanism design approach. *American Economic Review* **93**(3) 729–747.
- Arikan, M., B. Ata, J. Friedewald, R. Parker. 2012. What drives the geographical differences in deceased donor organ procurement in the United States? Working paper, University of Kansas.
- Ata, B., A. Skaro, S. Tayur. 2012. OrganJet: Overcoming geographical disparities in access to deceased donor kidneys in the United States. Working paper, Northwestern University.
- Baccara, Mariagiovanna, Allan Collard-Wexler, Leonardo Felli, Leeat Yariv. 2013. Child-adoption matching: Preferences for gender and race. Forthcoming in *American Economic Journal: Applied Economics*.
- Billar, B., S. Ghosh. 2006. Multivariate input processes. B. L. Nelson, S. G. Henderson, eds., *Handbooks in Operations Research and Management Science*, vol. 13. Elsevier Science, Amsterdam, 123–153.
- Coles, Peter, John Cawley, Phillip B. Levine, Muriel Niederle, Alvin E. Roth, John J. Siegfried. 2010. The job market for new economists: A market design perspective. *The Journal of Economic Perspectives* **24**(4) 187–206.
- Fisher, A. 2003. Still not quite as good as having your own: Toward a sociology of adoption. *Annual Review of Sociology* **29** 335–361.
- Flango, V., C. Flango. 1995. How many children were adopted in 1992. *Child Welfare* **74** 967–982.
- Gale, D., L. Shapley. 1962. College admissions and the stability of marriage. *The American Mathematical Monthly* **69**(1) 9–15.
- Goodwin, M. 2010. *Baby markets: money and the new politics of creating families*. Cambridge University Press.
- Hanna, M. D., R. G. McRoy. 2011. Innovative practice approaches to matching in adoption. *Journal of Public Child Welfare* **5**(1) 45–66.
- Hitsch, G. J., A. Hortaçsu, D. Ariely. 2010. What makes you click?—mate preferences in online dating. *Quantitative Marketing and Economics* **8**(4) 393–427.
- Howard, J., S. Brazin. 2011. Never too old: Achieving permanency and sustaining connections for older youth in foster care. Tech. rep., Evan B. Donaldson Adoption Institute, New York, NY.
- Kennan, J. 2006. A note on discrete approximations of continuous distributions. Unpublished manuscript.
- Landes, E, R Posner. 1978. The economics of the baby shortage. *Journal of Legal Studies* **7** 323–348.
- Lee, S. 2007. Preferences and choice constraints in marital sorting: Evidence from Korea. Unpublished manuscript.
- Lee, S., M. Niederle, H. Kim, W. Kim. 2011. Propose with a rose? signaling in internet dating markets. Tech. rep., National Bureau of Economic Research.

- Niederle, M., A. E. Roth. 2003. Unraveling reduces mobility in a labor market: Gastroenterology with and without a centralized match. *Journal of Political Economy* **111**(6) 1342–1352.
- Rees, M. A., J. E. Kopke, R. P. Pelletier, D. L. Segev, M. E. Rutter, A. J. Fabrega, J. Rogers, O. G. Pankewycz, Janet Hiller, A. E. Roth, T. Sandholm, M. U. Ünver, R. A. Montgomery. 2009. A nonsimultaneous, extended, altruistic-donor chain. *New England Journal of Medicine* **360**(11) 1096–1101.
- Rossetti, M. D. 2008. Java simulation library (JSL): an open-source object-oriented library for discrete-event simulation in Java. *International Journal of Simulation and Process Modelling* **4**(1) 69–87.
- Roth, A. E. 2010. Market design: Understanding markets well enough to fix them when they're broken. Tech. rep., NSF White Paper.
- Roth, A. E., M. Sotomayor. 1990. *Two-sided matching: A study in game-theoretic modeling and analysis*. Cambridge University Press.
- Roth, Alvin E., Elliott Peranson. 1999. The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American Economic Review* **89** 748–780.
- Roth, Alvin E., Tayfun Sönmez, M. Utku Ünver. 2004. Kidney exchange. *Quarterly Journal of Economics* **119**(2) 457–488.
- Roth, Alvin E., Tayfun Sönmez, M. Utku Ünver. 2005. Pairwise kidney exchange. *Journal of Economic Theory* **125** 151–188.
- Roth, Alvin E., Tayfun Sönmez, M. Utku Ünver, Francis Delmonico, Susan Saidman. 2006. Utilizing list exchange and nondirected donation through ‘Chain’ paired kidney donations. *American Journal of Transplantation* **6**(11) 2694–2705.
- Rushton, A., H. Minnis. 2002. Residential and foster family care. M. Rutter, E. Taylor, eds., *Residential and foster family care..* Oxford: Blackwell, 359–372.
- Sönmez, Tayfun, M. Utku Ünver. 2013. Market design for kidney exchange. Zvika Neeman, Alvin E. Roth, Nir Vulkan, eds., *Handbook of Market Design*. Oxford University Press, 91–137.
- U.S. Census Bureau. 2003. Adopted children and stepchildren: 2000. URL <https://www.census.gov/prod/2003pubs/censr-6.pdf>.
- U.S. General Accounting Office. 1999. Foster care: Effectiveness of independent living services.
- Zax, David. 2013. Falling for the job. *Time* **181**(12) B1.