Natural Interviewing Equilibria in Matching Settings

Joanna Drummond
University of Toronto
Canada
jdrummond@cs.toronto.edu

Omer Lev
University of Toronto
Canada
omerl@cs.toronto.edu

Allan Borodin
University of Toronto
Canada
bor@cs.toronto.edu

Kate Larson
University of Waterloo
Canada
kate.larson@uwaterloo.ca

ABSTRACT

While matching markets are ubiquitous, much of the work on stable matching assumes that both sides of the market are able to fully specify their preferences. However, as the size of matching markets grows, this assumption is unrealistic, and so there is interest in understanding how agents may use *interviews* to refine their preferences over alternatives. In this paper we study a market where one side (*e.g.* hospital residency programs) maintains a common preference master list, while the other side (*e.g* residents) have individual preferences which they can refine by conducting a limited number of interviews. The question we study is *How should residents choose their interview sets, given the choices of others?* We describe a payoff function for this imperfect information game, and show that this game always has a pure strategy equilibrium. Moreover, for certain structures of residents' utility there is a unique Bayesian equilibrium in which residents interview assortatively: with k interviews, each resident group $r_{kj+1}, ..., r_{kj+k}$ interviews with hospitals $h_{kj+1}, ..., h_{kj+k}$. For Borda-based linear utility functions, this equilibrium only exists when two interviews are allowed. We show this equilibrium varies for other utility functions, including exponential, and show general results regarding when this equilibrium does and does not exist.

1. INTRODUCTION

Real world matching problems are ubiquitous and cover many domains. One of the most widely studied matching problems is the canonical *stable matching problem (SMP)* [11]. Finding a stable matching is key in many real-world matching markets including college admissions, school choice, reviewer-paper matching, various labor-market matching problems [21], and, famously, the residency matching problem, where residents are matched to hospital programs via a centralized matching program (such as the National Residency Matching Program, NRMP, in the United States) [24].

This notion of stability, where no one in the market has both the incentive and ability to change their partner, has been empirically shown to be very valuable for real-world markets. For example, centralized mechanisms that produced a stable match tended to halt unraveling in residency matching programs, while unstable mechanisms tended to be abandoned [24]. Many matching markets that produce stable matches implement the Deferred Acceptance (DA) mechanism, introduced in Gale and Shapley's seminal paper [11].

However, to guarantee stability, stable matching mechanisms assume that participants are able to rank all their options. Assuming that participants do not have any information burden or interviewing budget is simply not the case in real-world markets: for example, in the NRMP in 2015, 27,293 positions were offered by 4,012 hospital programs [22], however residents tend to apply to an average of only 11 programs, spending between \$1,000 to \$5,000 [2]. This implies that, even if resident-proposing Deferred Acceptance (RP-DA) is the mechanism used, residents must be strategic about what hospital programs they choose to interview with, as they cannot be matched to a program with which they do not interview. Furthermore, by not carefully choosing with whom to interview, residents face the possibility of not being matched at all. There is significant evidence of this happening, as an aftermarket (SOAP) exists for the NRMP; with SOAP having matched 1,129 positions to residents in 2015, or 4.14% of the initial available positions [22]. We thus wish to study *interviewing equilibria*, not stability, for matching markets.

In spite of there being many examples where it is not feasible for participants to specify full preferences over all alternatives, there has been only limited work which has addressed participants' strategic considerations (notable exceptions include [6, 5, 16]). There similarly is little work investigating how people people choose their interviews in practice, though there is some work that suggests people tend to interview *assortatively* (i.e., in tiers): the best candidates apply to the best schools/hospitals, and the worst candidates apply to the worst schools/hospitals (e.g., [1]).

In this paper, using the residency matching problem as a motivating example, we initiate a study of the equilibrium behavior of participants who must decide with whom to interview, knowing they are participating in a centralized matching market running the resident-proposing deferred acceptance algorithm. In particular, under the assumption that hospitals maintain a *master list*, a commonly known fixed ranking over all residents, and that residents

Appears at: 4th Workshop on Exploring Beyond the Worst Case in Computational Social Choice (EXPLORE 2017). Held as part of the Workshops at the 16th International Conference on Autonomous Agents and Multiagent Systems. May 8th-9th, 2017. Sao Paulo, Brazil.

can interview with at most k hospitals, we study which subset of hospitals residents will choose to interview and then rank. Many real-world matching markets use master lists; for example, university entrances in Turkey and China are determined by test scores [12, 25], as is high-school choice in Mexico City and Ghana [7, 1]. We further note that stating our problem using master lists also provides results for other problems: this problem can be re-contextualized as a serial dictatorship mechanism with known picking order [3].

We first formalize a payoff function for any resident in this game and show that a pure strategy equilibrium always exists under general conditions on the distributions and valuation functions from which residents' underlying preferences are drawn. We then turn to investigating when assortative interviewing forms an equilibrium, under various assumptions regarding residents' preferences. We instantiate residents' preferences as drawn from a ϕ-Mallows model (i.e., resident's idiosyncratic preferences are described by a noisy universal ranking). Under this setting, we provide a condition that is sufficient (though not necessary) to guarantee assortative interviewing. We further instantiate agents' valuation functions using classes of scoring rules from the social choice literature [4], for which there exists some evidence suggesting they may approximate the structure of participants preferences [17, 19]. We study the interplay between valuation-function structure, interview-budget size and assortative interviewing. For small interviewing budgets (of size 2 or 3), assortative interviewing may be an equilibrium depending on the valuation functions of residents and if the dispersion is not too large. However, for larger interviewing budgets our results indicate that for a large segment of resident preference structures, assortative interviewing is *not* an equilibrium.

2. RELATED RESEARCH

While there is a large body of research on the problem of finding stable matchings for various markets and market conditions (including when master lists are present, e.g. [13]), there has been significantly less work on the interviewing problem in which we are interested. Interviews are information-gathering activities and one research direction has looked at interviewing policies which attempt to minimize the number of interviews conducted while ensuring that a stable matching is found. Rastegari *et al.* showed that while finding the minimal interviewing policy is NP-hard in general, there are special cases where a polynomial-time algorithm exists [23]. Drummond and Boutilier looked at a similar problem, using minimax regret and heuristic approaches for interviewing policies [10] . Neither of these papers study strategic issues arising when agents get to choose with whom they wish to interview.

Motivated by the college admissions problem, Chade and co-authors have looked at how students may strategically apply to colleges, where they assume that there is an agreed-upon ranking of the colleges, but that students' quality or caliber is determined by a noisy signal [6, 5]. This work investigates how students decide where to apply in a decentralized market. We instead focus on centralized matching markets which result in stable matchings. Coles *et al.* [8] discuss signalling in matching markets. They assume that agents' preferences are distributed according to some (restricted) distributions, known *a priori*, and each agent knows their own preferences. Firms can make at most one job offer, and workers can send one *signal* to a firm indicating their interest, paralleling, in some sense, a very restricted interviewing problem. Under this setting, firms can often do better than simply offering their top candidate a job, though there are also examples where signalling may be harmful [14]. Again, the market structure in these works is quite different than the centralized matching markets we are interested in.

The work most closely related to the problem in this paper is by Lee and Schwarz [16]. They studied an interviewing game where firms and workers (or hospitals and residents) interview with each other in order to be matched. They formulate a two-stage game where firms were required to first choose workers to interview for some fixed cost. The interview action reveals both workers' and firms' preferences, which are then revealed to a market mechanism running (firm-proposing) DA. They showed that if there is no coordination then firms' best response is picking k workers at random to interview. However, if firms can coordinate then it is best for them to each select k workers so that there is perfect overlap (forming a set of disconnected complete bipartite interviewing subgraphs). This result relies heavily on the assumption that all firms and workers are *ex-ante* homogeneous, with agents' revealed preferences being idiosyncratic and independent. This assumption is very strong; for the results to hold either agents have effectively no information about their preferences before they interview, or the market must be perfectly decomposable into homogeneous sub-markets that are known before the interviewing process starts. In this paper we study a similar interviewing game, but use a different (and we believe, more realistic) set of assumptions on the structure and knowledge of preferences.

3. MODEL

There are n residents and n hospital programs. The set of residents is denoted by $R = \{r_1, ..., r_n\}$; the set of hospital programs is denoted by $H = \{h_1, ..., h_n\}$. We are interested in *one-to-one matchings* which means that residents can only do their residency at a single hospital, and that hospitals can accept at most one resident. A *matching* is a function $\mu : R \cup H \to R \cup H$, such that $\forall r \in R$, $\mu(r) \in H \cup \{r\}$, and $\forall h \in H$, $\mu(h) \in R \cup \{h\}$. If $\mu(r) = r$ or $\mu(h) = h$ then we say that r or h is unmatched. A matching μ is *stable* if there does not exist some $(r, h) \in R \times H$, such that $h \succ_r \mu(r)$ and $r \succ_h \mu(h)$.

Both hospitals and residents have (strict) preferences over each other, and we let $H_\succ$ and $R_\succ$ denote the sets of all possible preference rankings over H and R respectively. We assume that hospitals have identical preferences over all residents, which we call the *master list*, $\succ_H$. Without loss of generality, let $\succ_H = r_1 \succ r_2 \succ \ldots \succ r_n$ where $r_i \succ_H r_j$ means that r_i is preferred to r_j, according to $\succ_H$. We further assume that the master list is common knowledge to all members of H and R. That is, all hospitals agree on the preference ranking over residents and each resident knows where they, and all others, rank in the list. While each resident, r, has idiosyncratic preferences over the hospitals, we assume that these are drawn *i.i.d.* from some common distribution D, and that this is common knowledge. If resident r draws preference ranking η from D, then $h_i \succ_\eta h_j$ means that h_i is preferred to h_j by r under η. We assume there is some common scoring function $v : H \times H_\succ \mapsto \mathbb{R}$, applied to rankings η drawn from D such that, given any $\eta \in H_\succ$ with $h_i \succ_\eta h_j$, $v(h_i, \eta) > v(h_j, \eta)$.

Critical to our model is the assumption that residents do not initially know their true preferences, but can refine their knowledge by conducting a number of *interviews*, not exceeding their interviewing budget k. We let $I(r_j) \subset H$ denote the interview set of resident r_j, and $|I(r_j)| \leq k$ for some fixed $k \leq n$. Once r_j has finished interviewing, r_j knows her preference ranking over $I(r_j)$. She then submits this information to the matching algorithm, resident-proposing deferred acceptance (RP-DA). The matching proceeds in rounds, where in each round unmatched residents propose to their next favorite hospital from their interview set to whom they have not yet proposed. Each hospital chooses its favorite resident from amongst the set of residents who have just proposed and its current match, and the hospital and its choice are then tentatively matched. This process continues until everyone is matched. The resulting matching, μ, is guaranteed to be stable, resident-optimal, and hospital-pessimal [11]. This matching is also guaranteed to be unique, as stable matching problems with master lists have unique stable solutions [13]. Thus our results directly hold for any mechanism that returns a stable match, including hospital-proposing deferred acceptance, and the greedy linear-time algorithm [13].

3.1 Description of the Game

We now describe the *Interviewing with a Limited Budget* game:

1. Each resident $r \in R$ simultaneously selects an interviewing set $I(r) \subset H$, based on their knowledge of D and the hospitals' master list $\succ_H$, where $|I(r)| \leq k$.

2. Each resident, r, interviews with hospitals in $I(r)$ and discovers their preferences over members of $I(r)$.

3. Each resident reports their learned preferences over $I(r)$ and reports all other hospitals as unacceptable. Each hospital reports the master list to a centralized clearinghouse, which runs resident-proposing deferred acceptance (RP-DA), resulting in the matching μ.

3.2 Payoff function for Interviewing with a Limited Budget

Let M be the set of all matchings, and let μ denote the ex-post matching resulting from all agents playing the *Interviewing with a Limited Budget* game. In order for resident r_j to choose their interview set $I(r_j) \subset H$, she has to be able to evaluate the payoff she expects to receive from that choice, where the payoff depends on both the actual preference ranking she expects to draw from D, the interview sets of the other residents, and the expected matching achieved from the mechanism as described. Crucially, we observe that r_j need only be concerned about the interview set of resident r_i when $r_i \succ_H r_j$. If $r_j \succ_H r_i$ then, because we run RP-DA, r_j would always be matched before r_i with respect to any hospital they both had in their interview set. Thus, we can denote r_j's expected payoff for choosing interview set S by: $u_{r_j}(S) = u_{r_j}(S|D, I(r_1), ..., I(r_{j-1}))$.

Given fixed interviewing sets $I(r_1), ..., I(r_{j-1})$, and some partial match $m = \mu_{|r_1,...,r_{j-1}}$, we must compute the probability that m happened via RP-DA. Let $m(r_i)$ denote who resident r_i is matched to under m. For any r_i, there is a set of rankings consistent with r_i being matched with $m(r_i)$ under RP-DA (and the hospitals' master list $\succ_H$). Denote this set as $T(r_i, m)$. Formally, $T(r_i, m) \subseteq H_\succ$ is:

$$T(r_i, m) = \{\xi \in H_\succ | \forall h' \in H \text{ s.t. } h' \in I(r_i) \wedge h' \succ_\xi m(r_i), \\ \exists r_a \text{ s.t. } r_a \succ_H r_i \wedge m(r_a) = h'\}$$

Given the interviewing sets of residents $r_1, \ldots, r_{j-1}$, the probability of partial match m is

$$P(m|I(r_1), ..., I(r_j)) = \prod_{r_i \in \{r_1,...,r_{j-1}\}} \sum_{\xi \in T(r_i,m)} P(\xi|D). \quad (1)$$

where $P(\xi|D)$ is the probability that some resident drew ranking $\xi \in H_\succ$ from D.

Using Eq. 1, we can now determine the probability that some hospital h is matched to r_j using RP-DA, when r_j has interviewed with set S, and has preference list η. We simply sum over all possible matches in which this could happen. Because RP-DA is resident optimal, and all hospitals have a master list, any hospital that r_j both interviews with and prefers to h must already be matched. We formally define the set of such matchings, $M^*(S, \eta, I(r_1), ..., I(r_{j-1}))$:

$$M^*(S, \eta, I(r_1), ..., I(r_{j-1}), h) = \\ \{m \in M | m(r_j) = h; \forall r_i \in \{r_1, ..., r_{j-1}\} m(r_i) \in I(r_i); \\ \text{and } \forall x \in S, \text{ if } x \succ_\eta h, \exists r_i \in \{r_1, ..., r_{j-1}\} \text{ s.t. } x \in I(r_i) \text{ and } m(r_i) = x\}$$

Thus, the probability that h is matched to r_j using RP-DA given η, S, and the interviewing sets for all residents preferred to r_j on the hospitals' master list is

$$P(\mu(h) = r_j|\eta, S, I(r_1), ..., I(r_{j-1})) = \sum_{m \in M^*(S,\eta,I(r_1),...,I(r_{j-1}),h)} P(m|I(r_1), ..., I(r_{j-1})). \quad (2)$$

For readability, we will frequently refer to $P(\mu(h) = r_j|\eta, S, I(r_1, ..., I(r_{j-1})$ as $P(\mu(h) = r_j|\eta, S)$. Finally, we have all of the building blocks to formally define the payoff function. Recall that $v(h, \eta)$ is the imposed utility function, dependent on η: for any given η, $v(h, \eta)$ is fixed. Then, our payoff function is:

$$u_{r_j}(S) = \sum_{h \in S} \sum_{\eta \in H_\succ} v(h, \eta) P(\eta|D) P(\mu(h) = r_j|\eta, S, I(r_1), ..., I(r_{j-1})) \quad (3)$$

Intuitively, what the payoff function in Eq. 3 does is weight the value for some given alternative by how likely r_j is to be matched to that item, given the interview sets of the "more desirable" residents, $r_1, \ldots, r_{j-1}$.

As an illustrative example, imagine there are two residents, r_1 and r_2, each of whom have interviewed with hospitals h_1 and h_2. Resident r_1 will be matched with whomever she most prefers, while r_2 will be assigned the other. The probability that r_2 will be assigned h_1 is simply the probability that r_1 drew ranking $h_2 \succ h_1$, while the probability that r_2 is matched to h_2 is the probability that r_1 drew ranking $h_1 \succ h_2$.

3.3 Probabilistic Preference Models

While our payoff-function formulation, described in the previous section, is general in that we do not instantiate it with a particular distribution function, we do assume that some distribution is used over the space of possible rankings of hospitals. The Mallows model is characterized by a reference ranking σ, and a dispersion parameter $\phi \in (0, 1]$,[1]

[1] A ϕ-Mallows model is not well defined for $\phi = 0$, but if all residents are guaranteed to draw the reference ranking, the equilibrium is trivial.

which we denote as $D^{\phi,\sigma}$. Let A denote the set of alternatives that we are ranking, and let $A_{\succ}$ denote the set of all permutations of A (the index $i \in [1,n]$ in $a_i \in A$ indicates rank in σ). The probability of any given ranking r is:

$$P(r|D^{\phi,\sigma}) = \frac{\phi^{d(r,\sigma)}}{Z}$$

Here d is Kendall's τ distance metric, and Z is a normalizing factor; $Z = \sum_{r' \in P(A)} \phi^{d(r,\sigma)} = (1)(1+\phi)(1+\phi+\phi^2)...(1+ ... + \phi^{|A|-1})$ [18].

As $\phi \to 0$, the distribution approaches drawing the reference ranking σ with probability 1; when $\phi = 1$, this is equivalent to drawing from the uniform distribution. The Mallows model (and mixtures of Mallows) have plausible psychometric motivations and are commonly used in machine learning [20, 15, 18]. Mallows models have also been used in previous investigations of preference elicitation schemes for stable matching problems (e.g., [9, 10]).

To prove our equilibria results, we will need additional results regarding properties of Mallows models. To the best of our knowledge, the following have not been stated previously, and may be of more general interest. Proofs omitted due to space constraints

LEMMA 1. *Given some Mallows model $D^{\phi,\sigma}$ with fixed dispersion parameter ϕ and reference ranking $\sigma = a_i \succ a_j$, then the probability that a ranking η is drawn from $D^{\phi,\sigma}$ such that $a_i \succ_\eta a_j$ is equal to drawing from some distribution $D^{\phi,\sigma'}$ where σ is a prefix of σ'. By symmetry, this proof also holds when σ is a suffix of σ'.*

COROLLARY 2. *Given any reference ranking σ and two alternatives a_i, a_{i+1}, $P(a_i \succ a_{i+1}|D^{\phi,\sigma}) = \frac{1}{1+\phi}$.*

COROLLARY 3. *Given any reference ranking σ and three alternatives a_i, a_{i+1}, a_{i+2} and some $\eta \in \{a_i, a_{i+1}, a_{i+2}\}_{\succ}$, then the probability that we draw some ranking β consistent with η is: $P(\beta|D^{\phi,\sigma}) = \frac{\phi^{d(\eta, a_i \succ a_{i+1} \succ a_{i+2})}}{(1+\phi)(1+\phi+\phi^2)}$.*

LEMMA 4. *The probability a_i will be ranked in place j is $\frac{\phi^{|j-i|}}{1+\phi+...+\phi^{n-1}}$.*

LEMMA 5. *Let $\eta \in D^{\phi,\sigma}$ in which $a_j \succ_\eta a_i$ for $i < j$, then $P(\eta) < \frac{\phi^{j-i}}{Z}$.*

4. GENERAL EQUILIBRIA FOR INTERVIEWING MARKETS WITH MASTER LISTS

We provide an equilibrium analysis for the game presented in Section 3. We first show that a pure equilibria for this game always exists, even under arbitrary distributions and scoring functions, but may be computationally infeasible to directly calculate. We then instantiate this model for various distributions and scoring functions, focusing on one family of distributions: the ϕ-Mallows model. We provide a necessary and sufficient condition for *assortative* interviewing under a Mallows model and then investigate what values of ϕ and k will result in assortative interviewing for various scoring functions.

4.1 General Equilibria for Interviewing Markets with Master Lists

We start our analysis by studying the most general form of the *Interviewing with a Limited Budget game*, and show that a pure strategy equilibrium always exists.

THEOREM 6. *A pure strategy always exists for the Interviewing with a Limited Budget game.*

PROOF. We wish to show that if every resident chooses their expected utility maximizing interviewing set, this forms a pure strategy. Given any resident r_j who is jth in the hospitals' rank ordered list, r_j's expected payoff function only depends on residents $r_1, ..., r_{j-1}$. As r_j knows that each other resident r_i is drawing from distribution D *i.i.d.*, she can calculate $r_1, ..., r_{j-1}$'s expected utility maximizing interview set, using Eq. 3. Her payoff function depends only on D and $I(r_1), ..., I(r_{j-1})$, both of which she now has. She then calculates the expected payoff for each $\binom{n}{k}$ potential interviewing sets, and interviews with the one that maximizes her expected utility. □

We note that Theorem 6 is an existence theorem and does not provide any additional insight into the equilibrium behavior, nor does it provide guidance as to how such an equilibrium might be computed. Our next result begins to provide some intuition as to equilibrium behavior. In particular it shows that if residents have interviewing budgets of size k and the equilibrium behavior for resident r_k is to interview assortatively (*i.e.* it chooses to interview with hospitals $h_1, \ldots, h_k$), then assortative interviewing is the equilibrium strategy for all residents.

PROPOSITION 7. *Given an interviewing budget of k interviews, some known distribution from which residents draw their preferences D and a scoring function v, if resident r_k's best response to all others interview assortatively is to interview assortatively, then assortative interviewing is an equilibrium for all residents.*

(Proof omitted due to space constraints)

4.2 Interviewing Equilibria Under Mallows Models with Master Lists

In this section we instantiate the distribution from which residents are drawing their preferences with a Mallows model in order to gain a deeper understanding of the results from the previous section. In particular, we provide a characterization of when assortative interviewing will form an equilibrium for this class of resident-preferences. Before proving our main result, we require some additional lemmas addressing characteristics of assortative interviewing in Mallows models.

All proofs are omitted due to space constraints.

LEMMA 8. *Given an interviewing budget of k interviews, a dispersion parameter ϕ, and a scoring function v, if resident r_k prefers interviewing with hospitals $\{h_1, ..., h_k\}$ to $\{h_1, ..., h_{k+1}\} \setminus \{h_j\}$ for all $h_j \in \{h_1, ..., h_k\}$, then for resident r_k, interviewing with $\{h_1, ..., h_k\}$ dominates interviewing with any other set of size k.*

We now provide a necessary and sufficient condition for assortative interviewing to hold when residents draw their preference from a Mallows model with dispersion ϕ. Let $P(h_i \text{ avail})$ denote the probability that hospital h_i is available for resident r_k (i.e., residents $r_1, ..., r_{k-1}$ are all matched to different alternatives). As we assume residents $r_1, ..., r_{k-1}$ interview assortatively, only one of $\{h_1, ..., h_k\}$ will be available.

Lemma 9. *Given an interviewing budget of k interviews, a dispersion parameter ϕ, and a scoring function v, if residents $r_1, ..., r_{k-1}$ all interview assortatively (i.e., with hospital set $S = \{h_1, ..., h_k\}$), satisfying the following inequality for all $h_j \in \{h_1, ..., h_k\}$ when $S' = S \setminus \{h_j\} \cup \{h_{k+1}\}$ is both sufficient and necessary to show that assortative interviewing is an equilibrium for resident r_k:*

$$P(h_j\ avail)\mathbb{E}(v(h_j)|D^{\phi,\sigma}) \geq$$
$$P(h_j\ avail)\mathbb{E}(v(h_{k+1})|D^{\phi,\sigma}) + \tag{4}$$
$$\sum_{\eta \in H_{\succ}} P(\eta|D^{\phi,\sigma}) \cdot [\sum_{h_i \in S'} P(h_i\ avail)\chi(h_{k+1} \succ_\eta h_i)v(h_{k+1},\eta)]$$

Where $\chi(h_i \succ_\eta h_j)$ is an indicator function that is 1 iff $h_i \succ_\eta h_j$, and 0 otherwise.

Theorem 10. *Given an interviewing budget of k interviews, a dispersion parameter ϕ, and a scoring function v, satisfying the inequality found in Lemma 9 for all $h_j \in \{h_1, ..., h_k\}$ is both sufficient and necessary to show that assortative interviewing is an equilibrium for all residents.*

Proof. This follows directly from combining Proposition 7 and Lemma 9. □

We now provide a more simplified condition for assortative interviewing, that is sufficient, though not necessary (and leave the proof to the appendix):

Lemma 11. *Given an interviewing budget of k interviews, a dispersion parameter ϕ, and a scoring function v, if residents $r_1, ..., r_{k-1}$ all interview assortatively (i.e., with hospital set $S = \{h_1, ..., h_k\}$), satisfying the following inequality for all $h_j \in \{h_1, ..., h_k\}$ when $S' = S \setminus \{h_j\} \cup \{h_{k+1}\}$ is sufficient to show that assortative interviewing is an equilibrium for resident r_k:*

$$P(h_j\ avail)\mathbb{E}(v(h_j)|D^{\phi,\sigma}) \geq$$
$$P(h_j\ avail)\mathbb{E}(v(h_{k+1})|D^{\phi,\sigma}) + \sum_{h_i \in S'} P(h_i\ avail)\mathbb{E}(v(h'_k)|D^{\phi,\sigma'})\frac{\phi}{Z(1-\phi)} \tag{5}$$

(Where σ' is equivalent to the reference ranking σ with one element h_i s.t. $h_j \succ_\sigma h_i$ removed, and h'_k is the kth item in σ'.)

Though we primarily discuss assortative interviewing as it is a technique commonly used in real-world interviewing markets, we note that the n/k complete disjoint bipartite subgraph equilibrium shown in Lee and Schwarz for uniform distributions on both sides of the market also holds when one side is drawing uniform iid (equivalently, a Mallows model with $\phi = 1.0$), and the other side has a master list.

Observation 12. *When residents draw iid from uniform, and hospitals have a master list, an equilibrium exists such that the interviewing graph forms n/k complete disjoint bipartite subgraphs. Moreover, any resident r_{ik+j} interviews with hospitals $\{h_{(j-1)k+1}, \ldots, h_{jk}\}$.*

5. ASSORTATIVE EQUILIBRIA FOR SMALL BUDGETS

We now discuss assortative equilibria when participants' interviewing budget is $k \leq 3$. We do so by instantiating specific scoring rules, and investigating under what circumstances assortative interviewing forms an equilibrium. We now formally define Borda, plurality, and exponential scoring rules, following definitions typically used in voting. We define all scoring rules with a multiplicative factor of 1, and an additive factor of 0, as these terms do not affect the analysis. For any slot s_i, $v(s_i) = n - i + 1$ in Borda, where n is the number of alternatives in the market. Under plurality, $v(s_1) = 1$, $v(s_i) = 0$ for all $i > 1$. We investigate a class of exponential functions that are dominated by the function $v(s_i) = (\frac{\varepsilon}{2})^{i-1}$, $1 > \varepsilon > 0$.

The proofs for the following two lemmas are omitted due to space constraints

Lemma 13. *If for a particular interviewer budget k, a dispersion parameter ϕ, the condition of Lemma 14 is satisfied for a plurality valuation function with a strict inequality, then there are exponential valuations which form an assortative equilibrium.*

In particular, any exponential valuation dominated by $(\frac{\varepsilon}{2})^{(i-1)}$ satisfies this condition, with $\varepsilon > 0$ determined by k.

Lemma 14. *A necessary and sufficient condition for assortative interviewing under plurality is:*

$$P(h_j\ avail) \geq \phi^{k-j+1} \tag{6}$$

This follows from instantiating plurality into Eq. 6, applying Lemmas 7 and 4, and simplifying.

5.1 Assortative Interviewing with Two Interviews

We provide direct proofs showing that assortative interviewing is an equilibrium for Borda and plurality. Exponential follows directly from Lemma 13.

Theorem 15. *Given plurality as residents' scoring function and a budget of $k = 2$ interviews, for a Mallows model with dispersion parameter ϕ such that $0 < \phi \leq 0.6180$, assortative interviewing forms an equilibrium.*

Proof. We begin by using the condition from Lemma 14. We provide the calculation for h_1; h_2 follows analogously (providing a bound of $0 < \phi \leq 0.7549$). We thus wish to show conditions on ϕ s.t. $P(h_1 \text{ avail}) \geq \phi^2$,when resident r_2 is choosing their interview set. For r_2, h_1 is available iff r_1 happened to draw a ranking over her preferences s.t. $h_2 \succ h_1$. Then, by Corollary 2, $P(h_1 \text{ avail}) = \frac{\phi}{1+\phi}$, implying we need to satisfy the equation $\frac{\phi}{1+\phi} \geq \phi^2$, which is true whenever $0 < \phi \leq 0.6180$. □

Theorem 16. *Given Borda as residents' scoring function and a budget of $k = 2$ interviews, for a Mallows model dispersion parameter ϕ such that $0 < \phi \leq 0.2650$, assortative interviewing forms an equilibrium.*

Proof. Because of Lemma 7, we only need to show that assortative interviewing is an equilibrium when $0 < \phi \leq 0.265074$ for resident r_2, and it will hold for all r_i. Furthermore, by Lemma 8, we only need to prove that $\{h_1, h_2\}$ dominates both $\{h_1, h_3\}$ and $\{h_2, h_3\}$ to show that it dominates all other possible interviewing sets of size 2.

We prove that choosing $\{h_1, h_2\}$ is better than choosing $\{h_2, h_3\}$, for all values of ϕ such that $0 < \phi \leq 0.265074$. We prove this by summing over all possible preference rankings that induce a specific permutation of the alternatives h_1, h_2, h_3. We then pair these summed permutations in such a manner that makes it easy to find a lower bound for $u_{r_2}(\{h_1,h_2\}) - u_{r_2}(\{h_2,h_3\})$. This lower bound is entirely in terms of ϕ, meaning that for any ϕ such that this bound is above 0, it will be above 0 for any market size n.

We look at three cases, pairing all possible permutations of h_1, h_2, h_3 as follows:
Case 1: all rankings η consistent with $h_2 \succ h_1 \succ h_3$ or η' consistent with $h_2 \succ h_3 \succ h_1$;
Case 2: all rankings η consistent with $h_1 \succ h_2 \succ h_3$ or η' consistent with $h_3 \succ h_2 \succ h_1$;
Case 3: all rankings η consistent with $h_1 \succ h_3 \succ h_2$ or η' consistent with $h_3 \succ h_1 \succ h_2$.

Note that as we have enumerated all possible permutations of h_1, h_2, h_3, these three cases generate every ranking in $H_\succ$. Furthermore, for any one of the three cases, we can iterate over only all possible rankings η that are consistent with the first member of the pair, and generate the ranking η' consistent with the second member of the pair by simply swapping two alternatives in the rank. Moreover, given some η, the number of discordant pairs in η' is simply the number in η, plus the number of additional discordant pairs between h_1, h_2, h_3 caused by swapping the two alternatives.

For clarity, let $u_{r_2}(\{h_1,h_2\}) - u_{r_2}(\{h_2,h_3\}) = U_1+U_2+U_3$, where U_1, U_2, U_3 correspond to our three cases. We also introduce the notation $P_{\mu(r_i)}(h)$ to denote the probability that r_i is matched to hospital h under matching μ. That is, $P_{\mu(r_i)}(h) = P(\mu(r_i) = h)$. The case proofs proofs are omitted due to space constraints

Once considering all cases, we combine them together:

$$\begin{aligned}
&u_{r_2}(\{h_1,h_2\}) - u_{r_2}(\{h_2,h_3\}) \geq \\
&\frac{\phi^2}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}(1-\phi) + \frac{2(\phi-\phi^3-\phi^4)}{(1+\phi)(1+\phi)(1+\phi+\phi^2)} \\
&- \frac{\phi}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}\left(\frac{\phi}{(1-\phi)^4} + \frac{1}{3(1-\phi)^3} + \frac{2}{3}\right)(1+\phi) \\
&+ \frac{\phi^2}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}(1-\phi) \qquad (7)
\end{aligned}$$

Thus, Eq. 7 gives us a lower bound for the difference in expected utility between $\{h_1,h_2\}$ and $\{h_2,h_3\}$ for resident r_2, for all n. Using numerical methods to approximate the roots of Eq. 7, we get that there is a root at 0, and a root at $\phi \approx 0.265074$.

As the calculations are analogous, we omit the discussion of their derivation, but it can be shown that:

$$\begin{aligned}
&u_{r_2}(\{h_1,h_2\}) - u_{r_2}(\{h_1,h_3\}) \geq \\
&\frac{1}{(1+\phi)(1+\phi+\phi^2)}\left[1+\phi-2\phi^2-2\phi^3-2\phi^3\left(\frac{\phi}{(1-\phi)^4}+\right.\right. \\
&\left.\left.\frac{1}{3(1-\phi)^3}+\frac{2}{3}\right)\right] \qquad (8)
\end{aligned}$$

Using numerical methods, it can be shown that this is positive for $0 < \phi < 0.413633$.

Thus, for the interval $0 < \phi \leq 0.265074$, we have shown that r_2's best move in this interval is to interview with $\{h_1,h_2\}$. Then, by Lemma 7, this is an equilibrium for all r_i as required. □

5.2 Assortative Interviewing with Three Interviews

Unlike when only two interviews are present, assortative interviewing is *not* an equilibrium under Borda when participants have a budget of 3 interviews. Under plurality (and exponential), assortative interviewing is still an equilibrium.

THEOREM 17. *Assortative interviewing is not guaranteed to be an equilibrium under the Borda valuation function, even for any ϕ.*

(Proof omitted due to space constraints)

Under Borda, an assortative interviewing equilibrium is not guaranteed to exist, even for any $1 \geq \phi > 0$. However, we now show that assortative interviewing is an equilibrium for plurality (and thus exponential) for $k = 3$:

THEOREM 18. *Given an interviewing budget of $k = 3$ interviews, and the plurality scoring function, assortative interviewing is an equilibrium for $0 < \phi \leq 0.4655$.*

PROOF. For $k = 3$, we simply check Eq. 6 from Lemma 14 with $h_j = h_1, h_2, h_3$. We find that the marginal contribution from h_1 is less than the marginal contribution of h_2 or h_3, and thus only present the calculation for h_1. We directly compute $P(h_1 \text{ avail})$, by multiplying the probability that r_1 did not take h_1, and multiplying it by the probability that r_2 did not take h_1, given that r_1 *also* did not take h_1. To calculate this we enumerate the probabilities of any possible rankings:

$$P(h_1 \text{ avail}) = P(\mu(r_1) \neq h_1)P(\mu(r_2) \neq h_1 | \mu(r_1) \neq h_1)$$

$$P(h_1 \text{ avail}) = \left(\frac{\phi+2\phi^2+\phi^3}{(1+\phi)(1+\phi+\phi^2)}\right)\left(\frac{\phi^2+2\phi^3}{(1+\phi+\phi^2)}\right)$$

Using numerical methods to find the roots of $P(h_1 \text{ avail}) - \phi^3$, we can show that Eq. 6 holds when $0 < \phi \leq 0.4655$. □

6. ASSORTATIVE EQUILIBRIA FOR LARGE BUDGETS

We begin by providing a few final results regarding properties of interviewing under Mallows models, including that when there is a setting for which there is no assortative equilibria for plurality, then there is no valuation function with assortative equilibria. We use this result to show that, for sufficiently small ϕ and a large enough budget of interviews ($k > 3$), assortative interviewing cannot be an equilibrium under any valuation function. We then provide a specific counterexample for *all* ϕ when $k = 4$ for plurality, implying there is no assortative equilibrium for any valuation function. This suggests that, for a wide category of resident valuation functions under Mallows, contrary to real-world behavior, assortative interviewing is not an equilibrium.

LEMMA 19. *Given a Mallows model with dispersion parameter ϕ, assortative interviewing for residents $r_1, ..., r_{k-1}$, and a hospital $h_i \in \{h_1, ..., h_k\}$ (i.e., the residents' interview set), then any profile $\eta_1, \ldots, \eta_{n-1} \in D^{\phi,\sigma}$ of $k-1$ preferences (for $r_1, \ldots, r_{k-1}$) such that h_i is available for r_k has a probability of: $P(r_1 = \eta_1, r_2 = \eta_2, \ldots, r_{k-1} = \eta_{k-1} | h_i \text{ avail}) < \frac{\phi^\gamma}{Z^{k-1}}$, where $\gamma = \sum_{j=1}^{k-i} j$.*

PROOF. In order for h_i to be available, there need to be $r'_{i+1}, \ldots, r'_k$ with preference orders $\eta_{i+1}, \ldots, \eta_k \in D^{\phi,\sigma}$ such that they were assigned hospitals $h_{i+1}, \ldots, h_k$. Hence, $h_{i+1} \succ_{\eta_{i+1}} h_i, \ldots, h_k \succ_{\eta_k} h_i$. According to Lemma 5, the probability for each of these events is at least $\frac{\phi}{Z}, \ldots, \frac{\phi^{k-i}}{Z}$ (respectively). Since they are independent of each other, and since the maximal probability for any particular $\eta \in D^{\phi,\sigma}$ is $\frac{1}{Z}$, the probability of a particular preference set occurring in which h_i is available is at least $\frac{\phi^\gamma}{Z^{k-1}}$. □

THEOREM 20. *If for a particular interviewer budget k, a dispersion parameter ϕ, when using plurality valuation there are no assortative equilibria due to h_1 violating Lemma 9's condition, then for that k and ϕ there is no assortative equilibria for any valuation function.*

PROOF. Looking at the condition of Lemma 9

$$P(h_j \text{ avail})\mathbb{E}(v(h_j)|D^{\phi,\sigma}) \geq$$
$$P(h_j \text{ avail})\mathbb{E}(v(h_{k+1})|D^{\phi,\sigma})+$$
$$\sum_{\eta\in H_\succ} P(\eta|D^{\phi,\sigma})\big[\sum_{h_i\in S'} P(h_i \text{ avail})\chi(h_{k+1}\succ_\eta h_i)v(h_{k+1},\eta)\big]$$

We again begin by expanding the value expectation ($\mathbb{E}$)This can be divided to n different inequalities:

$$P(h_j \text{ avail})P(h_j \text{ in } s_1)v(s_1) \geq v(s_1)[P(h_j \text{ avail})P(h_{k+1} \text{ in } s_1)+$$
$$\sum_{\substack{\eta\in H_\succ| \\ h_{k+1} \text{ in } s_1}} P(\eta|D^{\phi,\sigma}) \sum_{h_i\in S'} P(h_i \text{ avail})\chi(h_{k+1}\succ_\eta h_i)]$$

$\vdots$

$$P(h_j \text{ avail})P(h_j \text{ in } s_{n-1})v(s_{n-1}) \geq$$
$$v(s_{n-1})[P(h_j \text{ avail})P(h_{k+1} \text{ in } s_{n-1})+$$
$$\sum_{\substack{\eta\in H_\succ| \\ h_{k+1} \text{ in } s_{n-1}}} P(\eta|D^{\phi,\sigma}) \sum_{h_i\in S'} P(h_i \text{ avail})\chi(h_{k+1}\succ_\eta h_i)]$$
$$P(h_j \text{ avail})P(h_j \text{ in } s_n)v(s_n) \geq v(s_n)P(h_j \text{ avail})P(h_{k+1} \text{ in } s_n)$$

We shall show that under the theorem's assumptions, none of these inequalities hold for h_1, and therefore the general inequality (Lemma 9) does not hold.

Note that for each inequality we can simply ignore $v(s_\ell)$ ($1 \leq \ell \leq n$), since they appear on both sides of the inequality. The assumption of theorem is that first inequality does not hold, i.e.,

$$P(h_1 \text{ avail})P(h_1 \text{ in } s_1) < P(h_1 \text{ avail})P(h_{k+1} \text{ in } s_1)+$$
$$\sum_{\substack{\eta\in H_\succ| \\ h_{k+1} \text{ in } s_1}} P(\eta|D^{\phi,\sigma}) \sum_{h_i\in S'} P(h_i \text{ avail})\chi(h_{k+1}\succ_\eta h_i)$$

As shown in Lemma 4, for any $1 < \ell \leq k$ the probability of h_1 being in any spot s_ℓ is monotonically decreasing with ℓ, while the probability of h_{k+1} being in spot s_ℓ is monotonically increasing with ℓ. Hence, $P(h_1 \text{ avail})P(h_1 \text{ in } s_1) > P(h_1 \text{ avail})P(h_1 \text{ in } s_\ell)$. Similarly, $P(h_1 \text{ avail})P(h_{k+1} \text{ in } s_1) < P(h_1 \text{ avail})P(h_{k+1} \text{ in } s_\ell)$. We analogously see that:

$$\sum_{\substack{\eta\in H_\succ| \\ h_{k+1} \text{ in } s_1}} P(\eta|D^{\phi,\sigma}) \sum_{h_i\in S'} P(h_i \text{ avail})\chi(h_{k+1}\succ_\eta h_i) <$$
$$\sum_{\substack{\eta\in H_\succ| \\ h_{k+1} \text{ in } s_\ell}} P(\eta|D^{\phi,\sigma}) \sum_{h_i\in S'} P(h_i \text{ avail})\chi(h_{k+1}\succ_\eta h_i)$$

Simply put, the LHS gets smaller, while the RHS increases. Hence, for $1 \leq \ell \leq k$:

$$P(h_1 \text{ avail})P(h_1 \text{ in } s_\ell) < P(h_1 \text{ avail})P(h_{k+1} \text{ in } s_\ell)+$$
$$\sum_{\substack{\eta\in H_\succ| \\ h_{k+1} \text{ in } s_\ell}} P(\eta|D^{\phi,\sigma}) \sum_{h_i\in S'} P(h_i \text{ avail})\chi(h_{k+1}\succ_\eta h_i)$$

By Lemma 4, for any $\ell > k$, $P(h_1 \text{ in } s_\ell) < P(h_{k+1} \text{ in } s_\ell)$ which gives us:

$$P(h_1 \text{ avail})P(h_1 \text{ in } s_\ell) < P(h_1 \text{ avail})P(h_{k+1} \text{ in } s_\ell) \implies$$
$$P(h_1 \text{ avail})P(h_1 \text{ in } s_\ell) < P(h_1 \text{ avail})P(h_{k+1} \text{ in } s_\ell)+$$
$$\sum_{\substack{\eta\in H_\succ| \\ h_{k+1} \text{ in } s_\ell}} P(\eta|D^{\phi,\sigma}) \sum_{h_i\in S'} P(h_i \text{ avail})\chi(h_{k+1}\succ_\eta h_i)$$

Starting with the assumption that assortative interviewing does not hold for plurality, we show that none of the inequalities above hold for any slot s_ℓ, and therefore that the condition in Lemma 9 does not hold for $j = h_1$ for any valuation function. □

THEOREM 21. *Given an interviewing budget of $k > 3$ interviews, there exists $0 < \varepsilon < 1$ s.t. for any scoring function v no assortative interviewing forms an equilibrium for dispersion parameter $\phi < \varepsilon$.*

PROOF. Thanks to Theorem 20, it is enough for us to shown there is no assortative equilibrium under plurality (and that h_1 violates Lemma 9's condition). We again begin with the simplification from Lemma 14: $P(h_j \text{ avail}) \geq \phi^{k-j+1}$. Thanks to Lemma 19, we know $P(h_j \text{ avail})$ is of the form:

$$P(h_j \text{ avail}) = \frac{X(k)}{Z^{k-1}}\phi^{\sum_{i=1}^{k-j} i} + \frac{X^1(k)}{Z^{k-1}}\phi^{1+\sum_{i=1}^{k-j} i} + \ldots + \frac{X^\ell(k)}{Z^{k-1}}\phi^{(k\sum_{i=1}^{k-j} i)-1} + \frac{1}{Z^{k-1}}\phi^{k\sum_{i=1}^{k-j} i} \quad (9)$$

($X(k), X^1(k), \ldots, X^\ell(k)$ are functions that calculate the number of different sets of possible preference orders for $r_1, \ldots, r_k$, with each set being of probability $\phi^{\sum_{i=1}^{k-j} i}$ for $X(k)$, $\phi^{1+\sum_{i=1}^{k-j} i}$ for $X^1(k)$, etc.)

When $\phi \to 0$, $Z^{k-1} \to 1$, and Equations 9 becomes $P(h_j \text{ avail}) \to X(k)\phi^{\sum_{i=1}^{k-j} i}$. In particular, there is ε', such that $P(h_1 \text{ avail}) < X(k)\phi^{(\sum_{i=1}^{k-j} i)-1}$, and there is $\varepsilon = \min(\varepsilon', \frac{1}{X(k)})$ such that for $\phi < \varepsilon$, for $k > 3$:

$$\phi^k \geq \phi^{(\sum_{i=1}^{k-j} i)-2} > X(k)\phi^{(\sum_{i=1}^{k-j} i)-1} > P(h_1 \text{ avail})$$

Contradicting our condition (Equation 6). □

Moreover, we show that for $k = 4$, assortative interviewing is *never* an equilibrium.

THEOREM 22. *Given an interviewing budget of $k = 4$ interviews and any scoring function, assortative interviewing is not an equilibrium for any dispersion parameter ϕ.*

PROOF. We begin by instantiating the plurality valuation function. By Theorem 20, if assortative interviewing is not an equilibrium for plurality, it is never an equilibrium for any scoring rule. As noted before Eq. 6 is tight, so if we compute the marginal contribution from some $h^* \in \{h_1, h_2, h_3, h_4\}$, and the contribution from h^* is strictly less than the contribution from h_5 for any ϕ, assortative interviewing is not an equilibrium for $k = 4$ and plurality. We find that the contribution from h_1 is less than the marginal contribution from h_4.

To calculate $P(h_1 \text{ avail})$, we simply iterate over all 6 possible allocations for r_1, r_2, r_3 such that h_1 is not taken, and directly calculate the probabilities of each ranking profile for r_1, r_2, r_3 that allows that to happen. In the interest of clarity, we only provide a symbolic representation. Let A be the set of all permutations of h_2, h_3, h_4, so that $(a_1, a_2, a_3) \in A$.

$$P(h_1 \text{ avail}) = \sum_{(a_1,a_2,a_3)\in A} P(\mu(r_1) = a_1)P(\mu(r_2) = a_2|\mu(r_1) = a_1) P(\mu(r_3) = a_3|\mu(r_1) = a_1, \mu(r_2) = a_2)$$

We instantiate the above equation using the probabilities of each potential match, and use numerical methods to show the function $P(h_1 \text{ avail}) - \phi^4$ is negative for any ϕ in $0 < \phi \leq 1$. □

7. DISCUSSION

We investigate equilibria for interviewing (for example, between residents and hospitals) with a limited budget when a master ranked list (say, of residents) is known. We provide a generic payoff function, that is indifferent to participants' interviewing budgets, preference distributions, and scoring functions. We show that a pure strategy interviewing equilibrium always exists.

We then focus on this game for different scoring rules (Borda, plurality and exponential scoring rules), when residents' preferences are independently drawn from the same Mallows distribution. We find evidence that, for all scoring rules investigated, interviewing budgets typically seen in real-world markets do not admit assortative interviewing equilibria, even though this is a strategy frequently played in these markets. We do find that this is an equilibrium strategy for small interviewing budgets, when residents' preferences are sufficiently "similar" (i.e., low dispersion). Moreover, this assortative equilibrium strategy is a naturally arising equilibrium in which the maximum number of residents are matched; namely, the residents interview assortatively in tiers, forming a bipartite graph interviewing graph structure with n/k disconnected complete components. A similar bipartite graph interviewing structure is present in the work of Lee and Schwarz [16]. However, this structure naturally arises in our model, and we characterize a very different preference space than the Lee and Schwarz paper, which investigates the impartial culture model (i.e., a Mallows model with $\phi = 1$, or uniform distribution). We also provide an equilibrium for impartial culture in markets with master lists.

We hypothesize the difference in behavior seen in real-world markets and the equilibria shown here could result from a variety of factors. First, we only investigate assortative interviewing under a ϕ-Mallows model. As discussed in Section 3.3, while under some circumstances the Mallows model is viewed as a realistic model, it is possible participants' preferences in these markets are not sufficiently described by such a model. Another critical modeling assumption in this work is that of master lists, though assortative interviewing behavior is seen both in matching markets with and without master lists.

We also assume perfectly rational actors. Both prospect theory and quantal response equilibria could explain the difference in real-world behavior and the equilibria shown here. Individuals tend to misjudge probabilities, overestimating small probabilities, and underestimating near-certainties. (Perhaps leading them to believe assortative interviewing is a best response).

We hypothesize that, like in decentralized matching markets, the structure of the interviewing equilibria will contain both "reach" and "safety" schools, where participants diversify their interviewing portfolio to get both the benefit of a desirable, unlikely option, and a likely, but less desirable option. We find some evidence of this equilibrium in small markets with Borda valuations. Figure 1, depicts a market with 4 hospitals, 4 residents, and 2 interviews ($n = 4$, $k = 2$) and shows the explicit trade-off between high-value unlikely alternatives, and more choice over alternatives. The figure shows the exact payoff for each interviewing set, for given dispersion ϕ. As ϕ increases, we explicitly see the trade-off between more choice, and a better expected payoff value for individual alternatives. For sufficiently large ϕ, choice dominates individual expectations so that for r_2, interviewing $\{h_3, h_4\}$ dominates interviewing with any other set. For small ϕ, interviewing with $\{h_1, h_2\}$ dominates interviewing with any other set.Interestingly, for $\phi \in [0.5, 0.6]$, r_2's best option is to split the difference, and interview with one hospital (h_3) he is guaranteed to get and one hospital (h_2) that will be available with sufficiently high probability, and has a higher expected value. This choice available to r_2 results in some of the "reach" behavior we see in college admissions markets; r_3's best response now is to interview with h_1, h_4 (i.e., a "reach" choice, and a "safe" bet). We hypothesize that this "reach" and "safety" behavior is not only present for larger ϕ in markets with small interviewing budgets, but also in markets with large interviewing budgets.

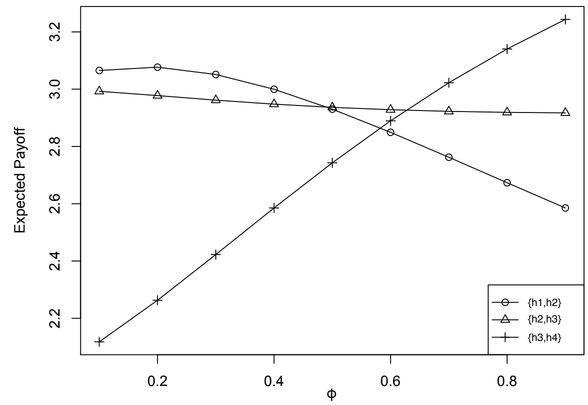

Figure 1: r_2's expected payoff for interviewing with various interviewing sets, as ϕ goes from 0 to 1, $n = 4$.

We hypothesize that results similar to the ones presented in this paper hold for different scoring functions and preference distributions (e.g., Plackett-Luce). Furthermore, the results presented here only investigate one-to-one matching markets. We believe that most of our results will directly hold for many-to-one markets where each hospital h has known capacity q_h. Another interesting future direction would be to relax the assumption that interviewing with any hospital has identical cost. In this regard, we wish to investigate equilibria when each resident has a known budget k, and each resident r has some known cost $c_r(h)$ for interviewing with hospital h; residents must then choose an interviewing set S s.t. $\sum_{h \in S} c_r(h) \leq k$. Perhaps the most important direction for future work is relaxing the master list assumption; we hypothesize that similar equilibria arise if preferences on both sides of the market are distributed according to a Mallows model with low dispersion. We also believe this work could lead to interesting questions in mechanism design, where the mechanism is a joint interviewing/matching mechanism, with a limited budget for interviews explicitly incorporated into the mechanism.

Acknowledgements

This work was supported in part by NSERC Discovery Grant RGPIN 7631.

REFERENCES

[1] K. F. Ajayi. School choice and educational mobility: Lessons from secondary school applications in Ghana. *University of California-Berkeley Working Paper*, 2011.

[2] K. D. Anderson, M. M. Dorough, C. R. Stein, S. A. Optenberg, I. M. Thompson, and others. The urology residency matching program in practice. *The Journal of Urology*, 163(6):1878–1887, 2000.

[3] S. Bade. Serial dictatorship: The unique optimal allocation rule when information is endogenous. *Theoretical Economics*, 10(2):385–410, May 2015.

[4] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

[5] H. Chade, G. Lewis, and L. Smith. Student portfolios and the college admissions problem. *The Review of Economic Studies*, 81(3):971–1002, 2014.

[6] H. Chade and L. Smith. Simultaneous search. *Econometrica*, 74(5):1293–1307, 2006.

[7] L. Chen and J. S. Pereyra. Self-selection in school choice: Theory and evidence from Mexico City high school match, 2015. Working paper.

[8] P. Coles, A. Kushnir, and M. Niederle. Preference signaling in matching markets. Technical report, National Bureau of Economic Research, 2010.

[9] J. Drummond and C. Boutilier. Elicitation and approximately stable matching with partial preferences. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI-13), Beijing*, 2013.

[10] J. Drummond and C. Boutilier. Preference elicitation and interview minimization in stable matchings. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI-14)*, pages 645–653, 2014.

[11] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69(1):9–15, 1962.

[12] I. E. Hafalir. Stability of marriage with externalities. *International Journal of Game Theory*, 37(3):353–369, 2008.

[13] R. W. Irving, D. F. Manlove, and S. Scott. The stable matching problem with master preference lists. *Discrete Applied Mathematics*, 156:2959–2977, 2008.

[14] A. Kushnir. Harmful signaling in matching markets. *Games and Economic Behavior*, 80:209–218, July 2013.

[15] G. Lebanon and Y. Mao. Non-parametric modeling of partially ranked data. *Journal of Machine Learning Research*, 9:2401–2429, 2008.

[16] R. S. Lee and M. Schwarz. Interviewing in two-sided matching markets. Working Paper 14922, National Bureau of Economic Research, Apr. 2009.

[17] G. F. Loewenstein, L. Thompson, and M. H. Bazerman. Social utility and decision making in interpersonal contexts. *Journal of Personality and Social Psychology*, 57(3):426–441, 1989.

[18] T. Lu and C. Boutilier. Learning Mallows models with pairwise preferences. In *Proceedings ICML 2011*, pages 145–152, Bellevue, WA, 2011.

[19] D. M. Messick and K. P. Sentis. Estimating social and nonsocial utility functions from ordinal data. *European Journal of Social Psychology*, 15:389–399, 1985.

[20] T. B. Murphy and D. Martin. Mixtures of distance-based models for ranking data. *Computational Statistics and Data Analysis*, 41:645–655, January 2003.

[21] M. Niederle, A. E. Roth, and T. Sonmez. Matching and market design. In S. N. Durlauf and L. E. Blume, editors, *The New Palgrave Dictionary of Economics (2nd Ed.)*, volume 5, pages 436–445. Palgrave Macmillan, Cambridge, 2008.

[22] N. R. M. Program. National resident matching program, results and data: 2015 main residency match®, 2015.

[23] B. Rastegari, A. Condon, N. Immorlica, and K. Leyton-Brown. Two-sided matching with partial information. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, pages 733–750. ACM, 2013.

[24] A. E. Roth. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70(4):1341–1378, July 2002.

[25] M. Zhu. College admissions in China: A mechanism design perspective. *China Economic Review*, 30:618–631, Sept. 2014.