Journal of Artificial Intelligence Research x (2017) x-x Submitted 5/17; published x/17

Natural Interviewing Equilibria in Matching Settings

Joanna Drummond JDRUMMOND@CS.TORONTO.EDU
Omer Lev OMERL@CS.TORONTO.EDU
Allan Borodin BOR@CS.TORONTO.EDU
University of Toronto, Department of Computer Science, 10 King's College Road, Toronto, ON M5S 3G4 Canada

Kate Larson KATE.LARSON@UWATERLOO.CA
University of Waterloo, Cheriton School of Computer Science, Waterloo, ON N2L 3G1 Canada

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 grow, 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 drawn from some distribution. We view the refinement of preferences as a probabilistic process that instantiates the prior beliefs about the ranking of hospitals. 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. We then focus on the interviewing problem when the preferences of residents are chosen from a Mallows model. Some empirical observations of matching with a master list (*e.g.* students considering high schools in Mexico) show that the side choosing the interviews will interview assortatively; that is, with k interviews, each resident group $r_{kj+1}, ..., r_{kj+k}$ interviews with hospitals $h_{kj+1}, ..., h_{kj+k}$. For $k = 2$ and $k = 3$, we characterize when such assortative interviewing results in a unique Bayesian equilibrium. Surprisingly, and contrary to some empirical observations as to how some real-world participants behave in matching markets, for $k = 4$, assortative interviewing is not a Bayesian equilibrium. We conjecture that the same (non equilibrium) is true for all $k > 3$ and show that this indeed holds for a sufficient small Mallows dispersion parameter $\phi > 0$.

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)* (Gale & Shapley, 1962). In this setting, we seek to find a one-to-one matching between two sets such that no two agents (each one from a different set) would prefer to be matched with each other, rather than their assigned matching. 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 (Niederle, Roth, & Sonmez, 2008), 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) (Roth, 2002).

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 (Roth, 2002). Many matching markets that produce stable matches implement the Deferred Acceptance (DA) mechanism, introduced in Gale and Shapley's seminal paper (Gale & Shapley, 1962).

However, in many actual matching markets, constraints on allowable outcomes or assumptions on the participants' underlying preferences need to be explicitly considered, and there has been growing interest in developing and using AI and multiagent systems techniques in matching markets. Such techniques are used for a variety of reasons; for example, to compactly represent preferences (*e.g.* (Gelain, Pini, Rossi, Venable, & Walsh, 2009; Pini, Rossi, & Venable, 2014)), to handle partial preferences (*e.g.* (Drummond & Boutilier, 2014; Rastegari, Condon, Immorlica, & Leyton-Brown, 2013)), to model and reason about quotas imposed on matching outcomes (Goto, Iwasaki, Kawasaki, Kurata, Yasuda, & Yokoo, 2016), and to consider distributional constraints (Kurata, Goto, Iwasaki, & Yokoo, 2017). For example, to guarantee stability, stable matching mechanisms assume that participants are able to rank as many options as they wish. However, stable matching mechanisms are frequently used in markets where the information burden placed upon participants may be quite severe. 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 (National Resident Matching Program, 2015), but residents tend to apply to an average of only 11 programs, spending between \$1,000 to \$5,000 (Anderson et al., 2000). 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 some 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 (National Resident Matching Program, 2015). We thus wish to study *interviewing equilibria* 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 Chade & Smith, 2006; Chade, Lewis, & Smith, 2014; Lee & Schwarz, 2009). Similarly, there 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.* Ajayi, 2011).

In this paper, using the residency matching problem as a motivating example, we initiate a study of the equilibrium behaviour 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 (*e.g.* according to grade-point average), and that residents can interview with at most k hospitals, we study which subset of hospitals residents will choose to interview and then rank. Following convention, we assume that it is the residents who are using interviews to refine their preferences, but note that our choice of which side has the master list (and which side chooses their interviews) is arbitrary and does not change our findings. Many real-world matching markets use master lists; for example, university entrances in Turkey and China are determined by test scores (Hafalir, 2008; Zhu, 2014), as is high-school choice in Mexico City and Ghana (Chen & Pereyra, 2015; Ajayi, 2011).[1] We will show that this *interview game* is a complex game in terms of understanding its equilibrium behaviour and more specifically the size k of the interview budget will significantly impact the nature of the resulting equilibria.

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 as a noisy universal ranking). Under this setting, we provide a condition that is necessary and sufficient to guarantee assortative interviewing. We further instantiate agents' valuation functions using classes of scoring rules from the social choice literature (Brandt et al., 2016), for which there exists some evidence suggesting they may approximate the structure of participants preferences (Loewenstein et al., 1989; Messick & Sentis, 1985). 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 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.* (Irving, Manlove, & Scott, 2008)), 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 (Rastegari et al., 2013). They also provide a model for minimal interviewing, and an MDP framework for minimal interviewing (with no fixed budget). Drummond and Boutilier looked at a similar problem, using minimax

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 (Bade, 2015).

regret and heuristic approaches for interviewing policies (Drummond & Boutilier, 2014). 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 (Chade & Smith, 2006; Chade et al., 2014). 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., 2010) 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 (Kushnir, 2013). 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 & Schwarz, 2009). 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 $H = \{h_1, ..., h_n\}$. 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 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 $r \succ_h \mu(h)$.

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 favourite hospital from their interview set to whom they have not yet proposed. Each hospital chooses its favourite 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 (Gale & Shapley, 1962). This matching is also guaranteed to be unique, as stable matching problems with master lists have unique stable solutions (Irving et al., 2008). Thus our results directly hold for any mechanism that returns a stable match, including hospital-proposing deferred acceptance, and the greedy linear-time algorithm (Irving et al., 2008).

3.1 Description of the Game

We now describe the *Interviewing with a Limited Budget* game. We attempt to formalize this game in a manner consistent with previous literature on interviewing, particularly with Rastegari et al. (2013):

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 matching $m = \mu_{|r_1,...,r_{j-1}}$, 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

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

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 share 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 (partial) matchings, $M^*(S, \eta, I(r_1), ..., I(r_{j-1}))$:

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

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

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

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:

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

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, just described, is general in that it can be instantiated using any scoring function and distribution over rankings, in this paper we are interested both in general results (*e.g.* utility function and distribution independent) and results under particular assumptions on both the scoring function classes and ranking distributions. In this section we introduce the preference ranking distribution we use, the ϕ-Mallows model, and discuss some of its properties.

The ϕ-Mallows model (or just Mallows model) is characterized by a reference ranking σ, and a dispersion parameter $\phi \in (0, 1]$,[2] 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; $Z = \sum_{r' \in A_\succ} \phi^{d(r,\sigma)} = (1)(1 + \phi)(1 + \phi + \phi^2)...(1 + ... + \phi^{|A|-1})$ (Lu & Boutilier, 2011).

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 (Murphy & Martin, 2003; Lebanon & Mao, 2008; Lu & Boutilier, 2011). Mallows models have also been used in previous investigations of preference elicitation schemes for stable matching problems as in (Drummond & Boutilier, 2013, 2014).

For certain results later in the paper, we rely on a number of properties of Mallows models. To the best of our knowledge, the following have not been stated previously, and may be of more general interest. The full proofs of these results can be found in the Appendix.

Intuitively, a Mallows model can be iteratively generated by repeated insertion, where the random insertion is weighted according to the dispersion parameter. Because of this, when comparing a small subset of elements in the whole ranking, the probability that any two given alternatives are in a specific order may not depend on the total number of alternatives. We discuss this more formally in the following lemmas and corollaries. Additionally, this repeated insertion procedure can be used to determine the probability

2. 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.

any given alternative will be placed in a certain slot in any given ranking: we simply look at the probability it gets inserted in that particular slot, after all other alternatives have been inserted.

We first observe that adding more alternatives to the beginning or end of a reference ranking does not change the probability of drawing two alternatives in a given order.

Lemma 1. *Given some Mallows model $D^{\phi,\sigma}$ with a fixed dispersion parameter ϕ and reference ranking σ in which $a_i \succ a_j$, 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 suffix or prefix of σ'.*

It is useful to instantiate the previous result to the case where two alternatives are adjacent to each other in the original ranking, a_i and a_{i+1}.

Corollary 2. *Given any reference ranking σ and two alternatives a_i, a_{i+1},*

$$P(a_i \succ a_{i+1} | D^{\phi,\sigma}) = \frac{1}{1+\phi}.$$

We similarly extend the previous corollary to include three consecutive items.

Corollary 3. *Given any reference ranking σ and alternatives a_i, a_{i+1}, a_{i+2} and some $\eta \in \{a_i, a_{i+1}, a_{i+2}\}_\succ$, the probability that some ranking β is drawn from $D^{\phi,\sigma}$ that is consistent with η is:*

$$P(\beta | D^{\phi,\sigma}) = \frac{\phi^{d(\eta, a_i \succ a_{i+1} \succ a_{i+2})}}{(1+\phi)(1+\phi+\phi^2)}$$

It is useful to know the probability that any one alternative will be in any particular position in a rank ordered list. We show that this is effectively equivalent to ordering all other alternatives, and then calculating the probability that we can put the alternative in question in its desired slot.

Lemma 4. *The probability that a_1 will be ranked in place j is $\frac{\phi^{j-1}}{1+\phi+\ldots+\phi^{n-1}}$. Equivalently, the probability that a_n will be ranked in place j is $\frac{\phi^{n-j}}{1+\phi+\ldots+\phi^{n-1}}$. Similarly, the probability a_j will be ranked in first place is $\frac{\phi^{j-1}}{1+\phi+\ldots+\phi^{n-1}}$.*

We also find it useful to bound the probability that any two alternatives will be "out of order" in any given ranking; this will later be used in some of our impossibility results.

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

Finally, we include an observation that follows from the Mallows' model definition:

Observation 6. *If $|j-i| > |j-i'|$, probability a_i is in place j is smaller than probability $a_{i'}$ is in place j. Similarly, probability a_j is in place i is smaller than probability a_j is in place i'.*

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 strategy equilibrium 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 7. *A pure strategy equilibrium 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. □

Note that this game has a very sequential nature: each resident's best response only depends on the $j-1$ agents that are ordered before her in the hospitals' master list. Thus, a large portion of the strategy space can be eliminated, as the behaviour of residents r_{j+1} to r_n do not affect r_j's payoff. We then continue solving for the best strategy by using iterated deletion of dominated strategies; r_1's best response is always to interview with $I(r_1) = \{h_1, \ldots, h_k\}$; this eliminates many dominated strategies for r_2, which in turn eliminates dominated strategies for r_3, and so on. Moreover, when there are no ties between the payoffs for interviewing with various sets for any given resident, one unique strategy per player will remain, thus resulting in a unique equilibrium.

We note that Theorem 7 is an existence theorem and does not provide any additional insight into the equilibrium behaviour, nor does it provide guidance as to how such an equilibrium might be computed.

We are interested in understanding whether and when a particular class of natural interviewing strategies form an equilibrium. In particular, if residents have interviewing budgets of size k, we ask the question *Will residents interview assortatively?*

Definition 1. *We say that an interviewing strategy profile is* assortative *iff, when residents have a budget of k interviews, each resident $r \in \{r_{jk+1}, \ldots, r_{jk+k}\}$ chooses to interview with the set of k hospitals $\{h_{jk+1}, \ldots, h_{jk+k}\}$.*

We now show that if assortative interviewing is a best response for resident r_k when all other residents interview assortatively, assortative interviewing is a best response for *every* resident r_i when all other residents interview assortatively. In other words, determining if assortative interviewing is a best response for r_k is sufficient to show that assortative interviewing is a best response for all residents (and is thus an equilibrium for this game).

Proposition 8. *Consider an interviewing budget of k interviews, some known distribution D from which all residents draw their preferences, a scoring function v, and a strategy profile for residents $r_1, \ldots, r_{k-1}$ such that they all interview assortatively. Then, if resident r_k's best response is to interview assortatively under this setting, it is a best response for any resident r_i to interview assortatively. Moreover, this then forms a unique equilibrium for this game in this setting.*

Proof. We introduce an indicator function to simplify notation for when a hospital is a resident's top available choice. For any hospital h and agent i, let $b^i(h, \eta) = 1$ iff h is available when r_i makes her choice, and is her most-desirable available alternative (*i.e.* $h \succ_\eta h_j$ for all other h_j available); and 0 otherwise. Directly following from the utility function, the utility of resident r_i when interviewing with hospitals $S \subset H$ can thus be written as:

$$u_{r_i}(S) = \sum_{h \in S} \sum_{\eta \in H_\succ} v(h, \eta) P(\eta, D) b^i(h, \eta)$$

We first discuss some resident r_i s.t. $i < k$ (*i.e.* resident r_i is more desirable than resident r_k). Because D can be described by a master list with noise, $S = \{h_1, \ldots, h_k\}$ is either equivalent to or dominates any other interviewing set for r_1 (as all alternatives are available to her); thus r_1 interviews with S. Suppose assortative interviewing is not an equilibrium; let r_i be the most desirable resident for which he prefers interviewing with $S' \neq S$, and suppose $r_i \succ_\sigma r_k$ (*i.e.* $i < k$). If all residents except r_i interview assortatively, $b^i(h, \eta) \geq b^k(h, \eta)$, with the inequality strict for $h \in S$. Hence, if $u_{r_i}(S) < u_{r_i}(S')$, then it follows that $u_{r_k}(S) < u_{r_k}(S')$: a contradiction. Thus residents $r_1, ..., r_{k-1}$ interview assortatively if r_k interviews assortatively.

Now, suppose that we have some resident r_j s.t. $k < j \leq 2k$ (*i.e.* resident r_j is less desirable than resident r_k). We now know that all residents $r_1, ..., r_k$ have interviewed assortatively. This implies that every hospital $h_1, ..., h_k$ are completely unavailable. This allows us to do an inductive argument: remove $r_1, ..., r_k$ and $h_1, ..., h_k$ from the market, and map r_j to its equivalent resident in $r_1, ..., r_k$. Thus, as shown above, r_j has incentive to interview assortatively, as required.

Note that, as all players have a strictly dominant strategy, this is a unique equilibrium for this game. □

4.2 Interviewing Equilibria Under Mallows Models with Master Lists

In the previous section we proved the existence of a pure strategy equilibria for the interviewing game, and provided some evidence of the existence of an assortative interviewing equilibrium under very general conditions. 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, without imposing any particular additional restrictions on the utility functions of the residents.

Before proving our main result in Theorem 11, we provide some observations and lemmas addressing characteristics of assortative interviewing in Mallows models. We first consider the situation where all residents draw the reference ranking, σ, with probability 1.[3] Any strategy profile such that each resident r_i interviews with hospital h_i is an equilibrium in this case. Thus, trivially, assortative interviewing forms an equilibrium.

For ease of notation, let $\Psi = \langle k, \phi, v \rangle$ be an instance of the Interviewing with a Limited Budget game with budget k, a Mallows model with dispersion parameter ϕ, and a scoring function v. We then show that if, for resident r_k, replacing any alternative $h_j \in \{h_1, \dots, h_k\}$ with alternative h_{k+1} is *not* an improvement to her expected utility, then interviewing with $\{h_1, \dots, h_k\}$ is her best response when she draws her preferences from a Mallows model. This allows us to greatly simplify the analysis: we must only investigate k possible interviewing sets, instead of $\binom{n}{k}$ possible interviewing sets to determine if assortative interviewing is the best strategy for r_k.

Lemma 9. *Given an Interviewing with a Limited Budget game $\Psi = \langle k, \phi, v \rangle$, if resident r_k's expected payoff from interviewing with hospitals $\{h_1, \dots, h_k\}$ is higher than her expected payoff from interviewing with hospitals $\{h_1, \dots, h_{k+1}\} \setminus \{h_j\}$ for all $j \in \{h_1, \dots, h_k\}$, then resident r_k's best response is to interview with $\{h_1, \dots, h_k\}$ (*i.e.* assortatively).*

Proof. Following the proof in Proposition 8, we use an indicator function to simplify when a hospital is a resident's top available choice. For any hospital h, let $b(h, \eta) = 1$ iff h is available for r_k, and $h \succ_\eta h_j$ for all other h_j available; and 0 otherwise. Directly following from the utility function, the utility of resident r_k when interviewing with hospitals $S = \{h_1, \dots, h_k\}$ can thus be written as:

$$u_{r_k}(S) = \sum_{h \in S} \sum_{\eta \in H_\succ} v(h, \eta) P(\eta, D^{\phi,\sigma}) b(h, \eta)$$

As we assume knowledge of the strategies for residents $r_1, ..., r_{k-1}$, we can calculate the probability that any given hospital is available. We thus can calculate the contribution of each hospital interview to the total utility, as $P(\eta, D^{\phi,\sigma}), v(h, \eta)$ are known a priori. Moreover, when $r_1, ..., r_k$ all interview with the same k hospitals, $b(h, \eta)$ is equivalent to $P(h \text{ avail})$: resident r_k gets whatever hospital $r_1, ..., r_{k-1}$ do not take.

Now, assume there exists some set S' of hospitals such that $u_{r_k}(S') > u_{r_k}(S)$. Define $\bar{S} = S \setminus S'$; denote the members of $\bar{S}$ as $h'_1, \dots, h'_l$. Also, note that h_{k+1} must be in $S' \setminus S$, as $\bar{S} \neq \emptyset$ by hypothesis, and h_{k+1} dominates all alternatives in $\{h_{k+1}, \dots, h_n\}$: h_{k+1} is available for r_k with probability 1 (as is all other alternatives not in S), and has higher expected value than any other h_j s.t. $h_{k+1} \succ_\sigma h_j$. Without loss of generality, let h'_1 be the hospital in $\bar{S}$ that minimizes the benefit gained from swapping some element in $\bar{S}$ with one of the more "desirable" elements in S'. More formally, h'_1 is the hospital in $\bar{S}$ that minimizes

$$y_1 = \sum_{\eta \in H_\succ} P(\eta | D^{\phi,\sigma}) b(h'_1, \eta) \big[v(h'_1, \eta) - v(h_{k+1}, \eta) \big]$$

3. We note that even though the Mallows model is not defined at $\phi = 0$, as $\phi \to 0$, the probability of drawing the reference ranking σ goes to 1.

y_1 is the value that is lost when h'_1 is the only available hospital from $h_1, \ldots, h_k$, and h_{k+1} must be chosen instead. The value added by choosing h_{k+1} instead of h'_1 is formally: $z_1 = \sum_{\eta \in H_\succ} P(\eta | D^{\phi,\sigma}) b(h_{k+1}, \eta) v(h_{k+1})$. Then, $u_{r_k}(S \cup \{h_{k+1}\} \setminus \{h'_1\}) = u_{r_k}(S) - y_1 + z_1$. If $y_1 \leq z_1$, the lemma is proven; for contradiction, suppose $z_1 - y_1 < 0$.

Without loss of generality, let h'_2 be the hospital in $\bar{S} \setminus \{h'_1\}$ that minimizes

$$\begin{aligned} y_2 &= \sum_{\eta \in H_\succ} P(\eta | D^{\phi,\sigma}) b(h'_2, \eta) \big[v(h'_2, \eta) - \max(v(h_{k+1}, \eta), v(h_{k+2}, \eta)) \big] \\ &= \sum_{\eta \in H_{\succ | h_{k+1} \succ h_{k+2}}} P(\eta | D^{\phi,\sigma}) b(h'_2, \eta) \big[v(h'_2, \eta) - v(h_{k+1}, \eta)) \big] \\ &\quad + \sum_{\eta \in H_{\succ | h_{k+2} \succ h_{k+1}}} P(\eta | D^{\phi,\sigma}) b(h'_2, \eta) \big[v(h'_2, \eta) - v(h_{k+2}, \eta)) \big] \end{aligned}$$

Again, y_2 is the benefit we get from h'_2, the alternative we are swapping out for h_{k+2}. The value added from h_{k+2} is $z_2 = \sum_{\eta \in H_\succ} P(\eta | D^{\phi,\sigma}) b(h_{k+2}, \eta) v(h_{k+2})$. Since h_{k+1} and h_{k+2} have the same probability of being available, but the expected value of $v(h_{k+1})$ is more than that of $v(h_{k+2})$, we know $z_2 < z_1$. Thanks to Corollary 2:

$$\sum_{\eta \in H_{\succ | h_{k+1} \succ h_{k+2}}} P(\eta | D^{\phi,\sigma}) b(h'_2, \eta) \big[v(h'_2, \eta) - v(h_{k+1}, \eta)) \big] = \frac{1}{1+\phi} y_2$$

Looking at the equivalent section of y_1:

$$\sum_{\eta \in H_{\succ | h_{k+1} \succ h_{k+2}}} P(\eta | D^{\phi,\sigma}) b(h'_1, \eta) \big[v(h'_1, \eta) - v(h_{k+1}, \eta) \big] > \frac{1}{1+\phi} y_1$$

but thanks to y_1 minimality:

$$\begin{aligned} &\sum_{\eta \in H_{\succ | h_{k+1} \succ h_{k+2}}} P(\eta | D^{\phi,\sigma}) b(h'_2, \eta) \big[v(h'_2, \eta) - v(h_{k+1}, \eta) \big] \\ &\quad > \sum_{\eta \in H_{\succ | h_{k+1} \succ h_{k+2}}} P(\eta | D^{\phi,\sigma}) b(h'_1, \eta) \big[v(h'_1, \eta) - v(h_{k+1}, \eta) \big] \end{aligned}$$

and therefore $y_2 > y_1$. Thus:

$$u_{r_k}(S \setminus \{h'_1, h'_2\} \cup \{h_{k+1}, h_{k+2}\}) = u_{r_k}(S) - y_1 + z_1 - y_2 + z_2 < u_{r_k}(S) - 2y_1 + 2z_1 < u_{r_k}(S)$$

Note that again, all other alternatives in $S \setminus S'$ must also have y_i such that $y_i > y_1$ and $z_i < z_1$, by the construction of y_1 and z_1. Let $l = |\bar{S}|$. Thus:

$$u_{r_k}(S') = u_{r_k}(S \setminus \bar{S}) + \sum_{i=1}^{l} z_i - y_i < u_{r_k}(S) - l y_1 + l z_1 < u_{r_k}(S)$$

This contradicts our assumption that $u_{r_k}(S') > u_{r_k}(S)$; thus, if such an S' exists, $y_1 \geq z_1$, and showing that S dominates $S \setminus \{h_j\} \cup \{h_{k+1}\}$ is sufficient for all $h_j \in S$. □

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 10. *Given an Interviewing with a Limited Budget game $\Psi = \langle k, \phi, v \rangle$, if residents $r_1, \ldots, r_{k-1}$ all interview assortatively (*i.e. *with hospital set $S = \{h_1, ..., h_k\}$), then assortative interviewing is a best response for resident r_k if and only if the following equality is satisfied for all $h_j \in \{h_1, ..., h_k\}$ when $S' = S \setminus \{h_j\} \cup \{h_{k+1}\}$:*

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

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.

Proof. By Lemma 9, showing that the marginal contribution from h_j is bigger than the marginal contribution from h_{k+1} is sufficient to show that S dominates any other interviewing set. Again, $S = \{h_1, ..., h_k\}$ and $S' = S \setminus \{h_j\} \cup \{h_{k+1}\}$. Using the payoff function in Section 3.2, this means that we want to find conditions such that the utility to r_k provided by h_j is larger than that of h_{k+1}:

$$\sum_{\eta \in H_\succ} v(h_j, \eta)P(\mu(h_j) = r_k|S, \eta, D^{\phi,\sigma})P(\eta|D^{\phi,\sigma}) \geq$$
$$\sum_{\eta \in H_\succ} v(h_{k+1}, \eta)P(\mu(h_{k+1}) = r_k|S', \eta, D^{\phi,\sigma})P(\eta|D^{\phi,\sigma}) \tag{4}$$

Note that, when interviewing with set S, the probability $\mu(h_j) = r_k$ is simply the probability that no resident in $r_1, ..., r_{k-1}$ chooses h_j. Thus, the left hand side of Eq. 4 simplifies to:

$$\sum_{\eta \in H_\succ} v(h_j, \eta)P(\mu(h_j) = r_k|S, \eta, D^{\phi,\sigma}P(\eta|D^{\phi,\sigma}) = P(h_j \text{ avail}) \sum_{\eta \in H_\succ} v(h_j, \eta)P(\eta|D^{\phi,\sigma})$$
$$= P(h_j \text{ avail})\mathbb{E}(v(h_j)|D^{\phi,\sigma}) \tag{5}$$

We then also wish to simplify the right hand side. Note that there are two cases in which resident r_k is matched with h_{k+1} when interviewing with set S': either h_j is the only hospital available (*i.e.* $r_1, ..., r_{k-1}$ have all been matched with $\{h_1, ..., h_k\} \setminus \{h_j\}$), or for some $h_i \in \{h_1, ..., h_k\} \setminus \{h_j\}$, h_i is available and under the ranking η in consideration, $h_{k+1} \succ_\eta h_i$. Again, $\chi(y)$ denote an indicator function, where $\chi(y) = 1$ iff y is true, and 0 otherwise. More formally, we express the RHS of the condition in Eq. 4 using the indicator function, and simplify:

$$\sum_{\eta \in H_\succ} P(\eta|D^{\phi,\sigma}) \cdot \big[v(h_{k+1}, \eta)P(h_j \text{ avail}) + \sum_{h_i \in S'} P(h_i \text{ avail})\chi(h_{k+1} \succ_\eta h_i)v(h_{k+1}, \eta)\big]$$
$$= P(h_j \text{ avail})\mathbb{E}(v(h_{k+1})|D^{\phi,\sigma})$$
$$+ \sum_{\eta \in H_\succ} P(\eta|D^{\phi,\sigma}) \cdot \big[\sum_{h_i \in S'} P(h_i \text{ avail})\chi(h_{k+1} \succ_\eta h_i)v(h_{k+1}, \eta)\big] \tag{6}$$

Combining the simplifications provided in Eqs. 5 and 6 completes the proof.

□

By combining the lemmas, we can show that we must only check k interviewing sets for resident r_k to prove that assortative interviewing forms an equilibrium for this game (i.e., it is each resident's best response to interview assortatively).

Theorem 11. *Given an Interviewing with a Limited Budget game $\Psi = \langle k, \phi, v\rangle$, satisfying the inequality found in Lemma 10 for all $h_j \in \{h_1, ..., h_k\}$ is both sufficient and necessary to show that all residents interviewing assortatively forms an equilibrium for this game.*

Proof. This follows directly from combining Proposition 8 and Lemma 10. □

We provide a more simplified condition for assortative interviewing, that is sufficient, though not necessary. This condition is easier to compute than the condition in Lemma 10,and thus may be valuable when verifying whether specific valuation functions admit assortative interviewing equilibria.

Lemma 12. *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 a best response for resident r_k:*

$$P(h_j \ avail)\mathbb{E}(v(h_j)|D^{\phi,\sigma}) \geq P(h_j \ avail)\mathbb{E}(v(h_{k+1})|D^{\phi,\sigma}) + \sum_{h_i \in S'} P(h_i \ avail)\mathbb{E}(v(h'_k)|D^{\phi,\sigma'})\frac{\phi}{Z(1-\phi)} \quad (7)$$

(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'.)

Proof. We begin from the sufficient and necessary condition stated in Lemma 10. Note that we can generate any ranking such that $h_{k+1} \succ h_i$ (for some given i) by iterating over all permutations of $H \setminus \{h_i\}$, and for each permutation, placing h_{k+1} in every slot above h_i. There are at most $n-1$ slots that h_i could be placed in (*i.e.* when h_{k+1} is drawn as the last element).

Let σ' be identical to the reference ranking σ, except with h_i removed. Rename every element after h_i such that it corresponds to its current index: in other words, $h'_j = h_{j+1}$ for all $j \geq i$. Let η' be some arbitrary ranking drawn from $D^{\phi,\sigma'}$. Let $H' = H \setminus \{h_i\}$. Remember, $S' = \{h_1, ..., h_{k+1}\} \setminus \{h_j\}$. Thus, we note that:

$$\sum_{\eta \in H_\succ} \sum_{h_i \in S'} P(h_i \text{ avail})\chi(h_{k+1} \succ_\eta h_i)v(h_{k+1}, \eta)P(\eta|D^{\phi,\sigma}) \leq \sum_{h_i \in S'} \left[P(h_i \text{ avail})\Big(\sum_{\eta' \in H'_\succ} v(h'_k, \eta')P(\eta'|D^{\phi,\sigma'})(\sum_{l=1}^{n} \frac{\phi^l}{Z})\Big)\right] \quad (8)$$

However, note that ϕ^l is a geometric series. We let $n \to \infty$, giving us:

$$\sum_{h_i \in S'} \big[P(h_i \text{ avail})\mathbb{E}(v(h'_k)|D^{\phi,\sigma'}) \sum_{l=1}^{n} \frac{\phi^l}{Z}\big] \leq \sum_{h_i \in S'} P(h_i \text{ avail})\mathbb{E}(v(h'_k)|D^{\phi,\sigma'}) \frac{\phi}{Z(1-\phi)} \tag{9}$$

Thus, because Eq 9 is an upper bound, it is sufficient to show the following, as required:

$$\begin{aligned} P(h_j \text{ avail})\mathbb{E}(v(h_j)|D^{\phi,\sigma}) \geq P(h_j \text{ avail})\mathbb{E}(v(h_{k+1})|D^{\phi,\sigma}) \\ + \sum_{h_i \in S'} P(h_i \text{ avail})\mathbb{E}(v(h'_k)|D^{\phi,\sigma'}) \frac{\phi}{Z(1-\phi)} \end{aligned} \tag{10}$$

□

While we have focussed on existence of assortative interviewing in this section, we do note that other interviewing equilibria may also exist. For example, if $\phi = 1$ in the Mallows model, then residents draw rankings from the uniform distribution. The outcome, first noted by Lee and Schwarz under a different model (Lee & Schwarz, 2009), residents and hospitals are divided into n/k subsets and matched inside those subsets, also forms an equilibrium.

Observation 13. *When residents draw their preferences iid from the uniform distribution, 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 $\{h_{(j-1)k+1}, \ldots, h_{jk}\}$.*

This follows from the condition in Lemma 10. As we are drawing from a uniform distribution, $\mathbb{E}(v(h_j)|D)$ is identical for any hospital h_j, eliminating all terms involving the valuation function, simply leaving probabilities that any alternative is available. Note that now r_1 is indifferent between any alternatives, as she has equal likelihood (probability 1) to get any of them; say she chooses $h_1, ..., h_k$. Then, h_2 prefers $h_{k+1}, ..., h_{2k}$, as he is indifferent between any alternatives that r_1 has *not* interviewed with. This process continues until the desired structure is formed.

5. Assortative Equilibria for Small Budgets

We now discuss assortative equilibria when participants' interviewing budget is $k \leq 3$. To ground the work we instantiate the scoring or utility functions of the residents using different classes of scoring rules. In particular, we look use three different scoring rules, inspired by the social choice literature, in order to better ascertain the effect of resident utility-structure on assortative equilibria.

The first function we consider is *plurality-based*, where $v(s_1) = 1$ and $v(s_i) = 0$ for all $i > 1$.[4] This utility function captures extreme situations where residents only get utility from being matched to their top choice. The second function we consider is *Borda-based.* In this function, residents' utility drops linearly in proportion to the rank of the alternative to

4. 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.

which they are matched. Formally, for any slot s_i, $v(s_i) = n-i+1$ where n is the number of alternatives (hospitals) in the market. Finally, we investigate a scoring function in between plurality and Borda. The *exponential* scoring function allows for utility to exponentially decrease as a resident is matched to a lower ranked alternative; $v(s_i) = (\frac{\epsilon}{2})^{i-1}$, $0 < \epsilon < 1$.

Our first result is a condition for when a resident with plurality-based scoring functions will interview assortatively.

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

$$P(h_j \text{ avail}) \geq \phi^{k-j+1} \tag{11}$$

Proof. We begin with the condition in Lemma 10:

$$P(h_j \text{ avail})\mathbb{E}(v(h_j)|D^{\phi,\sigma}) > \tag{12}$$

$$P(h_j \text{ avail})\mathbb{E}(v(h_{k+1})|D^{\phi,\sigma}) + \sum_{\eta \in H_{\succ}} P(\eta|D^{\phi,\sigma}) \cdot \big[\sum_{h_i \in S'} P(h_i \text{ avail})\chi(h_{k+1} \succ_\eta h_i)v(h_{k+1}, \eta)\big] \tag{13}$$

We instantiate this condition for the plurality function, noting that $v(h, \eta) > 0$ iff h is top-ranked in η. This allows us to greatly simplify Eq. 13:

$$\begin{aligned} P(h_j \text{ avail})\mathbb{E}(v(h_{k+1})|D^{\phi,\sigma}) + \sum_{i=1}^{j-1} P(h_i \text{ avail})\mathbb{E}(v(h_{k+1})|D^{\phi,\sigma}) \\ + \sum_{i=j+1}^{k} P(h_i \text{ avail})\mathbb{E}(v(h_{k+1})|D^{\phi,\sigma}) = \sum_{i=1}^{k} P(h_i \text{ avail})\mathbb{E}(v(h_{k+1})|D^{\phi,\sigma}) \end{aligned} \tag{14}$$

But, again, as the expected value for any h_j is simply the probability that h_j is in s_1 this further simplifies to:

$$P(h_{k+1} \text{ in } s_1) \sum_{i=1}^{k} P(h_i \text{ avail}) = P(h_{k+1} \text{ in } s_1)$$

Note that $\sum_{i=1}^{k} P(h_i \text{ avail}) = 1$ as all residents $r_1, ... r_{k-1}$ have been matched with exactly $k-1$ hospitals in $h_1, ..., h_k$, leaving exactly one hospital left with probability 1. Applying Lemma 4:

$$\begin{aligned} P(h_j \text{ avail})\frac{\phi^{j-1}}{1 + ... + \phi^{n-1}} \geq& \frac{\phi^k}{1 + ... + \phi^{n-1}} \\ P(h_j \text{ avail}) \geq& \phi^{k-j+1} \end{aligned} \tag{15}$$

□

We note that there is a strong relationship between the strategic behaviour of plurality-based residents and exponential-based residents. In particular, if assortative interviewing is an equilibrium for plurality, then there exists some set of exponential valuation functions that likewise admit an assortative interviewing equilibrium.

Lemma 15. *If for a given interviewer budget k and dispersion parameter ϕ, the condition of Lemma 14 is satisfied for a plurality valuation function with a strict inequality, then there exist exponential valuations under which assortative interviewing is an equilibrium.*

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

Proof. Looking at the condition of Lemma 10

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

We will first expand the value expectation ($\mathbb{E}$):

$$P(h_j \text{ avail})\sum_{i=1}^{n} P(h_j \text{ in } s_i)v(s_i) \tag{16}$$

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

$$+ \ldots + \sum_{\substack{\eta \in H_{\succ}| \\ h_{k+1} \text{ in } s_{n-1}}} P(\eta|D^{\phi,\sigma}) \sum_{h_i \in S'} P(h_i \text{ avail})\chi(h_{k+1} \succ_\eta h_i)v(s_{n-1}) \tag{17}$$

Note that for any $1 \leq \ell \leq n$,

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

Thus, combining Eq. 17 and Lemma 10, it is sufficient to show the following holds whenever plurality admits an assortative interviewing equilibrium:

$$P(h_j \text{ avail})P(h_j \text{ in } s_1)v(s_1) \geq P(h_j \text{ avail})P(h_{k+1} \text{ in } s_1)v(s_1) + \sum_{\ell=2}^{n} v(s_\ell) \tag{18}$$

We assume that for plurality valuation, the condition has a strict inequality. In other words:

$$P(h_j \text{ avail})P(h_j \text{ in } s_1) > P(h_j \text{ avail})P(h_{k+1} \text{ in } s_1)$$

Hence, there is an $\bar{\epsilon} \leq 1$ such that for all $1 \leq j \leq k$,

$$P(h_j \text{ avail})P(h_j \text{ in } s_1) - \bar{\epsilon} > P(h_j \text{ avail})P(h_{k+1} \text{ in } s_1)$$

Now, for $\epsilon < \frac{\bar{\epsilon}}{2}$, examine the valuation function $v(s_\ell) = \epsilon^{\ell-1}$. Note that $\sum_{\ell=2}^{n} \epsilon^{\ell-1} \leq \sum_{\ell=1}^{\infty} \epsilon^{\ell} = \frac{\epsilon}{1-\epsilon} \leq 2\epsilon$. This simplifies such that it satisfies Eq. 18, as required:

$$P(h_j \text{ avail})P(h_j \text{ in } s_1) > P(h_j \text{ avail})P(h_{k+1} \text{ in } s_1) + 2\epsilon \geq P(h_j \text{ avail})P(h_{k+1} \text{ in } s_1) + \sum_{\ell=2}^{n} v(s_\ell)$$

□

5.1 Assortative Interviewing with Two Interviews

We start by studying the case where residents are only allowed to interview with 2 hospitals. We show that for sufficiently small dispersion, ϕ, in the Mallows model from which residents are drawing their preferences, assortative interviewing is an equilibrium for plurality-based, Borda-based, and exponential scoring functions. Furthermore, we show that the equilibrium is sensitive to both the dispersion and the structure of the scoring functions.

Theorem 16. *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 $0 < \phi \leq 0.6180$. □

Though we do not formally state it, combining Theorem 16 and Lemma 15 shows that for exponential, when $0 < \phi < 0.6180$, there exists an ε such that if residents' scoring function is an exponential function dominated by $(\frac{\varepsilon}{2})^{(i-1)}$ with $\varepsilon > 0$, assortative interviewing is an equilibrium for that ϕ.

We now similarly show that when $k = 2$, assortative interviewing is also an equilibrium for Borda. We again directly compute the expected payoffs for the interviewing sets in question, finding that $\{h_1, h_2\}$ has the highest expected payoff (and is thus a best response).

Theorem 17. *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. We begin by noting that, because of Lemma 8, 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 9, 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, $P_{\mu(r_i)}(h) = P(\mu(r_i) = h)$.

Case 1. Because we have fixed $h_2 \succ h_1 \succ h_3$ or $h_2 \succ h_3 \succ h_1$, we know exactly what r_2's match will be. As we know r_1's interviewing set ($\{h_1, h_2\}$), and the distribution r_1's preferences are drawn *i.i.d.*, we know the likelihood that either h_1 or h_2 is available; by Lemma 2, $P(\mu(r_1) = h_1) = \frac{1}{1+\phi}$. Using this information, the payoff function, and the definition of η, η', we find a lower bound:

$$U_1 = \sum_{\eta \in P(H)^{h_1 \succ h_0 \succ h_2}} Pr_{\mu(r_0)}(h_1) \Big[(v(h_0, \eta) - v(h_2, \eta))Pr(\eta|D^{\phi,\sigma}) + (v(h_0, \eta') - v(h_2, \eta'))Pr(\eta'|D^{\phi,\sigma})\Big]$$

$$U_1 \geq P_{\mu(r_1)}(h_2)(1)(1-\phi)P(h_2 \succ h_1 \succ h_3) = \left(\frac{\phi}{1+\phi}\right)\left(\frac{\phi}{(1+\phi)(1+\phi+\phi^2)}\right)(1-\phi) \tag{19}$$

Case 2. We fix $h_1 \succ h_2 \succ h_3$ or $h_3 \succ h_2 \succ h_1$. This case is analogous to Case 1:

$$U_2 = \sum_{\eta \in P(H)^{h_0 \succ h_1 \succ h_2}} Pr_{\mu(r_0)}(h_0) \Big[(0)Pr(\eta|D^{\phi,\sigma}) + (v(h_1, \eta') - v(h_2, \eta'))Pr(\eta'|D^{\phi,\sigma})\Big]$$
$$+ Pr_{\mu(r_0)}(h_1) \Big[(v(h_0, \eta) - v(h_2, \eta))Pr(\eta|D^{\phi,\sigma}) + (v(h_0, \eta') - v(h_2, \eta'))Pr(\eta'|D^{\phi,\sigma})\Big]$$

$$U_2 \geq P(h_1 \succ h_2 \succ h_3)\frac{2}{1+\phi}(\phi - \phi^3 - \phi^4) \tag{20}$$

Case 3. We fix $h_1 \succ h_3 \succ h_2$ or $h_3 \succ h_1 \succ h_2$. Again, we look at pairs of rankings η, η', where η is consistent with $h_1 \succ h_3 \succ h_2$, and η' is identical to η, except $\text{rank}(h_1, \eta) = \text{rank}(h_3, \eta')$, and $\text{rank}(h_3, \eta) = \text{rank}(h_1, \eta')$.

Then, as before, we sum over all possible rankings consistent with $h_1 \succ h_3 \succ h_2$, but we break this into two subcases, so that $U_3 = U_{3a} + U_{3b}$:

$$U_{3a} = \sum_{\eta \in H_{\succ|h_1 \succ h_3 \succ h_2}} P_{\mu(r_1)}(h_1)[(v(h_2, \eta) - v(h_3, \eta))P(\eta|D^{\phi,\sigma}) + (v(h_2, \eta') - v(h_3, \eta'))P(\eta'|D^{\phi,\sigma})]$$

$$U_{3b} = \sum_{\eta \in H_{\succ|h_1 \succ h_3 \succ h_2}} P_{\mu(r_1)}(h_2)[(v(h_1, \eta) - v(h_3, \eta))P(\eta|D^{\phi,\sigma}) + (v(h_1, \eta') - v(h_3, \eta'))P(\eta'|D^{\phi,\sigma})]$$

Case U_{3b} is similar to Cases 1 and 2:

$$U_{3b} = \sum_{\eta \in P(H)^{h_0 \succ h_2 \succ h_1}} Pr_{\mu(r_0)}(h_1)[(v(h_0, \eta) - v(h_2, \eta))\frac{\phi^{d(\eta,\sigma)}}{Z} + (v(h_2, \eta) - v(h_0, \eta))\frac{\phi^{d(\eta,\sigma)+1}}{Z}$$

$$U_{3b} \geq \frac{\phi}{\phi+1}(1-\phi)P(h_1 \succ h_3 \succ h_2) \tag{21}$$

Case U_{3a}, however, is different from all other cases, in that *all* terms are negative. We note that U_{3a} as above is a monotonically decreasing function in terms of n. Thus, if U_{3a} converges as $n \to \infty$, we have found a lower bound for all n. Using this technique, we show the following bound holds:

$$U_{3a} \geq P_{\mu(r_1)}(h_1)\frac{-\phi}{(1+\phi)(1+\phi+\phi^2)}\Big(\frac{\phi}{(1-\phi)^4}+\frac{1}{3(1-\phi)^3}+\frac{2}{3}\Big)(1+\phi) \tag{22}$$

We have considered all cases, and can now combine them together. We add the bounds for U_1 (Eq. 19), U_2 (Eq. 20), U_{3a} (Eq. 22), and U_{3b} (Eq. 21). We simplify using Corollaries 2 and 3, giving us:

$$\begin{aligned} u_{r_2}(\{h_1,h_2\})-u_{r_2}(\{h_2,h_3\}) \geq & \frac{\phi^2}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}(1-\phi)+\frac{2(\phi-\phi^3-\phi^4)}{(1+\phi)(1+\phi)(1+\phi+\phi^2)} \\ & -\frac{\phi}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}\Big(\frac{\phi}{(1-\phi)^4}+\frac{1}{3(1-\phi)^3}+\frac{2}{3}\Big)(1+\phi) \\ & +\frac{\phi^2}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}(1-\phi) \end{aligned} \tag{23}$$

Thus, Eq. 23 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. 23, 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:

$$u_{r_2}(\{h_1,h_2\})-u_{r_2}(\{h_1,h_3\}) \quad \geq \frac{1}{(1+\phi)(1+\phi+\phi^2)}\Big[1+\phi-2\phi^2-2\phi^3-2\phi^3\Big(\frac{\phi}{(1-\phi)^4}+\frac{1}{3(1-\phi)^3}+\frac{2}{3}\Big)\Big] \tag{24}$$

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 8, this is an equilibrium for all r_i as required. □

5.2 Assortative Interviewing with Three Interviews

Interestingly, when residents can interview with up to three hospitals, assortative interviewing continues to be an equilibrium for plurality-based and exponential scoring functions but is no longer an equilibrium if residents have Borda-based scoring functions.

We begin with the counter-example for Borda and $k = 3$. In particular, assortative interviewing is not an equilibrium for a market with 4 residents, 4 hospitals, and 3 interviews. We prove this by directly computing the marginal value for interviewing with h_1 instead of interviewing with h_4. In our example, for *all* $\phi > 0$ the expected marginal value for interviewing with h_4 is better than interviewing with h_1, assortative interviewing cannot be an equilibrium.

Theorem 18. *Assortative interviewing is not always an equilibrium under the Borda valuation function for any* $0 < \phi \leq 1$.

Proof. We provide a counterexample for $n = 4$, $k = 3$. Suppose residents r_1 and r_2 interview assortatively, both interviewing with $S = \{h_1, h_2, h_3\}$. We show that for resident r_3, interviewing with interviewing set $S' = \{h_2, h_3, h_4\}$ dominates interviewing with $S = \{h_1, h_2, h_3\}$ for all ϕ.

By Lemma 9, it is sufficient to show that if the marginal value in interviewing with h_4 dominates the marginal value in interviewing with h_1 (as these two sets only differ by these two items), then interviewing with $\{h_2, h_3, h_4\}$ dominates $\{h_1, h_2, h_3\}$. We thus instantiate Equation 4 for $n = 4$, $k = 3$, S, and S' as above for resident r_3. Note that $Z = (1+\phi)(1+\phi+\phi^2)(1+\phi+\phi^3)$. Let $\mathbb{E}(u(h_i, S))$ denote the expected marginal value in interviewing alternative h_i in set S; remember $v(s_i) = 5 - i$.

$$\mathbb{E}(u(h_1, S)) = \sum_{\eta \in H_{\succ}} v(h_1, \eta) P(\mu(h_1) = r_3 | S, \eta, D^{\phi,\sigma}) P(\eta | D^{\phi,\sigma}) \tag{25}$$

$$\mathbb{E}(u(h_4, S')) = \sum_{\eta \in H_{\succ}} v(h_4, \eta) P(\mu(h_4) = r_3 | S', \eta, D^{\phi,\sigma}) P(\eta | D^{\phi,\sigma}) \tag{26}$$

As before, Equation 25 is simply the probability that h_1 is available times the expected value of h_1. As noted, $\mathbb{E}(v(h_1)|D^{\phi,\sigma}) = \sum_{i=1}^{4} P(h_1 \text{ in } s_i) \cdot v(s_i) = \sum_{i=1}^{4} P(h_1 \text{ in } s_i) \cdot (5-i)$. However, using Lemma 4, we know that $P(h_1 \text{ in } s_i) = \frac{\phi^{i-1}}{1+\phi+\phi^2+\phi^3}$, giving:

$$\mathbb{E}(u(h_1, S)) = P(h_1 \text{ avail}) \mathbb{E}(v(h_1)|D^{\phi,\sigma}) = P(h_1 \text{ avail}) \frac{4 + 3\phi + 2\phi^2 + \phi^3}{1 + \phi + \phi^2 + \phi^3} \tag{27}$$

Let $P(h_i \text{ taken})$ denote the probability that either r_1 is matched to h_i, or r_2 is matched to h_i (*i.e.* h_i is taken by the time we get to resident r_3). Also let $P(\mu(r_3) = h_4 | h_4 \text{ in } s_i)$ denote the probability that r_3 is matched to h_4 if h_4 is in slot s_i in r_3's ranking. This is easily calculable by enumerating over the subset of possible rankings such that this occurs, given that r_1 and r_2 have already taken certain alternatives. Then, using Lemma 4 again and an analogous approach as above, we show:

$$\begin{aligned}
\mathbb{E}(u(h_4, S')) &= \sum_{i=1}^{4} v(s_i) P(h_4 \text{ in } s_i) P(\mu(r_3) = h_4 | h_4 \text{ in } s_i) \\
&= \frac{4\phi^3}{1+\phi+\phi^2+\phi^3} + \frac{3}{Z}\big(\phi^2 + \phi^3 + P(h_2 \text{ taken})(\phi^3 + \phi^4) + P(h_3 \text{ taken})(\phi^4 + \phi^5)\big) \\
&+ \frac{2}{Z}\big(P(h_2 \text{ taken})(\phi + \phi^2) + P(h_3 \text{ taken})(\phi^2 + \phi^3) + P(h_1 \text{ avail})(\phi^3 + \phi^4)\big) \\
&+ \frac{P(h_1 \text{ avail})}{1+\phi+\phi^2+\phi^3}
\end{aligned} \tag{28}$$

As we assume that residents r_1 and r_2 both interview with S, the probability that h_1 is available, or h_2 (resp. h_3) is taken is the same across both $\mathbb{E}(u(h_1, S))$ and $\mathbb{E}(u(h_4, S'))$. We instantiate these as follows, by determining the probability that r_1 is matched to some hospital h_j other than h^*, and enumerate the probabilities of all rankings such that r_2 is matched to some hospital $h'_j \neq h^*$ given that r_1 is matched to h_j:

$$P(h_1 \text{ avail}){=}P(\mu(r_1){=}h_2|S,D^{\phi,\sigma})(\tfrac{\phi^2}{1+\phi+\phi^2+\phi^3}+\tfrac{\phi^2+\phi^3+\phi^4+2\phi^5+\phi^6}{Z})$$
$$+P(\mu(r_1){=}h_3|S,D^{\phi,\sigma})(\tfrac{\phi}{1+\phi+\phi^2+\phi^3}+\tfrac{\phi^3+2\phi^4+2\phi^5+\phi^6}{Z})$$
$$P(h_2 \text{ taken}){=}P(\mu(r_1){=}h_2|S,D^{\phi,\sigma}){+}P(\mu(r_1){=}h_3|S,D^{\phi,\sigma})(\tfrac{\phi}{1+\phi+\phi^2+\phi^3}+\tfrac{\phi^3+2\phi^4+2\phi^5+\phi^6}{Z})$$
$$+P(\mu(r_1){=}h_1|S,D^{\phi,\sigma})(\tfrac{\phi}{1+\phi+\phi^2+\phi^3}+\tfrac{1+\phi+\phi^2+\phi^3+\phi^4+\phi^5}{Z})$$
$$P(h_3 \text{ taken}){=}P(\mu(r_1){=}h_3|S,D^{\phi,\sigma}){+}P(\mu(r_1){=}h_2)(\tfrac{\phi^2}{1+\phi+\phi^2+\phi^3}+\tfrac{\phi^2+\phi^3+\phi^4+2\phi^5+\phi^6}{Z})$$
$$+P(\mu(r_1){=}h_1|S,D^{\phi,\sigma})(\tfrac{\phi^2}{1+\phi+\phi^2+\phi^3}+\tfrac{1+\phi+\phi^2+\phi^3+\phi^4+\phi^5+\phi^6}{Z})$$

We note that it is also possible to calculate exact values for the probability that r_1 is matched to h_1, h_2, h_3. We do this by calculating the probability that alternative is first, or the probability that alternative is second, and h_4 is first:

$$P(\mu(r_1) = h_1|S, D^{\phi,\sigma}) = P(h_1 \text{ in } s_1) + P(h_1 \text{ in } s_2 \text{ and } h_4 \text{ in } s_1) = \frac{1}{1+\phi+\phi^2+\phi^3} + \frac{\phi^3+\phi^4}{Z}$$
$$P(\mu(r_1) = h_2|S, D^{\phi,\sigma}) = P(h_2 \text{ in } s_1) + P(h_2 \text{ in } s_2 \text{ and } h_4 \text{ in } s_1) = \frac{\phi}{1+\phi+\phi^2+\phi^3} + \frac{\phi^4+\phi^5}{Z}$$
$$P(\mu(r_1) = h_3|S, D^{\phi,\sigma}) = P(h_3 \text{ in } s_1) + P(h_3 \text{ in } s_2 \text{ and } h_4 \text{ in } s_1) = \frac{\phi^2}{1+\phi+\phi^2+\phi^3} + \frac{\phi^5+\phi^6}{Z}$$

Note that, by combining the equations for the probabilities, we are left with two equations that are dependent only on ϕ. Moreover, after instantiating $\mathbb{E}(u(h_1, S))$ and $\mathbb{E}(u(h_4, S')$ as above, we note that both functions are continuous on the interval $(0, 1]$. Using numerical analysis techniques, it can be shown that there are no zeros for the function $\mathbb{E}(u(h_1, S)) - \mathbb{E}(u(h_4, S'))$ on the interval $(0, 1]$, and the function is negative on the interval $(0, 1]$ providing the counterexample as required. □

By directly computing expected payoffs, we now show that assortative interviewing is an equilibrium for plurality (and thus exponential) for $k = 3$:

Theorem 19. *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. 11 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:

$$P(h_1 \text{ avail}) = P(\mu(r_1) \neq h_1)P(\mu(r_2) \neq h_1|\mu(r_1) \neq h_1)$$
$$P(h_1 \text{ avail}) = (\frac{\phi + 2\phi^2 + \phi^3}{(1+\phi)(1+\phi+\phi^2)})(\frac{\phi^2 + 2\phi^3}{(1+\phi+\phi^2)})$$

Using numerical methods to find the roots of $P(h_1 \text{ avail}) - \phi^3$, we can show that Eq. 11 holds when $0 < \phi \leq 0.4655$.

□

6. Assortative Equilibria for Large Budgets

We begin by showing that when there is a setting for which there is no assortative equilibria for plurality, then there is no scoring function with assortative equilibria. We use this result to show that, for sufficiently small dispersion parameter ϕ and for $k > 3$ interviews, assortative interviewing cannot be an equilibrium under any scoring function. We then provide a specific counterexample for *all* ϕ when $k = 4$ for plurality, implying there is no assortative equilibrium for any scoring function. This suggests that, for a wide category of resident valuation functions under a Mallows distribution, contrary to some real-world behaviour, assortative interviewing is not an equilibrium.[5]

We provide one additional lemma regarding a bound on the availability of any given alternative h_i at the time resident r_k is being matched by the mechanism to her favourite remaining hospital. This probability is dependent on ϕ: for any hospital h_i such that $i < k$, as $\phi \to 1$, the probability h_i is available goes to $\frac{1}{k}$; as $\phi \to 0$, this probability goes to 0. Instead of looking at the probability directly, we look at the probability that a preference profile will admit a stable match such that h_i is available, and bound that.

Lemma 20. *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 $\eta_1, \dots, \eta_{n-1} \in D^{\phi,\sigma}$ of $k-1$ preferences (for $r_1, \dots, r_{k-1}$) such that h_i is available for r_k has probability $P(r_1 = \eta_1, r_2 = \eta_2, \dots, r_{k-1} = \eta_{k-1}) < \frac{\phi^\gamma}{Z^{k-1}}$, where $\gamma = \sum_{j=1}^{k-i} j$ and Z is the normalizing factor for a Mallows model.*

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

We further note that showing that plurality fails assortative interviewing is a strong indication that other valuation functions will also not admit assortative interviewing equilibria. In some sense, because plurality only provides a payoff when agents get their most preferred alternative, this benefits assortative interviewing: everyone wants a chance at the alternatives with the highest probability of being first in the ranking (that still have non-zero chance of being available). Thus, if h_1's marginal utility for being included in the interviewing set is *less* than h_{k+1}'s under plurality, it will be less under any other scoring rule.

Theorem 21. *Fix an instance of the Interviewing with a Limited Budget game $\Psi = \langle k, \phi, plurality \rangle$. If hospital h_1 causes the condition in Lemma 10 to be falsified (*i.e. *$\{h_2, \dots, h_{k+1}\}$ has a better expected payoff than $\{h_1, \dots, h_k\}$), then for k and ϕ, assortative interviewing is not an equilibrium for any valuation function.*

5. We note that the definition of assortative investigated here is fairly restrictive, but has desirable properties. In Section 7 we provide an example showing that even a relaxed version of assortative interviewing may not be an equilibrium in this setting.

Proof. Looking at the condition of Lemma 10

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

We again begin by expanding the value expectation ($\mathbb{E}$), as we did in Eq. 17: This can be divided to n different inequalities:

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

We shall show that under the theorem's assumptions, none of these inequalities hold for h_1, and therefore the general inequality (Lemma 10) 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 noted in Observation 6, 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 Observation 6, for any $\ell > k$, $P(h_1 \text{ in } s_\ell) < P(h_{k+1} \text{ in } s_\ell)$ which gives us:

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

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 10 does not hold for $j = h_1$ for any valuation function. □

Intuitively, there is a tradeoff between the likelihood that a hospital will be available for resident r_k by the time it is her turn to be matched, and the expected value of that hospital. Both of these are strongly tied to the dispersion parameter ϕ of the Mallows model residents are drawing from: as the dispersion parameter grows, the difference in expected value of any given hospital goes to 0. As the dispersion parameter gets small (*i.e.* goes to 0), the expected value of any hospital h_i goes to the value of its slot in expectation, $v(s_i)$. However, the likelihood it is taken by some higher ranked r_j (*i.e.* with $j < i$) also approaches 1. The following theorem addresses the latter case: for sufficiently small dispersion, even though the expected value of a hospital is high, the likelihood it will be available is so low that residents are disincentivized from choosing to interview with it.

We first show that for $k = 4$, assortative interviewing is not an equilibrium for *any* $\phi < 1$ and any scoring rule. As it does not seem as though there is anything special about four interviews, we further conjecture that likewise assortative interviewing is not an equilibrium for $k > 4$. We can show that for $k > 4$ and ϕ sufficiently small that assortative interviewing is not 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 $0 < \phi < 1$.*

Proof. By Theorem 21, if assortative interviewing is not an equilibrium for plurality, it is never an equilibrium for any scoring rule. As noted before Eq. 11 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 $0 < \phi < 1$. □

We now consider the case of $k > 4$.

Theorem 23. *Given an interviewing budget of $k > 4$ interviews, there exists $0 < \varepsilon < 1$ scuh that for any scoring function v no assortative interviewing forms an equilibrium for dispersion parameter $0 < \phi < \varepsilon$.*

Proof. Thanks to Theorem 21, it is enough for us to shown there is no assortative equilibrium under plurality (and that h_1 violates Lemma 10's condition). We again begin with the simplification from Lemma 14: $P(h_j \text{ avail}) \geq \phi^{k-j+1}$. Once again, appealing to Lemma 20, we know $P(h_j \text{ avail})$ is of the form:

$$P(h_j \text{ avail}) = \frac{X(k)}{Z^{k-1}}\phi^{\sum_{i=1}^{k-j} i} + \frac{X^1(k)}{Z^{k-1}}\phi^{1+\sum_{i=1}^{k-j} i} + \ldots + \frac{X^\ell(k)}{Z^{k-1}}\phi^{(k\sum_{i=1}^{k-j} i)-1} + \frac{1}{Z^{k-1}}\phi^{k\sum_{i=1}^{k-j} i} \tag{29}$$

($X(k), X^1(k), \ldots, X^\ell(k)$ 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 29 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 11). □

It seems quite unlikely that for $k > 4$, assortative interviewing is an equilibrium. Intuitively, if it is an equilibrium it should be for low ϕ: this is when the expected value of hospital h_i is very close to $v(s_i)$. However, this is also when residents $r_1, ..., r_{k-1}$ are all most likely to be matched with hospitals $h_1, ..., h_{k-1}$. We leave open the possibility that there may exist some δ such that when $0 < \varepsilon < \phi < \delta \leq 1$, assortative interviewing is an equilibrium for plurality.

7. Reach and Safety Strategies for a Small Budget

Our analysis has shown that assortative interviewing equilibria are rather exceptional and essentially can only be guaranteed for a very small number of interviews. This suggests that there may not be a simple characterization of interviewing equilibria. Consider the case for $k = 2$ interviews and the Borda scoring rule where we only guarantee assortative interviewing for some sufficiently small dispersion parameter ϕ. To gain better insight into the strategic behaviour of the residents as a function of ϕ, we calculated the exact values of ϕ where the interviewing equilibria changes in small markets. In doing so, we see that the structure of the interviewing equilibria 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.

Figure 1 depicts a market with 4 hospitals, 4 residents, and 2 interviews ($n = 4$, $k = 2$). The figure shows what sets are being chosen by the different residents for any dispersion

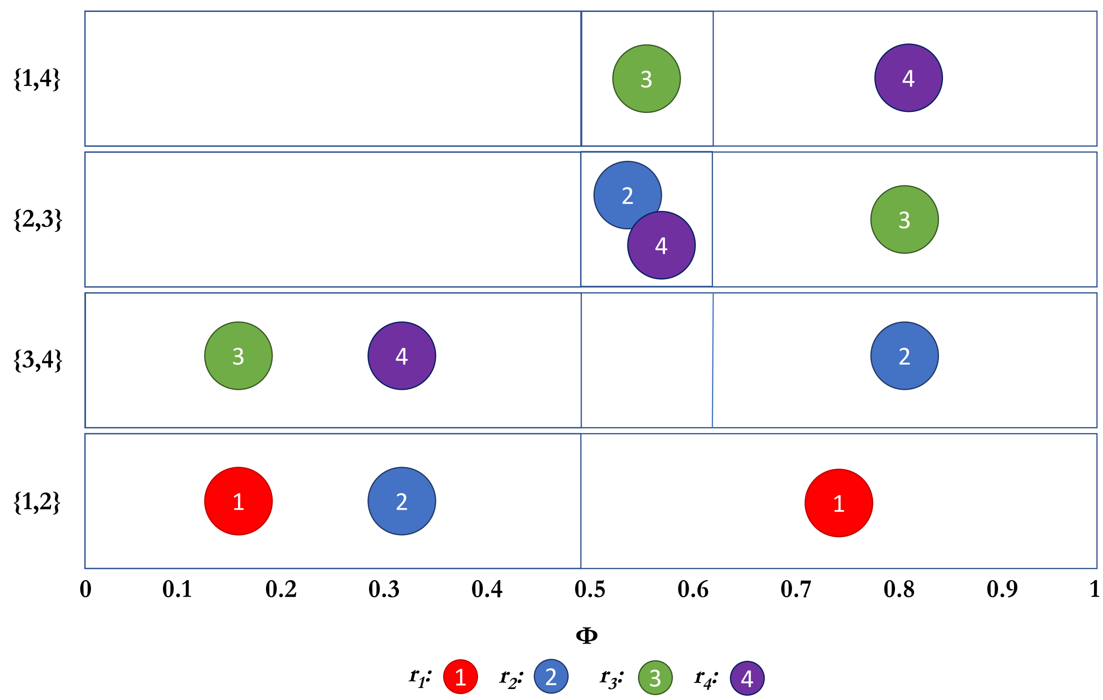

Figure 1: Interviewing sets of residents as a function of ϕ.

ϕ. As ϕ increases, we explicitly see the trade-off between a safer choice, and a better expected payoff value for individual alternatives. For small ϕ, as the theoretical results showed, assortative interviewing is optimal, and r_2 chooses $\{h_1, h_2\}$, while r_3 and r_4 choose $\{h_3, h_4\}$.

Interestingly, for $\phi \in [0.5, 0.62]$, 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" vs. safe behaviour we see in college admissions markets; namely, r_3's best response now is to interview with h_1, h_4 (*i.e.* a "reach" choice, and a "safe" bet), while r_4, being left without any truly "safe" option, aims slightly higher than its rank. As ϕ grows and approaches 1, any ordering of hospitals is as likely as another, making r_2's choice $\{h_3, h_4\}$, which are as likely as any to be highly ranked, and both are available. The desire to avoid choosing hospitals that are already chosen by too many other agents also drives r_3 and r_4 to $\{h_2, h_3\}$ and $\{h_1, h_4\}$, respectively; that is, they both want to avoid competing with r_1 and r_2. We hypothesize that this "reach" and "safety" behaviour is present for small ϕ in markets with larger interviewing budgets.

8. Conclusions and Future Directions

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 instantiate the payoff functions using different scoring functions (plurality-based, exponential, and Borda-based) when residents' preferences are drawn independently from the same Mallows model distribution. While assortative interviewing is an equilibrium when interviewing budgets are small and residents' preferences are sufficiently similar (*i.e.* the dispersion parameter in the Mallows model is small), in general it is not an equilibrium. This was a surprising result since assortative interviewing is observed in certain matching markets, and, when it is an equilibrium, supports several highly desirable properties such as maximizing the number of matched residents. In particular, if residents interview assortatively, then they naturally form a bipartite graph interviewing structure with n/k disconnected complete components. Under very different modelling assumptions (*i.e.* the impartial culture model), Lee and Schwartz showed the existence of a similarly structured equilibrium (Lee & Schwarz, 2009), and so it was somewhat surprising that the existence of this equilibria was so highly dependent on both scoring-function structure and distribution from which the underlying preferences were drawn.

There are numerous future research questions raised by our results. First, while we believe that the space of scoring functions used in this paper was broad in its scope, we always assumed that residents' underlying ranked preferences were drawn from a distribution generated by the ϕ-Mallows model. While the ϕ-Mallows model is standard in the literature, it is possible that other preference distributions (*e.g.* Plackett-Luce) may better support assortative interviewing. Second, the analysis did rely on the assumption that one side of the market maintained a master list. While master-lists do occur in real-world matching markets, lifting this assumption may broaden the results. However, the removal of the master-list assumption would complicate the analysis significantly, increasing the complexity of the payoff function formulation. Furthermore, we could consider modifying our definition of an interview set. Currently we assume that residents could interview up to k hospitals for free, but an alternative model to consider would be to allow each resident r to have a budget b_r, and incur a cost, $c_r(h)$, when interviewing hospital h, with the constraint that if S is the set of hospitals interviewed by resident r, then $\sum_{h \in S} c_r(h) \leq b_r$.

Our long-term research goal is to better understand the extent to which "natural equilibria" exist in matching games, and how such equilibria correspond with observed behaviour in actual markets. While assortative interviewing is often not an equilibrium, it is possible that some form of "nearly assortative interviewing" will more generally be an equilibirum. For example, our definition of assortative interviewing is very strict and there may be ways to relax the definition in meaningful ways that better capture interesting behaviour (*e.g.* being assortative for "safety" programs while allowing for a "reach" program). Furthermore, we are interested in techniques that could reduce the cognitive burden placed on participants in matching markets, while also reducing inefficiencies. For example, there may be ways to leverage research on preference elicitation for matching markets (*e.g.* (Drummond & Boutilier, 2014)) with matching market design so as to guide participants to interview with the appropriate programs so as to improve the overall quality of the match.

References

Ajayi, K. F. (2011). School choice and educational mobility: Lessons from secondary school applications in Ghana. Tech. rep., University of California-Berkeley Working Paper.

Anderson, K. D., Dorough, M. M., Stein, C. R., Optenberg, S. A., Thompson, I. M., & others (2000). The Urology Residency Matching Program in Practice. *The Journal of Urology*, *163*(6), 1878–1887.

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

Brandt, F., Conitzer, V., Endriss, U., Lang, J., & Procaccia, A. D. (Eds.). (2016). *Handbook of Computational Social Choice*. Cambridge University Press.

Chade, H., Lewis, G., & Smith, L. (2014). Student Portfolios and the College Admissions Problem. *The Review of Economic Studies*, *81*(3), 971–1002.

Chade, H., & Smith, L. (2006). Simultaneous Search. *Econometrica*, *74*(5), 1293–1307.

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

Coles, P., Kushnir, A., & Niederle, M. (2010). Preference Signaling in Matching Markets. Tech. rep., National Bureau of Economic Research.

Drummond, J., & Boutilier, C. (2013). Elicitation and Approximately Stable Matching with Partial Preferences. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI-13), Beijing*.

Drummond, J., & Boutilier, C. (2014). Preference Elicitation and Interview Minimization in Stable Matchings. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI-14)*, pp. 645–653.

Gale, D., & Shapley, L. S. (1962). College Admissions and the Stability of Marriage. *American Mathematical Monthly*, *69*(1), 9–15.

Gelain, M., Pini, M. S., Rossi, F., Venable, K. B., & Walsh, T. (2009). Male Optimal and Unique Stable Marriages with Partially Ordered Preferences. In *Proceedings of the International Conference on Collaborative Agents-Research and Development (CARE-09)*, pp. 44–55.

Goto, M., Iwasaki, A., Kawasaki, Y., Kurata, R., Yasuda, Y., & Yokoo, M. (2016). Strategyproof matching with regional minimum and maximum quotas. *Artificial Intelligence*, *235*, 40–57.

Hafalir, I. E. (2008). Stability of Marriage with Externalities. *International Journal of Game Theory*, *37*(3), 353–369.

Irving, R. W., Manlove, D. F., & Scott, S. (2008). The stable matching problem with master preference lists. *Discrete Applied Mathematics*, *156*, 2959–2977.

Kurata, R., Goto, M., Iwasaki, A., & Yokoo, M. (2017). Controlled school choice with soft bounds and overlapping types. *Journal of Artificial Intelligence Research*, *58*, 153–184.

Kushnir, A. (2013). Harmful Signaling in Matching Markets. *Games and Economic Behavior*, *80*, 209–218.

Lebanon, G., & Mao, Y. (2008). Non-Parametric Modeling of Partially Ranked Data. *Journal of Machine Learning Research*, *9*, 2401–2429.

Lee, R. S., & Schwarz, M. (2009). Interviewing in Two-Sided Matching Markets. Working paper 14922, National Bureau of Economic Research.

Loewenstein, G. F., Thompson, L., & Bazerman, M. H. (1989). Social Utility and Decision Making in Interpersonal Contexts. *Journal of Personality and Social Psychology*, *57*(3), 426–441.

Lu, T., & Boutilier, C. (2011). Learning Mallows Models with Pairwise Preferences. In *Proceedings ICML 2011*, pp. 145–152 Bellevue, WA.

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

Murphy, T. B., & Martin, D. (2003). Mixtures of Distance-based Models for Ranking Data. *Computational Statistics and Data Analysis*, *41*, 645–655.

National Resident Matching Program (2015). National Resident Matching Program, Results and Data: 2015 Main Residency Match®. Tech. rep., National Resident Matching Program, Washington, DC.

Niederle, M., Roth, A. E., & Sonmez, T. (2008). Matching and Market Design. In Durlauf, S. N., & Blume, L. E. (Eds.), *The New Palgrave Dictionary of Economics (2nd Ed.)*, Vol. 5, pp. 436–445. Palgrave Macmillan, Cambridge.

Pini, M. S., Rossi, F., & Venable, K. B. (2014). Stable matching problems with soft constraints. In *Proceedings of the Thirteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, pp. 1511–1512 Paris, France.

Rastegari, B., Condon, A., Immorlica, N., & Leyton-Brown, K. (2013). Two-sided Matching with Partial Information. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, pp. 733–750. ACM.

Roth, A. E. (2002). The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics. *Econometrica*, *70*(4), 1341–1378.

Zhu, M. (2014). College Admissions in China: A Mechanism Design Perspective. *China Economic Review*, *30*, 618–631.

Appendix A. Proofs From Section 3.3

Proof. (**Lemma 1**) Suppose σ is a prefix of σ'. Then, let σ be some ranking with p elements, including elements a_i and a_j. Let σ' be a ranking of $p+1$ elements with σ as its prefix, and an additional element a_p added at the end. We prove this by starting from the definition of $P(a_i \succ a_j|D^{\phi,\sigma'})$, and using algebraic manipulations to show this is equivalent to the definition of $P(a_i \succ a_j|D^{\phi,\sigma})$.

$$P(a_i \succ a_j|D^{\phi,\sigma'}) = \frac{\sum_{\eta' \in \{a_0,\ldots,a_{p-1},a_p\}_{\succ}^{a_i \succ a_j}} \phi^{d(\eta',\sigma')}}{1(1+\phi)...(1+...+\phi^{p-1}+\phi^{p})} \tag{30}$$

However, because a_i, a_j are in ranking σ, the only difference between summing over the set of all rankings in $\{a_0, ..., a_p\}_{\succ}^{a_i \succ a_j}$ and $\{a_0, ..., a_{p-1}\}_{\succ}^{a_i \succ a_j}$ is that there are p times as many rankings, one for each permutation generated by $\{a_0, ..., a_{p-1}\}_{\succ}$, each one with a_p in a different place (and thus a different Kendall-τ distance). Fixing some $\eta \in \{a_0, ..., a_{p-1}\}_{\succ}$, if a_p is in the last rank position (as it is in σ'), the distance is simply $d(\eta, \sigma)$. If a_p is in the second-to-last position, we have now added in an additional discordant pair, so the distance is $d(\eta, \sigma) + 1$. Using this, we generate the following:

$$\begin{aligned}
P(a_i \succ a_j | D^{\phi,\sigma'}) &= \frac{\sum_{\eta \in \{a_0,...,a_{p-1}\}_{\succ}^{a_i \succ a_j}} \sum_{l=0}^{p} \phi^{d(\eta,\sigma)+l}}{1(1+\phi)...(1+...+\phi^p)} \\
&= \frac{\left[\sum_{\eta \in \{a_0,...,a_{p-1}\}_{\succ}^{a_i \succ a_j}} \phi^{d(\eta,\sigma)}\right]\left[\sum_{l=0}^{p} \phi^l\right]}{1(1+\phi)...(1+...+\phi^p)} \\
&= \frac{\left[\sum_{\eta \in \{a_0,...,a_{p-1}\}_{\succ}^{a_i \succ a_j}} \phi^{d(\eta,\sigma)}\right](1+...+\phi^p)}{1(1+\phi)...(1+...+\phi^{p-1})(1+...+\phi^p)} = \frac{\sum_{\eta \in \{a_0,...,a_{p-1}\}_{\succ}^{a_i \succ a_j}} \phi^{d(\eta,\sigma)}}{1(1+\phi)...(1+...+\phi^{p-1})} \\
&= P(a_i \succ a_j | D^{\phi,\sigma})
\end{aligned}$$

By symmetry, this also holds when σ is a suffix of σ'. □

Proof. (**Corollary 2**) Consider $\sigma = a_i \succ a_{i+1}$, a reference ranking with two elements in it. Then, the set of all potential rankings such that $a_i \succ a_{i+1}$ under $D^{\phi,\sigma}$ is solely the ranking $a_0 \succ a_1$. By the definition of the Mallows model, this ranking has probability $\frac{1}{1+\phi}$. We add some arbitrary prefix σ' to σ and some arbitrary suffix σ'' to σ to create a new reference ranking γ. By Lemma 1, the probability that some η is drawn from $D^{\phi,\gamma}$ such that $a_i \succ_\eta a_{i+1}$ is $\frac{1}{1+\phi}$ as required. □

Proof. (**Corollary 3**) Consider $\sigma^* = a_i \succ a_{i+1} \succ a_{i+2}$, a reference ranking with three elements in it. The set of all potential rankings under D^{ϕ,σ^*} such that $a_i \succ a_{i+1} \succ a_{i+2}$ is solely that ranking. Using the same argument as in Lemma 1, we note that creating some new reference ranking $\gamma = \sigma' \succ \sigma^* \succ \sigma''$ and drawing from $D^{\phi,\gamma}$ does not change the likelihood that we draw a ranking consistent with $a_i \succ a_{i+1} \succ a_{i+2}$.

Therefore, the probability that we draw a ranking β consistent with some permutation η of a_i, a_{i+1}, a_{i+2} under the distribution $D^{\phi,\gamma}$ is simply the probability that we drew η under the distribution D^{ϕ,σ^*}, which is $\frac{\phi^{d(\eta,\sigma^*)}}{(1+\phi)(1+\phi+\phi^2)}$. □

Proof. (**Lemma 4**) This is equivalent to generating the set of all $(n-1)!$ possible rankings excluding alternative a_1 (a_n), and then adding a_1 (a_n) in place j. Whatever the ranking, adding a_1 (a_n) in place j adds $j-1$ ($n-j$) to each possible ranking's Kendall's τ distance from $\sigma \setminus \{a_1\}$ ($\sigma \setminus \{a_n\}$) , making the distance from σ grow by exactly $j-1$ ($n-j$). Similarly, adding a_j in first place adds $j-1$ to the distance from $\sigma \setminus \{a_j\}$, increasing the distance from σ by $j-1$.

However, we also added in an additional element, and must include that in the normalization factor Z. The normalization factor for $n-1$ alternatives is $(1+\phi)(1+\phi^2)...(1+...+\phi^{n-2})$. The normalization factor for n elements is identical, but multiplied by $1+...+\phi^{n-1}$.

□

Proof. (**Lemma 5**) For a_ℓ, $i > \ell > j$. if $a_\ell \succ_\eta a_i$, this adds 1 to the Kendall τ distance of η from σ (due to $a_i \succ_\sigma a_\ell$). But if $a_i \succ_\eta a_\ell$, this means that $a_j \succ_\eta a_\ell$, again adding 1 to the Kendall τ distance of η from σ.

So the Kendall τ distance of η from σ is at least $\sum_{\ell=i}^{j-1} 1 = j - i$, and therefore, $P(\eta) < \frac{\phi^{j-i}}{Z}$. □