Natural Interviewing Equilibria in Matching Settings

Allan Borodin, Joanna Drummond

University of Toronto, Toronto, Ontario M5S 3G4, Canada

Kate Larson

University of Waterloo, Waterloo, Ontario N2L 3G1, Canada

Omer Lev

Ben-Gurion University of the Negev, Beersheba 8410501, Israel

Abstract

A common assumption in matching markets is that both sides fully specify their preferences. However, with many participants, this becomes unreasonable. Facing numerous alternatives, many of which they are unfamiliar with, agents focus on alternatives they consider likely, studying them more carefully.

Using the setting of hospitals and residents with Deferred Acceptance, we examine Nash equilibria arising when hospitals have a master list of residents (e.g., by grade), while residents have some uncertainty, and need to choose a subset of hospitals to *interview* in so they can rank them. Assuming residents' preferences are drawn from a Mallows distribution, we show assortative equilibrium (k top residents interview with k top hospitals, etc.) arise only with small interview sizes. Surprisingly, they do not happen in larger interview sizes, even when residents' preferences are almost identical. We examine simulations on possible outcome equilibrium, showing residents may be pursuing a reach/safety strategy.

Email addresses: `bor@cs.toronto.edu` (Allan Borodin), `jdrummond@cs.toronto.edu` (Joanna Drummond), `kate.larson@uwaterloo.ca` (Kate Larson), `omerlev@bgu.ac.il` (Omer Lev)

URL: `http://www.cs.toronto.edu/~bor/` (Allan Borodin), `https://cs.uwaterloo.ca/~klarson/` (Kate Larson), `http://bgu.ac.il/~omerlev` (Omer Lev)

Keywords: matching, interviewing, equilibria

1. Introduction

Since Gale-Shapley's groundbreaking work [1], matching algorithms, in which elements of disjoint sets are matched to one another based on each agent's preferences of being matched to the different elements of the other set, have proliferated in a variety of useful settings. These range from matching children to schools to matching refugees to countries [2]. The main goal of most of these algorithms is to create *stable* matchings, i.e., assignments which are stable and in which no agent has a way to change their behavior in a way that will improve the outcome for them.

We will take our inspiration and terminology from medical residencies [3] which are one of the main settings implementing matching algorithms in the real world. That is, matching medical students to residencies in hospitals. Such systems are prevalent in many countries, such as the US, Canada, and others. For example, the National Residency Matching Program (NRMP), an American program for matching medical residents to hospitals, offered in 2015 27,293 positions in 4,012 hospital programs [4].

However, many of the proposed matching mechanisms make assumptions that do not hold in the real world. Many actual settings include partial preferences [5, 6], quotas imposed on matching outcomes [7], distributional constraints [8], and computational constraints, for which compact representations of preferences are useful (e.g., [9, 10]).

In this paper we will focus on a particular problem, arising when the number of options in front of residents is large (as noted above, US residents needed to choose from over 4000 positions, and they apply to only 11, on average! [11]), and they have uncertainty regarding the best hospitals for them. That is, residents may have a vague intuition of which hospitals are better than others[1], but their

[1] There are publicly available rankings, such as the US News and World Report ranking.

own ranking may be influenced by specific, personal considerations (e.g., the personal chemistry with the people in the hospital) . The common way to deal with such a case is to *interview* with a set of hospitals, allowing the resident to figure out their "real" ranking between the possible hospitals. However, this requires the resident to choose the set of hospitals they will interview, based on the limited information they have.

The selection of this interviewing set gives rise to strategic concerns, and widely used mechanisms – which are strategyproof when assuming every resident knows their full ranking of the hospitals – are no longer strategyproof [23, 24]. Every resident will try, naturally, to maximize the outcome according to their own welfare function, but where will we end up, and what will be the resulting stable state? To examine the Nash equilibrium strategies, we will look at a particular matching mechanism – one of the most widely used ones – Resident-Proposing Deferred Acceptance (RP-DA). We will also assume that hospitals have a ranking of residents (e.g., according to their GPA or exam grades, which are used to determine acceptance throughout the world [12, 13, 14, 15]), and, of course, prefer higher ranked residents over lower ones (we do not need to assume residents have this list, but we do assume they know their own ranking).

One possible strategy – indeed, the one that the authors of this paper initially assumed would be a Nash equilibrium in many cases – is an *assortative* one, in which hospitals and residents are stratified: highly ranked residents interview at well-regarded hospitals; medium residents interview at medium hospitals; and low ranked residents interview at low ranked hospitals. Such an equilibrium seems almost "natural" in construct – residents are divided into cohorts, and each cohort interviews in hospitals with equivalent quality. Such an equilibrium is desirable, as it makes sense for residents to interview where there is a good chance they will get in. We know some applicants in matching settings try to interview above their level and fail [16], while if they interview below their level they sell themselves short. Moreover, such a structure, should it exist in Nash equilibria allows for easier analysis and for focusing on only analyzing these cohorts.

Indeed, there is some anecdotal evidence such a strategy is pursued in some matching settings [15]. A different possible strategy (which also has some anecdotal evidence that it is used [16, 17]) is a *reach/safety* one, in which residents apply to hospitals “above” their approximate ranking (i.e., they try to “reach” to better hospitals) as well as hospitals where they are almost guaranteed to get in, so about their ranking or slightly below (i.e., “safe” hospitals).

Of course, while we use the terminology of hospitals and residents, our results hold for any setting in which one side cannot fully rank the other side and has uncertainty regarding its ranking, necessitating a preliminary decision on which options to focus on ranking. This can happen for students interviewing at schools, universities inviting candidates for a job opening[2], or prospective PhD. students choosing which potential advisors to meet.

Our Contribution. We explore the structure of Nash equilibria when hospitals have a joint list ranking of residents, while residents are unsure regarding their rankings of the hospitals, and need to select k hospitals to interview in, which result in a ranking of those k hospitals. We assume residents have a similar utility valuation (i.e., the value they get from getting their 1st choice, 2nd choice, etc.). We focus, in particular, on the case where residents' hospital ranking are sampled from a Mallows' distribution. This means, broadly, that there is a widely shared basic (*ground truth*) ranking σ and a parameter ϕ. The parameter determines, in a sense, the likelihood of choosing a preference order significantly different from σ. We define the Mallow's model in Section 3.3. As can be seen from the definition, as ϕ approaches 1, every possible preference ranking becomes almost equally likely. And as ϕ approaches 0, it becomes more and more probable that the preference ranking will not deviate far from the ground truth ranking.

Beyond several existence proofs for pure equilibria and conditions for assor-

[2]In this case, the candidates may have a shared ranking over universities (e.g., from one of the international rankings), and the universities are the “resident”-equivalents, choosing which potential candidates they wish to invite for a job talk.

tative equilibria, we show that for a very small interviewing set ($k = 2, 3$), there are utility functions for which, as long as the probability of having a wildly different ranking than the ground truth is not very high, equilibria will be assortative. However, if the sets are larger, this no longer holds. Even though when rankings are identical to the ground truth an assortative strategy is a Nash equilibrium, that is not true even if rankings are infinitesimally close to the ground truth, regardless of the utility functions. We also show how the equilibrium looks like for several choices of larger interview sets ($k = 4, 6$), and we are able to observe reach/safety strategies being an equilibrium in some cases, even in cases where agents are close to the ground truth (e.g, the worst resident also applying to the top hospital).

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 et al. [18]), there has been significantly less work on the interviewing problem which we deal with. One research direction looked at interviewing policies that attempt to minimize the number of interviews conducted while ensuring that a stable matching is found. Rastegari et al. [6] showed that while finding the minimal interviewing policy is NP-hard in general, there are special cases where a polynomial-time algorithm exists. They also provide a model for minimal interviewing, and an MDP framework for minimal interviewing (with no fixed quota). Drummond and Boutilier [5] looked at a similar problem, using minimax regret and heuristic approaches for interviewing policies. However, neither of these papers examined strategic issues arising when agents get to choose where they wish to interview.

Several papers [19, 20, 21, 22] addressed a limited/fixed set of interviews, but unlike this paper, they assume uncorrelated preferences (that is, every hospital/resident has its own ranking of residents/hospitals, independent of others[3]),

[3]In addition to the independent, uniform distribution of preferences, Immorlica and Mah-

which we believe is less realistic, though that is also a particular case of our model. In any case, their assumptions allow for truthfulness to be a Nash equilibrium or a highly probable best-response, which is not applicable in our model. Kadam [20] considers the Nash equilibrium both when all preferences are uncorrelated, or exactly the same, but does not allow a small variance (their model also includes hospitals ruling out candidates after interviews, which our model does not support). In two linked papers [23, 24], Haeringer and Klijn [23] showed that when limiting the number of interviews, the Nash equilibria becomes less efficient, and possibly not stable. Then, in an experiment, Calsamiglia et al. [24] show that in various cases, when you limit people's ability to show their full preferences they can become less truthful (for example, increasing the rank of a school they are likely to be able to enter), and their behavior creates various issues for the mechanism. He and Magnac [25] show experimentally the effects of an interview cost (of sorts) in decreasing the matching quality, noting that a low cost both decreases loads as well as maintaing quality.

Motivated by the college admissions problem, Chade et al. [27, 26] 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. That work investigates how students decide where to apply in a decentralized market (as well as the applying students being those with the fixed list). Unlike them, we focus on centralized matching markets which result in stable matchings. Recently, Shorrer [28] analyzed a case in which residents do not know their own ranking, and the rankings of the hospitals is fixed (similar to the setting in Chade et al. [27, 26]). While this setting is quite different from ours (we focus, in a sense, on how varied the ranking of hospitals is between agents, by tweaking the Mallows' distribution ϕ parameter), it is interesting that the results have some similarity to our own:

dian [19] also examines the case in which each hospital has a fixed probability of being selected, and preferences are formed by continuous sampling. Beyhaghi et al. [21] divides both hospitals and residents into two groups, and preferences are uniform within each group.

that without outside options, mixing "safe" and "reach" options makes sense. Indeed, in his case, this structure is a result of players not being sure of their own ranking by others, while for us it is the combination of not knowing one's own preferences, as well as not knowing other people's preferences. So while the safe/reach structure would disappear in the Shorrer [28] model if one knew their own ranking, that is not the case in our setting.

Coles et al. [29] discuss signaling 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 signaling may be harmful [30]. 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 this paper is Lee and Schwarz [31]. 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 arguably more realistic) set of assumptions on the structure and knowledge of preferences. The timeline of events for this game is given in Section 3.1.

3. Model

There are n residents and n hospital programs[4]. The set of residents is denoted by $R = \{r_1, \ldots, r_n\}$; the set of hospital programs is denoted by $H = \{h_1, \ldots, h_n\}$. Both hospitals and residents have (strict) preferences over each other, and 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*, i.e., residents can only do their residency at a single hospital, and hospitals can accept at most one resident.[5] A *matching* is a 1-1 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. We assume that residents prefer to be assigned to any hospital over not being matched, and hospitals prefer to have any resident over not filling the position at all. 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 $r \succ_h \mu(h)$.

Hospitals have identical preferences over all residents, which we call the *master list*, $\succ_{ML}$[6]. Without loss of generality, let $\succ_{ML} = r_1 \succ r_2 \succ \ldots \succ r_n$, where $r_i \succ_{ML} r_j$ means that r_i is preferred to r_j according to $\succ_{ML}$.

Each resident, r, has idiosyncratic preferences over the hospitals[7], which we

---

[4]The assumption that there are an equal number of residents and hospitals is without loss of generality. If there are more residents than hospitals, then the lowest ranked residents will not obtain any interview and can therefore be ignored. If there are more hospitals than residents, we can add "dummy" residents having the lowest ranks and the matching mechanism can ignore the match of any dummy resident.

[5]This is a simplifying assumption that eases understanding of the equations. Generally, our results hold without this assumption as well.

[6]This can be thought of as a list based on grades. As in our model, these are usually known to hospitals, allowing them to rank candidate residents.

[7]This may be based on location, relationship status, etc.

assume are drawn *i.i.d.* from some common distribution D, and that this is common knowledge as well. 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 also assume residents are aware of their own ranking in the master list.

Finally, we will assume (as in Coles and Shorrer [32]) 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 refine their information by conducting a number of *interviews*, not exceeding their interviewing quota 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 $k < n$. Once r_j has finished interviewing, r_j knows their preference ranking over $I(r_j)$. This information is then submitted 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 [1]. This matching is also guaranteed to be unique, as stable matching problems with master lists have unique stable solutions [18]. Thus our results directly hold for any mechanism that returns a stable matching, including hospital-proposing deferred acceptance and the greedy linear-time algorithm [18].

3.1. Description of the Game

We now describe the *Interviewing with a Limited Quota* game. We attempt to formalize this game in a manner consistent with previous literature on interviewing, particularly with Rastegari et al. [6]. The game follows the following timeline.

1. A master list ranking all residents becomes known to all hospitals. Residents know their own ranking.
2. 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 $\succ_{ML}$, where $|I(r)| \leq k$.
3. Each resident r interviews with hospitals in $I(r)$ and learns their own preference over members of $I(r)$.
4. 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. Note that thanks to the DA mechanism's strategyproofness, there is no reason for residents to misreport their preferences once they have learned their personal ranking as they have already strategized in choosing their interview set. This contrasts with the mechansims studied in [23, 24] where students strategize without interviews in choosing their schools.

Example 1. *Suppose $k = 2$ and we have 4 hospitals – h_1, h_2, h_3, h_4, and 4 residents – r_1, r_2, r_3, r_4. All hospitals know the residents' quality (r_1 being the best, followed by r_2, then r_3, and r_4 is the worst), and every resident knows their position in the hospitals' ranking. Suppose residents have two possible rankings of hospitals: with probability 0.5 a resident's ranking of hospitals is $h_1 \succ h_2 \succ h_3 \succ h_4$, and with probability 0.5, it is $h_2 \succ h_1 \succ h_4 \succ h_3$.*

Residents r_1 and r_2 can choose to interview at h_1 and h_2, while residents r_3 and r_4 can choose to interview at hospitals h_3 and h_4. Such a choice is both assortative and stable – both r_1 and r_2 know they will never prefer h_3 and h_4 over the hospitals they interview in; and because of this, both r_3 and r_4 know hospitals h_1 and h_2 will surely be taken already by the time it is their turn to interview, so no point in interviewing there. In this case, there is no other stable choice.

If the probability of any ordering is as likely as any other, then many other *interviewing strategies are stable, including non-assortative ones. For example, r_1 and r_3 interviewing at h_2 and h_4 while r_2 and r_4 interview at h_1 and h_3.*

3.2. Payoff function for Interviewing with a Limited Quota

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 Quota* game. In order for resident r_j to choose their interview set $I(r_j) \subset H$, they have to be able to evaluate the payoff they expect to receive from that choice, where the payoff depends on both the actual preference ranking they expect to draw from distribution 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 $r_i \succ_{ML} r_j$. If $r_j \succ_{ML} r_i$ 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: $u_{r_j}(S) = u_{r_j}(S|D, I(r_1), \dots, I(r_{j-1}))$.

Given fixed interviewing sets $I(r_1), \dots, I(r_{j-1})$, and some partial matching $m = \mu_{|r_1,\dots,r_{j-1}}$, we compute the probability (with respect to the realized preferences of the residents) that matching m happened via RP-DA. Let $m(r_i)$ denote which hospital 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 $\succ_{ML}$). That is, we are looking for all preference orders such that for every hospital resident i interviewed at, those he prefers over the one he got were matched (by m) to residents ranked above him in the master list. Such preference orders are consistent with the matching m. Denote this set as $T(r_i, m)$. Formally, $T(r_i, m) \subseteq H_\succ$ is:

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

Note that $T(r_i, m) \neq \emptyset$ for all r_i and $T(r_1, m)$ is the set of all rankings over hospitals.

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

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

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

Using Equation 1, we 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 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. That is, we define M^* as the set of matchings that given the interview sets of residents $1, \ldots, j-1$, resident j's preference η and their interview set S, assign resident j with hospital h and any hospitals resident j ranks above it are all assigned to resident $1, \ldots, j-1$, which also interviewed them. Formally, the set of such (partial) matchings is:

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

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

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

For readability, we refer to $P(\mu(h) = r_j|\eta, S, I(r_1), \ldots, I(r_{j-1}))$ as $P(\mu(h) = r_j|\eta, S)$. We now have all 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:

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

Intuitively, what the payoff function in Equation 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, $r_1, \ldots, r_{j-1}$.

Example 2. *Suppose there are two residents, r_1 and r_2, each of whom interviewed with hospitals h_1 and h_2. Resident r_1 will be matched with whomever they most prefer, while r_2 will be assigned the other. The probability that r_2 will be assigned h_1 is 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, as 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 and results under particular assumptions and constraints 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 [33]) is characterized by a reference ranking σ, and a dispersion parameter $\phi \in (0, 1]$, 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 (for $i \in \{1, \ldots, n\}$, $a_i \in A$ indicates the alternative's rank in σ). The probability of any given ranking η is:

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

Here d is Kendall's τ distance metric[8], and Z is a normalizing factor; $Z = \sum_{\eta \in A_{\succ}} \phi^{d(\eta,\sigma)} = (1)(1+\phi)(1+\phi+\phi^2)\ldots(1+\ldots+\phi^{|A|-1})$ [34].

As $\phi \to 0$, the distribution approaches drawing the reference ranking σ with probability 1; when $\phi = 1$, the Mallows distribution is equivalent to drawing from the uniform distribution. Hence, ϕ marks, in a sense, the likelihood of choosing a preference order significantly different from σ. The Mallows model (and mixtures of Mallows) have plausible psychometric motivations and are commonly used in machine learning [35, 36, 34]. Mallows models have also been used in previous investigations of preference elicitation schemes for stable matching problems as in Drummond and Boutilier [37, 5].

Intuitively, a Mallows model can be iteratively generated by repeated insertions of alternatives in a growing preference set, where the particular insertion point 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. Additionally, this repeated insertion procedure can be used to determine the probability 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. These insights on the Mallows model, which to the best of our knowledge have not been previously stated, are captured in the following results. The proofs appear in the appendix.

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 ordering n agents, in which $a_i \succ a_j$ $(1 \leq i, j \leq n)$, the probability that a ranking η is drawn from $D^{\phi,\sigma}$ such that $a_i \succ_\eta a_j$ is equal to*

[8]The Kendall τ distance between two ordering of m items is the number of pairwise disagreements between them (e.g. in one $a \succ b$ and in the other $b \succ a$).

drawing from some distribution $D^{\phi,\sigma'}$ where σ is a suffix or prefix of σ' (that is, there is σ, an ordering of n agents, and σ', an ordering of n' agents ($n' > n$), and σ' can be divided into σ, an ordering of the first/last n agents, and an ordering of the last/first $n' - n$ agents).

In particular, we instantiate Lemma 1 to the case where two alternatives are adjacent to each other in the original ranking, a_i and a_{i+1}.

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

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

We similarly extend Corollary 1 to include three consecutive items.

Corollary 2. *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 2. *The probability that a_1 will be ranked in place j is $\frac{\phi^{j-1}}{1+\phi+\ldots+\phi^{n-1}}$. Furthermore, 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 $\frac{\phi^{j-1}}{1+\phi+\ldots+\phi^{n-1}}$.*

It is possible to bound the probability that any two alternatives will be "out of order" in any given ranking;

Lemma 3. *Let $\eta \in D^{\phi,\sigma}$ be such that $a_j \succ_\eta a_i$ for some $i < j$, then $P(\eta) < \frac{\phi^{j-i}}{Z}$.*

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

Observation 1. *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. Equilibria for Interviewing Markets with General Preferences and Master Lists

We provide an equilibrium analysis for the *Interviewing with a Limited Quota Game*. We first show that a pure strategy equilibrium for this game always exists, even under arbitrary distributions and scoring functions. We further explore our model by assuming that preference rankings are drawn from the ϕ-Mallows model, and we then analyze when and how the Mallows parameter ϕ, different scoring functions, and quota sizes k, support assortative interviewing

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 Quota game*, and show that a pure strategy equilibrium always exists.

Theorem 1. *A pure strategy equilibrium always exists for the Interviewing with a Limited Quota 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 $r_1, \ldots, r_{j-1}$. As r_j knows that each other resident r_i is drawing from distribution D *i.i.d.*, they can calculate $r_1, \ldots, r_{j-1}$'s expected utility maximizing interview set, using Equation 3. Their payoff function depends only on D and $I(r_1), \ldots, I(r_{j-1})$, all of which they now have. They then calculate the expected payoff for each $\binom{n}{k}$ potential interviewing sets, and interview with the one that maximizes their expected utility. $\square$

Note that this game is sequential in nature: each resident r_j's best response only depends on the $j-1$ agents that are ranked higher than them 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 does not affect r_j at all. We then continue solving for the best strategy by using iterated deletion of dominated strategies.

Let us number the hospitals such that r_1's best response is to interview with $I(r_1) = \{h_1, \ldots, h_k\}$ (and subsequently, $h_{k+1}, \ldots, h_{2k}$ are those that r_1 would choose should $h_1, \ldots, h_k$ be gone). Note that this numbering of hospitals is determined by the scoring function and the distribution D from which all residents draw i.i.d. their actual preferences. Note also that in the absence of other residents, every r_i would have the same ordering of all hospitals (we will assume there is a consistent way to break ties so that the ordering is a total ordering). For a Mallows model, this would be the reference ranking σ. Knowing that r_1 will be preferred over all other residents by every hospital, the known interviewing set for r_1 will then eliminate many strategies for r_2, which in turn eliminates 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 1 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. The only known way to calculate it is directly (brute-force), which is computationally infeasible

We are interested in understanding whether and when a particular class of natural interviewing strategies form an equilibrium. In particular, if residents have interviewing quotas of size k, we ask the question *Will residents interview assortatively?* And if not, what different strategies will they pursue?

Definition 1. *When residents have a quota of k interviews, we say that an interviewing strategy profile is* assortative *iff for $j = 0, 1, 2, \ldots, \frac{n}{k} - 1$, each resident $r \in \{r_{jk+1}, \ldots, r_{jk+k}\}$ chooses to interview with the set of k hospitals*

$\{h_{jk+1}, \ldots, h_{jk+k}\}$.[9]

This is a very strict definition of assortative interviewing that one might initially believe should hold when one side (i.e., the hospitals) have a master list and there is some publicly known reference ranking. A weaker concept might be called *weakly assortative interviewing* or consecutive interviewing, where every resident r_i will interview with some "appropriate tier" of hospitals $\{h_j, \ldots, h_{j+k-1}\}$ for some j and furthermore that j is "close" to i. Our theoretical results will only consider the strict definition in Definition 1. However, our computational studies will also show that residents will often deviate even from weak assortative interviewing.

We begin our theoretical analysis by deriving conditions that ensure assortative interviewing. Namely, we will prove that we can focus mainly on the behaviour of r_k (k being the interviewing quota), under some conditions. We do so by showing that if assortative interviewing is a best response for resident r_k if all other residents $i < k$ interview assortatively, then assortative interviewing is a best response for *every* resident r_i ($i < k$) 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 the first k residents (and is thus an equilibrium for them in this game).

Proposition 1. *Consider an interviewing quota 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 $r_1, \ldots, r_k$ to interview assortatively. Moreover, this then forms a unique equilibrium for $r_1, \ldots, r_k$ in this setting.*

[9] We are assuming for convenience that k divides n. When k does not divide n, there will be some remaining $k' < k$ residents that will interview with the remaining k' hospitals: $h_{\lfloor \frac{n}{k} \rfloor k+1}, \ldots, h_n$.

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 their choice (i.e., $r_1, \ldots, r_{i-1}$ have not been allocated h), and is their 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:

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

Since for r_1, it is always true that $b^1(h, \eta) = 1$ for any desired h (since r_1 goes first, no $h \in H$ has been allocated by another $r \in R$), suppose it will interview in a set of k hospitals $\{h_1, \ldots, h_k\}$ (the numbering according to r_1's choices as determined by the distribution D). We are concerned with the best response strategy of r_k which only depends on the strategies of r_i for $i < k$. Suppose there is no assortative equilibrium, and let $r_i, i < k$, be the resident with the lowest index for which it is better off interviewing in set $S' \neq \{h_1, \ldots, h_k\}$. Then $b^i(h, \eta) \geq b^k(h, \eta)$, with the inequality being strict for some $h \in \{h_1, \ldots, h_k\}$. Note that for any $h \notin \{h_1, \ldots, h_k\}$, $b^i(h, \eta) = 1$.

Hence, if $u_{r_i}(\{h_1, \ldots, h_k\}) < u_{r_i}(S')$, this means if all agents $r_1, \ldots, r_{k-1}$ are being assortative (so $b^k(h, \eta) = 1 = b^i(h, \eta)$ for $h \in S' \setminus \{h_1, \ldots, h_k\}$), $u_{r_k}(\{h_1, \ldots, h_k\}) < u_{r_k}(S')$. That is, if it is not beneficial for r_i to be assortative, it would not be beneficial for r_k to be assortative if $r_1, \ldots, r_{k-1}$ are assortative.

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

If the proposition applies to the first k residents (and hospitals), this means all of the hospitals $h_1, \ldots, h_k$ are occupied by one of residents $r_1, \ldots, r_k$. We can simply remove these hospitals and residents, and ask ourselves if the distribution is such that Proposition 1 applies to residents $r_{k+1}, \ldots r_{2k}$ and hospitals $h_{k+1}, \ldots, h_{2k}$. If it does, we continue inductively, checking for if the proposition applies for each batch of k hospitals and residents. If it does, each group of k residents interviews assortatively. Being able to examine only the first k agents allows us to simplify the notations of our proofs.

4.2. Interviewing Equilibria Under Mallows Models with Master Lists

Having proven the existence of a pure strategy equilibria for the interviewing game, we instantiate the distribution from which residents are drawing their preferences to a Mallows model, to gain a deeper understanding of the results (and their limitations). In particular, we provide a characterization for when assortative interviewing forms an equilibrium without imposing any particular additional restrictions on the utility functions of the residents.

Before proving our main result in Theorem 2, 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.[10] 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 Quota game with quota k, a Mallows model with dispersion parameter ϕ, and a scoring function v. We show that if, for resident r_k, replacing any alternative $h_j \in \{h_1, \ldots, h_k\}$ with alternative h_{k+1} is *not* an improvement to their expected utility, then interviewing with $\{h_1, \ldots, h_k\}$ is their best response when they draw their 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 4. *Given an Interviewing with a Limited Quota game $\Psi = \langle k, \phi, v\rangle$, if resident r_k's expected payoff from interviewing with hospitals $\{h_1, \ldots, h_k\}$ (when*

[10]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.

residents $r_1, \ldots, r_{k-1}$ have interviewed with them as well) is higher than their expected payoff from interviewing with hospitals $\{h_1, \ldots, h_{k+1}\} \setminus \{h_j\}$ for all $j \in \{h_1, \ldots, h_k\}$, then resident r_k's best response is to interview with $\{h_1, \ldots, h_k\}$ (i.e., assortatively).

Proof. The idea behind the proof is that if there is a set of hospitals that are better than interviewing assortatively, since no other resident prior to r_k interviews there, the hospitals in this set that are outside of $\{h_1, \ldots, h_k\}$ have an ordering. That is, the expected utility from adding h_{k+1} is larger than that of adding h_{k+2}, since, in expectation h_{k+1} is likely to be ranked higher by the resident than h_{k+2}. Therefore, taking out the hospital with the least expected utility from $\{h_1, \ldots, h_k\}$ and adding h_{k+1} in its stead should already be beneficial, since any other hospital added to the interviewing set will remove a hospital with a higher utility (than the one removed for h_{k+1}), and replace it with lesser utility hospital (since an hospital from $h_{k+2}, \ldots, h_n$ has smaller expected utility). Therefore, if there is a set that is better than assortative, it should show up already when replacing some hospital in $\{h_1, \ldots, h_k\}$ by h_{k+1}.

More formally, following the proof in Proposition 1, 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 $S = \{h_1, \ldots, h_k\}$ 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 $r_1, \ldots, r_{k-1}$, 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 $P(\eta, D^{\phi,\sigma})$ and $v(h, \eta)$ are known a priori. Moreover, when $r_1, \ldots, r_k$ all interview with the same k hospitals, $b(h, \eta)$ is equivalent to the probability that hospital h is available for r_k (which we denote by $P(h \text{ avail})$): resident r_k gets whatever hospital

$r_1, \ldots, r_{k-1}$ 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 $h'_1, \ldots, h'_l$. Also, note that h_{k+1} must be in $S' \setminus S$, as $\bar{S} \neq S$ and h_{k+1} dominates all alternatives in $\{h_{k+1}, \ldots, h_n\}$: h_{k+1} is available for r_k with probability 1 (as are 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

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

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 interviewing in h_{k+1} instead of h'_1 is formally: $z_1 = \sum_{\eta \in H_\succ} P(\eta | D^{\phi,\sigma}) b(h_{k+1}, \eta) v(h_{k+1}, \eta)$. 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; Otherwise, we assume $z_1 - y_1 < 0$ and establish a contradiction.

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 1:

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

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:

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

and therefore $y_2 > y_1$. Thus:

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

Note that due to similar considerations, 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 $r_1, \dots, r_{k-1}$ are all matched to different alternatives). As we assume residents $r_1, \dots, r_{k-1}$ interview assortatively, only one of $\{h_1, \dots, h_k\}$ will be available.

Lemma 5. *Given an Interviewing with a Limited Quota game $\Psi = \langle k, \phi, v\rangle$, if residents $r_1, \ldots, r_{k-1}$ all interview assortatively (i.e., with hospital set $S = \{h_1, \ldots, h_k\}$), then assortative interviewing is a best response for resident r_k if and only if the following inequality is satisfied for all $h_j \in \{h_1, \ldots, h_k\}$ when $S' = S \setminus \{h_j\} \cup \{h_{k+1}\}$:*

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

where

$$\chi(h_i \succ_\eta h_j) = \begin{cases} 1, & \textit{if } h_i \succ_\eta h_j \\ 0, & \textit{otherwise} \end{cases}$$

Proof. By Lemma 4, 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. 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} :

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

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

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

We now 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., $r_1, \ldots, r_{k-1}$ have all been matched with $\{h_1, \ldots, h_k\} \setminus \{h_j\}$), or for some $h_i \in \{h_1, \ldots, h_k\} \setminus \{h_j\}$, 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 Equation 4 using the indicator function, and simplify:

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

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

Combining the simplifications provided in Equations 5 and 6 completes the proof. □

By combining the lemmas, we show that we need only check k interviewing sets for resident r_k to prove that assortative interviewing forms an equilibrium for this game.

Theorem 2. *Given an Interviewing with a Limited Quota game* $\Psi = \langle k, \phi, v \rangle$, *then satisfying the inequality found in Lemma 5 for all* $h_j \in \{h_1, \ldots, h_k\}$ *is both sufficient and necessary to show that all residents interviewing assortatively form an equilibrium for this game.*

Proof. For the first k residents, this follows directly from combining Proposition 1 and Lemma 5. The theorem would be correct if we could apply this proposition and lemma iteratively, one group of k hospitals and residents at a time. Thanks to the Mallows distribution's properties, we can: If r_k's best response was assortative, we know that all the residents $r_1, \ldots, r_k$ interviewed assortatively, thus all hospitals $h_1, \ldots, h_k$ are taken. This means that the same equations that told us that r_k's best response (to $r_1, \ldots, r_{k-1}$) was assortative tell us that r_{2k}'s best response (to $r_{k+1}, \ldots, r_{2k-1}$) is assortative: Since a switch between h_1 and h_2 has the same probability as switching between h_{k+1} and h_{k+2}, if Proposition 1 and Lemma 5 can be applied once on hospitals and residents $1, \ldots, k$, they can be applied again for $k+1, \ldots, 2k$, as all equations remain the same, due to the practical "disappearance" of the hospitals $h_1, \ldots, h_k$ for agents $r_{k+1}, \ldots, r_{2k}$ (thus their order can be ignored). Now that we have shown that the first two groups of k residents interview assortatively, we can use the same argument iteratively for the next k residents, and so on. □

We provide a simplified condition for assortative interviewing that is sufficient though not necessary. This condition is easier to compute than the condition in Lemma 5, and thus will be valuable later on, when verifying whether specific valuation functions admit assortative interviewing equilibria.

Lemma 6. *Given an interviewing quota of k interviews, a dispersion parameter ϕ, and a scoring function v, if residents $r_1, \ldots, r_{k-1}$ all interview assortatively (i.e., with hospital set $S = \{h_1, \ldots, h_k\}$), then satisfying the following inequality for all $h_j \in \{h_1, \ldots, h_k\}$ 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 5. 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, $S' = \{h_1, \ldots, h_{k+1}\} \setminus \{h_j\}$. Thus, we note that:

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

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 Equation 9 is an upper bound, it is sufficient to show the following, as required:

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

□

While we have focused on the existence of assortative interviewing, we 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. As first noted by Lee and Schwarz [31] under a different model, when residents and hospitals are divided into n/k subsets and matched inside these subsets, this also forms an equilibrium.

Example 3. *Suppose residents draw their preferences i.i.d. from the uniform distribution. 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. This allows us to use the condition in Lemma 5. When r_1 makes their choice, they are indifferent between any alternatives, as they have equal likelihood (probability 1) to get any of them. Suppose they chooses $h_1, \ldots, h_k$. Then, r_2 prefers $h_{k+1}, \ldots, h_{2k}$, as they are indifferent between any alternatives that r_1 has* not *interviewed with. This continues until all hospitals have one interviewer. Once again, the next resident is indifferent between all alternatives, and can choose to interview like r_1, the next one like r_2, and so on. This results that for $j \leq k$, any resident r_{ik+j} interviews with hospitals $\{h_{(i-1)k+1}, \ldots, h_{ik}\}$.*

5. Assortative Equilibria for Small Interviewing Quotas

We now discuss assortative equilibria when participants' interviewing quota 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 consider three different scoring rules, inspired by the social choice literature [38, 39, 40], in order to better ascertain the effect of resident utility-structure on assortative equilibria.

Let s_i be the i^{th} ranked hospital in a residents ranking η. The first function we consider is *plurality-based*, where $v(s_1) = 1$ and $v(s_i) = 0$ for all $i > 1$.[11] 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, the residents' utility drops linearly in proportion to the rank of the alternative to 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 (Coles and Shorrer [32] also examine such a scoring function).

[11] 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.

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; $v(s_i) = (\frac{\epsilon}{2})^{i-1}$, $0 < \epsilon < 1$.

These scoring rules can be viewed as conveying something about the utility of residents. For example, plurality indicates residents only want to get their top choice and do not care for anything else. Borda, on the other hand, indicates a gradual, linear, decrease in the value of the hospital. Similar framing can be done for any scoring rule. Such a framing implies, of course, how willing would a resident be to take a risk, considering they may end up with less-desirable choice. If the difference between getting one's top choice and getting the least-desirable choice is miniscule, one might be willing to take a far-fetched chance on getting the top choice, since the potential damage is small [41].

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

Lemma 7. *A necessary and sufficient condition for assortative interviewing under plurality is:*

$$P(h_j \ avail) \geq \phi^{k-j+1} \tag{10}$$

Proof. We begin with the condition in Lemma 5:

$$P(h_j \text{ avail})\mathbb{E}(v(h_j)|D^{\phi,\sigma}) > \tag{11}$$

$$P(h_j \text{ avail})\mathbb{E}(v(h_{k+1})|D^{\phi,\sigma}) + \sum_{\eta \in H_\succ} P(\eta|D^{\phi,\sigma}) \cdot \big[\sum_{h_i \in S'} P(h_i \text{ avail})\chi(h_{k+1} \succ_\eta h_i)v(h_{k+1}, \eta)\big] \tag{12}$$

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 simplify Equation 12:

$$P(h_j \text{ avail})\mathbb{E}(v(h_{k+1})|D^{\phi,\sigma}) + \sum_{i=1}^{j-1} P(h_i \text{ avail})\mathbb{E}(v(h_{k+1})|D^{\phi,\sigma}) + \sum_{i=j+1}^{k} P(h_i \text{ avail})\mathbb{E}(v(h_{k+1})|D^{\phi,\sigma}) = \sum_{i=1}^{k} P(h_i \text{ avail})\mathbb{E}(v(h_{k+1})|D^{\phi,\sigma}) \tag{13}$$

But, again, as the expected value for any hospital h is simply the probability that h 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 $r_1, \ldots, r_{k-1}$ have been matched with exactly $k-1$ hospitals in $h_1, \ldots, h_k$, leaving exactly one hospital left with probability 1.

Applying Lemma 2 to both sides of the inequality (recall that $\mathbb{E}(v(h_j)|D^{\phi,\sigma})$ is simply $P(h_j \text{ in } s_1)$):

$$\begin{aligned} P(h_j \text{ avail}) \frac{\phi^{j-1}}{1 + \ldots + \phi^{n-1}} \geq & \frac{\phi^k}{1 + \ldots + \phi^{n-1}} \\ P(h_j \text{ avail}) \geq & \phi^{k-j+1} \end{aligned} \tag{14}$$

□

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 8. *If for a given interviewer quota k and dispersion parameter ϕ, the condition of Lemma 7 is satisfied for a plurality valuation function with strict inequality, then there exist exponential valuations under which assortative interviewing is an equilibrium.*

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

Proof. Looking at the condition of Lemma 5

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

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

$$P(h_j \text{ avail}) \sum_{i=1}^{n} P(h_j \text{ in } s_i) v(s_i) \tag{15}$$

$$\geq P(h_j \text{ avail}) \sum_{i=1}^{n} P(h_{k+1} \text{ in } s_i) v(s_i)$$

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

$$+ \ldots + \sum_{\substack{\eta \in H_{\succ} | \\ h_{k+1} \text{ in } s_{n-1}}} P(\eta | D^{\phi,\sigma}) \sum_{h_i \in S'} P(h_i \text{ avail}) \chi(h_{k+1} \succ_\eta h_i) v(s_{n-1}) \tag{16}$$

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 Equation 16 and Lemma 5, it is sufficient to show the following holds whenever plurality admits an assortative interviewing equilibrium:

$$P(h_j \text{ avail}) P(h_j \text{ in } s_1) v(s_1) \geq P(h_j \text{ avail}) P(h_{k+1} \text{ in } s_1) v(s_1) + \sum_{\ell=2}^{n} v(s_\ell) \tag{17}$$

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 Equation 17, 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 two 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 3. *Using plurality as the residents' scoring function and a quota 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 7 for h_1. 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 their preferences s.t. $h_2 \succ h_1$. Then, by Corollary 1, $P(h_1 \text{ avail}) = \frac{\phi}{1+\phi}$, implying we need to satisfy the equation $\frac{\phi}{1+\phi} \geq \phi^2$, which is true whenever $0 < \phi \leq 0.6180$. Doing the same for h_2 provies a bound of $0 < \phi \leq 0.7549$, so we take the tighter bound of 0.618. □

Though we do not formally state it, combining Theorem 3 and Lemma 8 shows that for exponential scoring functions, 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 4. *Given Borda as residents' scoring function and a quota of* $k = 2$ *interviews, for a Mallows model dispersion parameter* ϕ *such that* $0 < \phi \leq 0.265074$*, assortative interviewing forms an equilibrium.*

Proof. We begin by noting that, because of Lemma 1, 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 4, 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 only over 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 of discordant pairs 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 1, $P(\mu(r_1) = h_1) = \frac{1}{1+\phi}$. Using this information, the payoff function, and the definition of η, η', we find a lower bound:

$$U_1 = \sum_{\eta \in P(H)^{h_2 \succ h_1 \succ h_3}} P_{\mu(r_1)}(h_2) \Big[(v(h_1, \eta) - v(h_3, \eta))P(\eta | D^{\phi,\sigma}) + (v(h_1, \eta') - v(h_3, \eta'))P(\eta' | D^{\phi,\sigma})\Big]$$

$$U_1 \geq P_{\mu(r_1)}(h_2)(1)(1-\phi)P(h_2 \succ h_1 \succ h_3) = \left(\frac{\phi}{1+\phi}\right)\left(\frac{\phi}{(1+\phi)(1+\phi+\phi^2)}\right)(1-\phi) \tag{18}$$

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:

$$U_2 = \sum_{\eta \in P(H)^{h_1 \succ h_2 \succ h_3}} P_{\mu(r_1)}(h_1) \Big[(0)P(\eta | D^{\phi,\sigma}) + (v(h_2, \eta') - v(h_3, \eta'))P(\eta' | D^{\phi,\sigma})\Big]$$
$$+ P_{\mu(r_1)}(h_2) \Big[(v(h_1, \eta) - v(h_3, \eta))P(\eta | D^{\phi,\sigma}) + (v(h_1, \eta') - v(h_3, \eta'))P(\eta' | D^{\phi,\sigma})\Big]$$
$$U_2 \geq P(h_1 \succ h_2 \succ h_3)\frac{2}{1+\phi}(\phi - \phi^3 - \phi^4) \tag{19}$$

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:

$$U_{3b} = \sum_{\eta \in P(H)^{h_1 \succ h_3 \succ h_2}} P_{\mu(r_1)}(h_2)[(v(h_1, \eta) - v(h_3, \eta))\frac{\phi^{d(\eta,\sigma)}}{Z} + (v(h_3, \eta) - v(h_1, \eta))\frac{\phi^{d(\eta,\sigma)+1}}{Z}$$
$$U_{3b} \geq \frac{\phi}{\phi+1}(1-\phi)P(h_1 \succ h_3 \succ h_2) \tag{20}$$

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:

$$U_{3a} \geq P_{\mu(r_1)}(h_1) \frac{-\phi}{(1+\phi)(1+\phi+\phi^2)} \left(\frac{\phi}{(1-\phi)^4} + \frac{1}{3(1-\phi)^3} + \frac{2}{3} \right)(1+\phi) \quad (21)$$

We have considered all cases, and can now combine them together. We add the bounds for U_1 (Equation 18), U_2 (Equation 19), U_{3a} (Equation 21), and U_{3b} (Equation 20). We simplify using Corollaries 1 and 2, giving us:

$$\begin{aligned} u_{r_2}(\{h_1,h_2\}) - u_{r_2}(\{h_2,h_3\}) \geq & \frac{\phi^2}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}(1-\phi) + \frac{2(\phi-\phi^3-\phi^4)}{(1+\phi)(1+\phi)(1+\phi+\phi^2)} \\ & - \frac{\phi}{(1+\phi)(1+\phi)(1+\phi+\phi^2)} \left(\frac{\phi}{(1-\phi)^4} + \frac{1}{3(1-\phi)^3} + \frac{2}{3} \right)(1+\phi) \\ & + \frac{\phi^2}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}(1-\phi) \end{aligned} \quad (22)$$

Thus, Equation 22 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 Equation 22, 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:

$$u_{r_2}(\{h_1,h_2\}) - u_{r_2}(\{h_1,h_3\}) \geq \frac{1}{(1+\phi)(1+\phi+\phi^2)} \left[1+\phi-2\phi^2-2\phi^3-2\phi^3 \left(\frac{\phi}{(1-\phi)^4} + \frac{1}{3(1-\phi)^3} + \frac{2}{3} \right) \right] \quad (23)$$

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 1, 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 r_3 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, and hence assortative interviewing cannot be an equilibrium.

Theorem 5. *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 4, 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 $v(s_i) = 5 - i$.

$$\mathbb{E}(u(h_1, S)) = \sum_{\eta \in H_\succ} v(h_1, \eta) P(\mu(h_1) = r_3 | S, \eta, D^{\phi,\sigma}) P(\eta | D^{\phi,\sigma}) \tag{24}$$

$$\mathbb{E}(u(h_4, S')) = \sum_{\eta \in H_\succ} v(h_4, \eta) P(\mu(h_4) = r_3 | S', \eta, D^{\phi,\sigma}) P(\eta | D^{\phi,\sigma}) \tag{25}$$

As before, Equation 24 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 2, we know that $P(h_1 \text{ in } s_i) =$

$\frac{\phi^{i-1}}{1+\phi+\phi^2+\phi^3}$, giving:

$$\mathbb{E}(u(h_1, S)) = P(h_1 \text{ avail})\mathbb{E}(v(h_1)|D^{\phi,\sigma}) = P(h_1 \text{ avail})\frac{4 + 3\phi + 2\phi^2 + \phi^3}{1 + \phi + \phi^2 + \phi^3} \quad (26)$$

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 2 again and an analogous approach as above, we adapt Equation 25:

$$\begin{aligned}
\mathbb{E}(u(h_4, S')) &= \sum_{i=1}^{4} v(s_i)P(h_4 \text{ in } s_i)P(\mu(r_3) = h_4|h_4 \text{ in } s_i) \\
&= \frac{4\phi^3}{1 + \phi + \phi^2 + \phi^3} \\
&\quad + \frac{3}{Z}\left(\phi^2 + \phi^3 + P(h_2 \text{ taken})(\phi^3 + \phi^4) + P(h_3 \text{ taken})(\phi^4 + \phi^5)\right) \\
&\quad + \frac{2}{Z}\left(P(h_2 \text{ taken})(\phi + \phi^2) + P(h_3 \text{ taken})(\phi^2 + \phi^3) + P(h_1 \text{ avail})(\phi^3 + \phi^4)\right) \\
&\quad + \frac{P(h_1 \text{ avail})}{1 + \phi + \phi^2 + \phi^3} \qquad (27)
\end{aligned}$$

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:

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

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:

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

By combining the equations for the probabilities we are left with two equations depending only on ϕ. Moreover, after instantiating $\mathbb{E}(u(h_1, S))$ and $\mathbb{E}(u(h_4, S')$ above, we note that both functions are continuous on the interval $(0, 1]$. Using numerical 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 show that assortative interviewing is an equilibrium for plurality (and thus exponential scoring functions) for $k = 3$:

Theorem 6. *Given an interviewing quota 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 Equation 10 from Lemma 7 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:

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

The first parenthesis is using Corollary 2, and the second the probability h_3 is preferred over h_2 using Corollary 1. Using numerical methods to find the roots of $P(h_1 \text{ avail}) - \phi^3$, we can show that Equation 10 holds when $0 < \phi \leq 0.5462$. □

This means that for a small reviewing set, interviewing is assortative not only when residents know their probability of having a ranking much different than $h_1, h_2, \ldots, h_n$ is low, but also when it approaches quite significant numbers (probability of being exactly truthful are smaller than $(\frac{2}{3})^n$). Of course, because of the small interview size, the probability of actually interviewing in those in which one is different from the ground truth is smaller than with a larger interviewing set. But the larger ϕ means it is quite a significant likelihood.

6. Assortative Equilibria for Large Interviewing Quotas

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.

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 their 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 9. *Given a Mallows model with dispersion parameter ϕ, assortative interviewing for residents $r_1, \ldots, r_{k-1}$, and a hospital $h_i \in \{h_1, \ldots, h_k\}$ (i.e., the residents' interview set), then any profile $\eta_1, \ldots, \eta_{k-1} \in D^{\phi,\sigma}$ of $k-1$ preferences (for $r_1, \ldots, r_{k-1}$) such that h_i is available for r_k has probability $P(r_1 = \eta_1, r_2 = \eta_2, \ldots, r_{k-1} = \eta_{k-1}) \geq \frac{\phi^\gamma}{Z^{k-1}}$, 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 $r'_{i+1}, \ldots, r'_k$ with preference orders $\eta_{i+1}, \ldots, \eta_k \in D^{\phi,\sigma}$ such that they were assigned hospitals $h_{i+1}, \ldots, h_k$. Hence, at the very least, $h_{i+1} \succ_{\eta_{i+1}} h_i, \ldots, h_k \succ_{\eta_k} h_i$. According to Lemma 3, the probability for each of these events is at most $\frac{\phi}{Z}, \ldots, \frac{\phi^{k-i}}{Z}$ (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 monotonic valuation functions will also not admit assortative interviewing equilibria. In some sense, intuitively, because plurality only provides a payoff when residents get their most preferred alternative, this benefits assortative interviewing: everyone wants a chance at the alternatives with the highest probability of being first in their 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 also be less under any other scoring rule.

Theorem 7. *Fix an instance of the Interviewing with a Limited Quota game* $\Psi = \langle k, \phi, plurality \rangle$. *If hospital* h_1 *causes the condition in Lemma 5 to be falsified (i.e.,* $\{h_2, \ldots, h_{k+1}\}$ *has a better expected payoff than* $\{h_1, \ldots, h_k\}$*), then for* k *and* ϕ *(the Mallows' model parameter), assortative interviewing is not an equilibrium for any valuation function.*

Proof. Looking at the condition of Lemma 5 (recall $S' = \{h_1, \ldots, h_k\} \setminus \{h_j\} \cup \{h_{k+1}\}$ for any $h_j \in \{h_1, \ldots, h_k\}$)

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

We again begin by expanding the value expectation ($\mathbb{E}$), as we did in Equation 16: This can be divided into n different inequalities:

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

We shall show that under the theorem's assumptions, none of these inequalities hold for h_1, and therefore the general inequality (Lemma 5) 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 the theorem, since we are using plurality, is that the 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 1 (end of Section 3), 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 1, for any $\ell > k$, $P(h_1 \text{ in } s_\ell) < P(h_{k+1} \text{ in } s_\ell)$ which gives us:

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

Starting with the assumption that assortative interviewing does not hold for plurality, we show that none of the inequalities above hold for any slot s_ℓ, and therefore that the condition in Lemma 5 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 their 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: 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.We then continue to show that for $k > 4$ and ϕ sufficiently small, assortative interviewing is not an equilibrium.

Theorem 8. *Given an interviewing quota of $k = 4$ interviews and any scoring function, assortative interviewing is not an equilibrium for any dispersion parameter $0 < \phi < 1$.*

Proof. By Theorem 7, if assortative interviewing is not an equilibrium for plurality due to h_1, it is never an equilibrium for any scoring rule. As noted before, Equation 10 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_5.

To calculate $P(h_1 \text{ avail})$, we simply iterate over all six 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 permutation of

h_2, h_3, h_4 be denoted as (a_1, a_2, a_3), and let A be the set of all such permutations (i.e., $(a_1, a_2, a_3) \in A$ is a particular permutation of h_2, h_3, h_4).

$P(h_1 \text{ avail}) =$

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

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

Proof. Due to Theorem 7, it is enough to show there is no assortative equilibrium under plurality (and that h_1 violates Lemma 5's condition). We use the simplification from Lemma 7: $P(h_j \text{ avail}) \geq \phi^{k-j+1}$, and we will show it does not hold. Appealing to Lemma 9, we know $P(h_j \text{ avail})$ is of the form:

$$P(h_j \text{ avail}) = \frac{X(k)}{Z^{k-1}}\phi^{\sum_{i=1}^{k-j} i} + \frac{X^1(k)}{Z^{k-1}}\phi^{1+\sum_{i=1}^{k-j} i} + \ldots + \frac{X^\ell(k)}{Z^{k-1}}\phi^{(k\sum_{i=1}^{k-j} i)-1} + \frac{1}{Z^{k-1}}\phi^{k\sum_{i=1}^{k-j} i} \quad (28)$$

$(X(k), X^1(k), \ldots, X^\ell(k)$ are functions that calculate the number of different sets of possible preference orders for $r_1, \ldots, r_k$, with each set being a particular distance from the ground truth σ, thus having the 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 $\phi \to 0$, $Z^{k-1} \to 1$, Equation 28 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 $k > 3$:

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

Contradicting our condition (Equation 10). □

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 $r_1, \dots, r_{k-1}$ are all most likely to be matched with hospitals $h_1, \dots, h_{k-1}$. 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 Interviewing Quota

Our analysis has shown that assortative interviewing equilibria are not the norm 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. In this section we empirically explore small interview quotas to better illustrate the impact of the Mallows model dispersion parameter on equilibria structure.

Consider the case for $k = 2$ interviews where (for the Borda scoring rule) 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 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

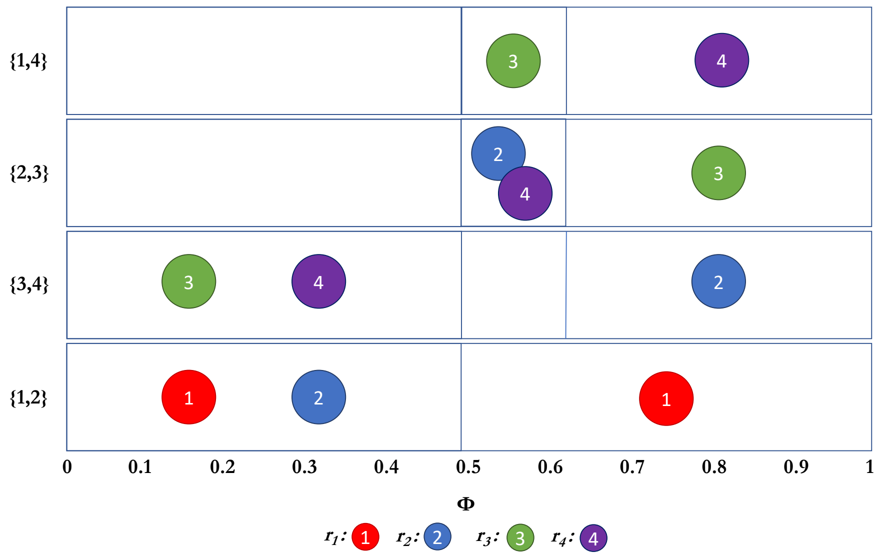

Figure 1: Interviewing sets of residents as a function of ϕ when using the Borda scoring function, with 4 participants, and interview set size of 2.

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 further 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 are available. The desire to avoid interviewing hospitals that are already chosen by many other residents 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 expand on these results and now consider the case of $n = 6$ residents with $k = 2, 3$ and 4 interviews per resident. Here we see in Figures 2,3, and 4, similar equilibrium strategies as for the $n = 4, k = 2$ case. For $k = 2$, and $\phi \leq 0.4$, we again see that assortative interviewing is an equilibrium. When $\phi = 0.5$, we

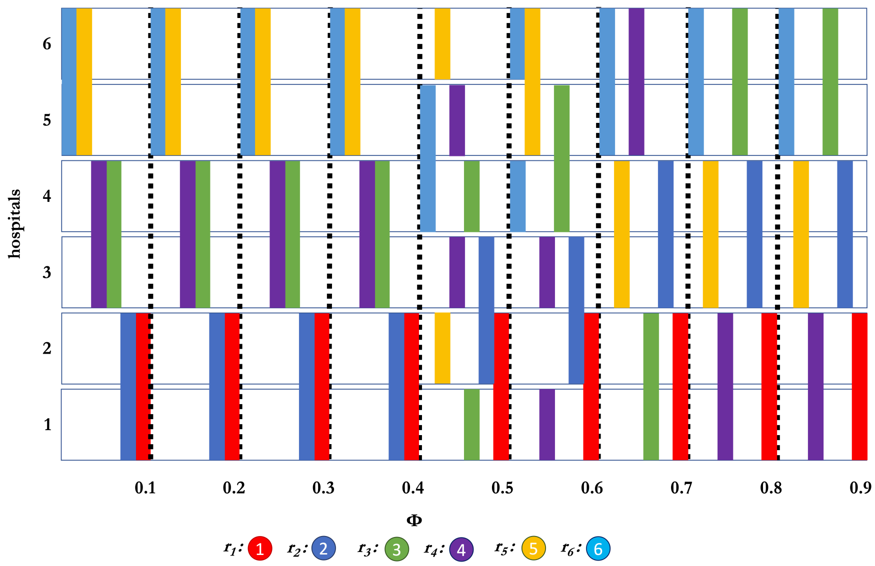

Figure 2: Interviewing sets of residents as a function of ϕ when using the Borda scoring function, with 6 participants, and interview set size of 2.

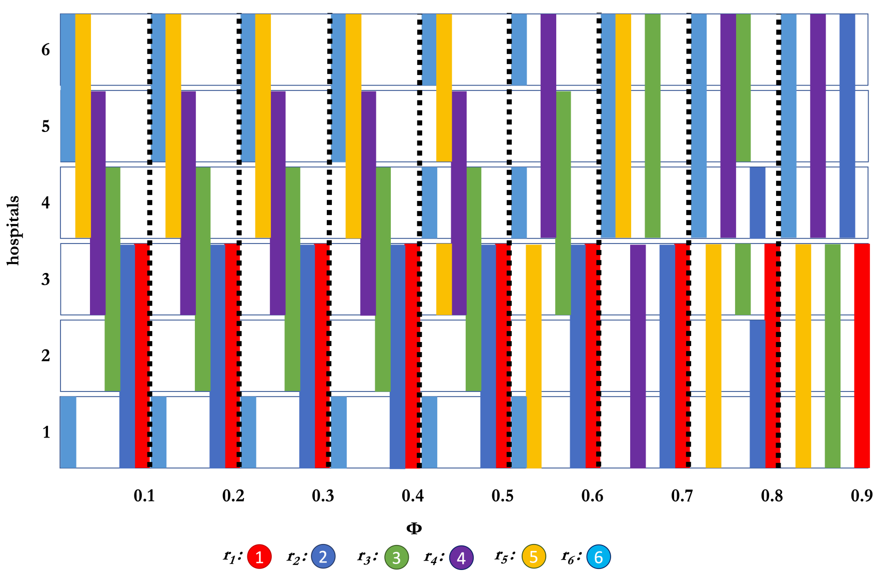

Figure 3: Interviewing sets of residents as a function of ϕ when using the Borda scoring function, with 6 participants, and interview set size of 3.

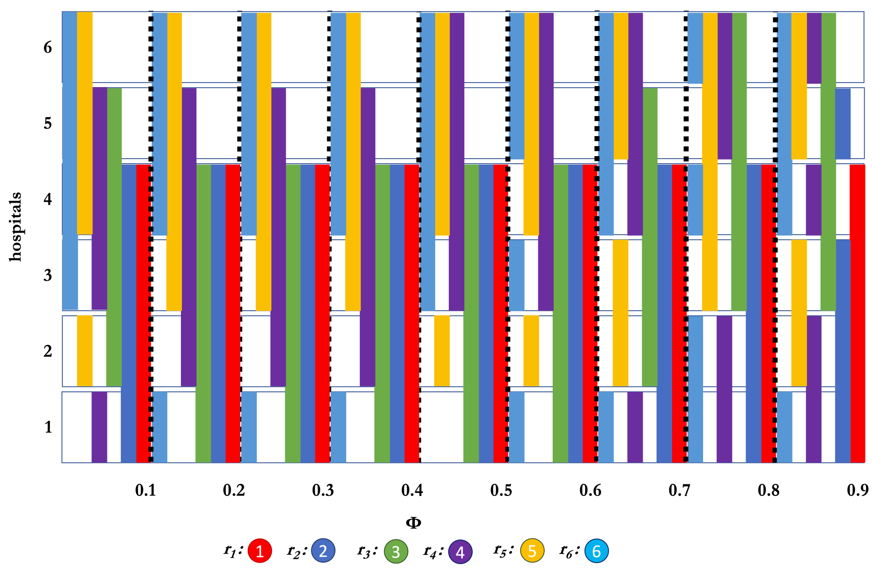

Figure 4: Interviewing sets of residents as a function of ϕ when using the Borda scoring function, with 6 participants, and interview set size of 4.

observe that the second resident departs from strict assortative interviewing in favor of a weak version of assortative interviewing and this in turn affects the other players, as, for example, the third resident applies what is basically a safety move (hospital 4, which is theirs if they want it) with a reach move (hospital 1, the top choice). Of some interest, for $\phi \geq 0.7$, all residents except r_1 use a weak assortative strategy, that is, they interview in sets of hospitals which are adjacent in rank, rather than splitting their interviews between radically different ranked hospitals.

Turning to $k = 3$ interviews per resident, we see as Theorem 5 claimed, that the third resident does not interview assortatively. While residents 3,4, and 5 are mostly weakly assortative (with the exception of the third resident and $\phi = 0.8$, where it tries a small reach choice), the sixth resident goes consistently for a reach and safety strategy, as it interviews in the top hospital as well. The resident's behaviour only changes when ϕ is large enough ($\phi > 0.6$), when the chance of the true ranking being different from the ground truth is much higher. Of interest, when $\phi = 0.9$, the second resident chooses hospitals 4,5,6 (even knowing that at least two of the hospitals in $\{1,2,3\}$ will be available. But when ϕ is sufficiently close to 1, the distribution is approaching the uniform distribution so that this resident might as well choose hospital 4,5,6 as they might very well be as desirable as 1,2,3 where the residents top choices might be taken.

Finally for $k = 4$, we see that resident 1 (as we know must happen) interviews assortatively for all settings of ϕ while other residents are much more willing to experiment. Not included in Figure 4 are further results, showing that even for very small ϕ ($\phi \leq 10^{-20}$), there are residents which are not even weakly assortative. We hypothesize that this "reach" and "safety" behaviour is present in markets with larger interviewing quotas.

8. Conclusions and Future Directions

We investigate equilibria for interviewing (for example, between residents and hospitals) with a limited quota when a master ranked list (say, of residents) is known. We provide a generic payoff (or utility) function that is indifferent to participants' interviewing quotas, 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 quotas 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. Furthermore, it seems natural that close to $\phi = 0$, assortative equilibrium is sensible, since at that value almost any agent is sure to take their exact ranking. Moreover, 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 Schwarz [31] showed the existence of a similarly structured equilibrium, and so it was somewhat surprising that the existence of this equilibria was so highly dependent on both the scoring-function structure and the distribution from which the underlying preferences were drawn. Beyond this, our simulations (in particular, Figures 3 and 4) seem to indicate even weaker forms of assortative interviewing (e.g., contiguous interviewing, in which interviews are a contiguous set; or a limited range of interviewing variance) do not seem to hold either.

There are numerous future research questions raised by our results, to which at least some of out technical results and techniques may also contribute. Most concretely, we hypothesize Theorem 9 could be replaced by extending Theorem 8 for all $k \geq 4$. Second, 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 relies 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 will obviously generalize the setting, and may invalidate our results. More specifically, 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 $\sum_{h \in S} c_r(h) \leq b_r$.

A 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 equilibrium. 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. One such possibility is for interviewing to be assortative for "safety" programs while allowing for one or a few "reach" programs. (See for example the strategy of resident 6 for small values of the Mallows' parameter in Figure 3.) 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 and Boutilier [5]) 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

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

[2] T. Andersson, Refugee matching as a market design application, in: J.-F. Laslier, H. Moulin, M. R. Sanver, W. S. Zwicker (Eds.), The Future of Economic Design, Springer, 2019, pp. 445–450.

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

[4] National Resident Matching Program, National Resident Matching Program, Results and Data: 2015 Main Residency Match®, Technical Report, National Resident Matching Program, Washington, DC, 2015.

[5] J. Drummond, C. Boutilier, Preference elicitation and interview minimization in stable matchings, in: Proceedings of the Twenty-Eighth AAAI

Conference on Artificial Intelligence (AAAI-14), AAAI Press, Québec City, Canada, 2014, pp. 645–653.

[6] B. Rastegari, A. Condon, N. Immorlica, K. Leyton-Brown, Two-sided matching with partial information, in: Proceedings of the Fourteenth ACM Conference on Electronic Commerce, ACM, Philadelphia, Pennsylvania, 2013, pp. 733–750.

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

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

[9] M. Gelain, M. S. Pini, F. Rossi, K. B. Venable, T. Walsh, Male optimal and unique stable marriages with partially ordered preferences, in: Proceedings of the International Conference on Collaborative Agents-Research and Development (CARE-09), Springer Berlin Heidelberg, Melbourne, Australia, 2009, pp. 44–55.

[10] M. S. Pini, F. Rossi, K. B. Venable, Stable matching problems with soft constraints, in: Proceedings of the Thirteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), International Foundation for Autonomous Agents and Multiagent Systems, Paris, France, 2014, pp. 1511–1512.

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

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

[13] M. Zhu, College admissions in China: A mechanism design perspective, China Economic Review 30 (2014) 618–631.

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

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

[16] M. Smart, Offer rates to different ethnic groups close to expected values, Analysis note 2015/05, UCAS Analysis and Research, 2015.

[17] J. Fortenbury, 3 types of schools you should consider when applying to colleges, USA Today, 2014.

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

[19] N. Immorlica, M. Mahdian, Marriage, honesty, and stability, in: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms, Vancouver, Canada, 2005, pp. 53–62.

[20] S. V. Kadam, Interviewing in matching markets, Working Paper, 2015.

[21] H. Beyhaghi, D. Sabán, É. Tardos, Effect of selfish choices in deferred acceptance with short lists, CoRR abs/1701.00849 (2017).

[22] H. Beyhaghi, É. Tardos, Randomness and fairness in two-sided matching with limited interviews, in: Proceedings of the 12th Innovations in Theoretical Computer Science (ITCS), volume 185, 2021, pp. 74:1–74:18.

[23] G. Haeringer, F. Klijn, Constrained school choice, Journal of Economic Theory 144 (2009) 1921–1947.

[24] C. Calsamiglia, G. Haeringer, F. Klijn, Constrained school choice: An experimental study, American Economic Review 100 (2010) 1860–1874.

[25] Y. He, T. Magnac, Application Costs and Congestion in Matching Markets, TSE Working Paper 17-870, Toulouse School of Economics, Tolouse, France, 2017.

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

[27] H. Chade, L. Smith, Simultaneous search, Econometrica 74 (2006) 1293–1307.

[28] R. I. Shorrer, Simultaneous Search: Beyond Independent Successes, Technical Report, Penn State University, 2019.

[29] P. Coles, A. Kushnir, M. Niederle, Preference Signaling in Matching Markets, Technical Report, National Bureau of Economic Research, 2010.

[30] A. Kushnir, Harmful signaling in matching markets, Games and Economic Behavior 80 (2013) 209–218.

[31] R. S. Lee, M. Schwarz, Interviewing in two-sided matching markets, The RAND Journal of Economics 48 (2017) 835–855.

[32] P. Coles, R. Shorrer, Optimal truncation in matching markets, Games and Economic Behavior 87 (2014) 591–615.

[33] C. L. Mallows, Non-null ranking models, Biometrika 44 (1957) 114–130.

[34] T. Lu, C. Boutilier, Learning Mallows models with pairwise preferences, in: Proceedings ICML 2011, Omnipress, Bellevue, Washington, 2011, pp. 145–152.

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

[36] G. Lebanon, Y. Mao, Non-parametric modeling of partially ranked data, Journal of Machine Learning Research 9 (2008) 2401–2429.

[37] J. Drummond, C. Boutilier, Elicitation and approximately stable matching with partial preferences, in: Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI-13), Beijing, AAAI Press, Beijing, China, 2013, pp. 97–105.

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

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

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

[41] C. Calsamiglia, F. M. Mora, A. Miralles, School Choice Design, Risk Aversion, and Cardinal Segregation, Discussion Paper Series 13464, Institute for the Study of Labor (IZA), Bonn, Germany, 2020.

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 $P(a_i \succ a_j | D^{\phi,\sigma})$.

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

However, because a_i, a_j are in ranking σ, the only difference between summing over the set of all rankings in $\{a_0, \ldots, a_p\}_{\succ}^{a_i \succ a_j}$ and $\{a_0, \ldots, a_{p-1}\}_{\succ}^{a_i \succ a_j}$ is that there for each permutation generated by $\{a_0, \ldots, a_{p-1}\}_{\succ}$, there are p permutations in $\{a_0, \ldots, a_p\}_{\succ}$, each one with a_p in a different place (and thus a different Kendall-τ distance). Fixing some $\eta \in \{a_0, \ldots, a_{p-1}\}_{\succ}$, 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:

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

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

Proof. (**Corollary 1**) 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 2**) 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 2**) 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 to the ranking (growing from $n-1$ to n), and must include that in the normalization factor Z. The normalization factor for $n-1$ alternatives is $(1+\phi)(1+\phi^2)\ldots(1+\ldots+\phi^{n-2})$. The normalization factor for n elements is identical, but multiplied by $1+\ldots+\phi^{n-1}$. □

Proof. (**Lemma 3**) 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, $P(\eta) < \frac{\phi^{j-i}}{Z}$. □