Natural Interviewing Equilibria for Stable Matching

Joanna Drummond, Allan Borodin, and Kate Larson

Abstract

Stable matching problems are ubiquitous, though much of the work on stable matching assumes that both sides of the market are able to fully specify their preference orderings. However, as the size of matching markets grow, this assumption becomes unrealistic, and so there has been interest in understanding how agents may use *interviews* to refine their preferences over subsets of alternatives. In this paper we study a market where one side (hospital residency programs) maintains a common preference master list, while the other side (residents) have idiosyncratic preferences which they can refine by conducting a limited number of interviews. The question we study is *How should residents choose their interview sets, given the choices of others?* We provide a payoff function for this imperfect information game, and find that this game always has a pure strategy equilibrium. Moreover, when residents are restricted to two interviews and their preferences are distributed according to a ϕ-Mallows model with low dispersion, there is a unique Bayesian equilibrium in which residents interview assortatively: each resident pair r_{2j}, r_{2j+1} interviews with hospitals h_{2j}, h_{2j+1}. We observe that with high dispersion, assortative interviewing is not an equilibrium.

1 Introduction

Real world matching problems are ubiquitous, and cover many domains. One of the most widely studied matching problems in economics is the canonical *stable matching problem (SMP)* [7]. Finding a stable matching is a goal of many real-world matching markets, including college admissions, school choice, reviewer-paper matching, various labor market matching problems [17], and, famously, the residency matching problem, where residents are matched to hospital programs via a centralized matching program (such as the National Residency Matching Program, NRMP, in the United States) [21].

This notion of stability, where no one in the market has both the incentive and ability to change their partner, has been empirically shown to be a very valuable property in real-world markets. Repeated improvement to stable matching mechanisms for the American medical market halted unraveling in that market and in other matching markets, centralized mechanisms that produced a stable match tended to halt unraveling, while unstable mechanisms tended to be abandoned [21]. Many of these markets implement the Deferred Acceptance (DA) mechanism, first introduced in Gale and Shapley's seminal paper [7].

However, in practice all of these mechanisms assume that participants provide their full preferences to the mechanism. Even if it is in participants' best interest to do so (as DA is strategy-proof for one side of the market), it is frequently infeasible for participants to list all alternatives they find acceptable. In the NRMP in 2015, there were 4,012 first year hospital programs that residents could apply to [19], however residents tend to apply to an average of only 11 programs, spending between $1,000 to $5,000 to do so [1].

This implies that, even if resident-proposing Deferred Acceptance (RP-DA) is the mechanism used, residents must be strategic about what hospital programs they choose to interview with, as they cannot be matched to a program they do not interview with. If there is too much competition for the hospitals they choose to interview with, residents risk not being matched at all. There is significant evidence of this happening, as an aftermarket

(SOAP) exists for the NRMP, having matched 1,666 programs to doctors in 2015, out of roughly 30,000 initial available positions [19].

While the residency matching market is one prominent example of this, these difficulties arise in other markets as well: university departments have limited budgets to interview faculty candidates, and candidates have limited time to interview; reviewers only have limited time to screen papers for their ability to review; high school students have limited funds to apply to colleges. However, little work has investigated (from a game-theoretic viewpoint) participants' strategic considerations in markets like these. Of the papers that have (e.g., [3, 2]), the markets studied have been *decentralized* matching markets. To the authors' knowledge, the only previous work that has investigated the strategic considerations participants in the market make when they know they will be matched via a centralized matching mechanism, such as DA, has assumed all participants in the market are ex-ante indifferent between all alternatives [11]. Our goal is to extend this work, and investigate when residents know their preferences are drawn according to some known (non-uniform) distribution, but, like Lee and Schwarz, do not know their idiosyncratic instantiation of that draw until after they choose their interviewing set.

To begin investigating this problem, we focus on using the residency matching problem as our motivating example, though we do make some simplifying assumptions when modeling to begin investigating the equilibria of this interviewing market. We assume that residents choose their interviewing sets (as they are the ones who have to pay to interview), that hospitals will find anyone who interviews with them acceptable, and that hospitals list their true preferences (i.e., there is no strategic behavior after the interviewing process is complete[1]). We also assume that residents only know their idiosyncratic preferences after having conducted their interviews. More importantly, we assume that there is a fixed ranking *a priori* over all residents that all hospitals share, and residents know their ranking (a "master list"). While a significant restriction, we note that this assumption mimics some real-world markets. For example, university entrance in Turkey is determined solely by a universal test score, and participants know their test score before applying to universities [8]. Chinese university matching markets likewise use a master list, and the market runs a centralized mechanism [22].

We first formalize a payoff function for any given resident in this game, and show that a pure strategy equilibrium for this game always exists (with very little restrictions on the distribution that residents' preferences are drawn from). We then investigate equilibria under one specific distribution, a ϕ-Mallows model. While we also restrict residents to two interviews to obtain initial results and gain insight, we conjecture the same conclusions hold for arbitrary, fixed number of interviews.[2] Under this model, there is some consistent, agreed upon reference ranking (e.g., the US News & World Report's Best College Rankings), and all residents' preferences are drawn according to a dispersion parameter ϕ. When dispersion is low (i.e., residents believe their preferences and others' preferences are very similar to the reference ranking), we show that it is an equilibrium for residents to interview assortatively in tiers: the best residents apply to the best hospitals, and the worst residents apply to the worst hospitals. This characterizes an equilibrium similar to that presented in Lee and Schwarz [11], however, our equilibrium arises naturally via the incentives in the market, and covers a very different range of ex-ante preference distributions than theirs does.

Furthermore, our results are consistent with the intuition that residents interview with the best hospitals that they think they have a good chance of being matched with. Interestingly, when preferences are not highly dispersed, we find no benefit for "reach" or "safety" choices, unlike previous work with decentralized matching mechanisms (e.g., [3, 2]). How-

---

[1] Residents have no incentive to be strategic after they choose their interviewing set, as the matching mechanism is RP-DA.

[2] Additionally, some real-world markets only have one interview (e.g., [22]).

ever, we also find evidence of a strong trade-off between participants favoring more choice over higher expected valuation of an alternative when residents' preferences are highly variable.

In Section 2, we further discuss previous literature, and define a ϕ-Mallows model distribution. In Section 3, we formally describe the game we investigate in this paper. Section 4 provides an equilibrium analysis, showing that under an arbitrary number of interviews and with an arbitrary preference distribution for residents, a pure equilibrium always exists. We then explicitly characterize that equilibrium for two interviews and residents' preferences drawn from a Mallows model with low dispersion. Our analysis suggests many extensions which we discuss in Section 5. We finish with a brief conclusion in Section 6.

2 Background

In this paper we investigate the stable matching problem, as standardly defined [7]. A *stable matching* is a matching that is individually rational (no one would rather go unmatched than be matched to the alternative they are matched to) and does not contain any *blocking pairs*. A blocking pair is a resident/hospital pair that both prefer to be matched with each other than their assigned match.

The Deferred Acceptance (DA) algorithm [7] is one of the most famous algorithms for solving stable matching problems. In this algorithm, one side of the market "proposes" to the other, which chooses to be tentatively matched to the best alternative out of all proposals received thus far. Resident-proposing DA (RP-DA) runs in polynomial time, and has several nice properties: the resulting matching is guaranteed to be stable and resident-optimal (and hence resident strategy-proof) [21].

While there has been great interest in finding stable matchings for various markets, little work has investigated how residents choose *who* they interview with, particularly in a game theoretic setting. Some work has investigated interviewing and partial preferences in stable matching problems. Rastegari et al. investigate minimal interviewing policies [20]. They show that, in general, finding such a policy is NP-hard, but under certain preference restrictions, an algorithm that finds a minimal interviewing policy exists. Note that their result relies on a few assumptions: first, they assume that participants give all known preference information (i.e., any information that is known without having to perform interviews). Importantly, they also require that the resulting matching is stable and resident-optimal under *all* possible completions of the currently known partial preferences. This guarantees that there are no potential blocking pairs after all interviews have been performed, no matter what agents' underlying preferences are. Thus, the number of interviews required is highly dependent on the amount of initial information agents provide; unless agents can provide a large amount of information initially to the algorithm, they will be forced to perform a large number of interviews. Thus, while this algorithm provides strong guarantees, participants in the market must still provide a great deal of information. Furthermore, Rastegari et al. do not investigate any incentive compatibility issues regarding their algorithm. Using heuristics, Drummond and Boutilier [6] investigate a similar problem and likewise do not investigate incentive compatibility.

Chade et al [2] investigate search frictions in the college admissions problem. This problem investigates a Bayesian approach to students deciding where to interview. Under this model, students know which college is more desirable, and colleges have noisy information about which students are most desirable. They then investigate equilibria under this model, finding a tractable separable solution. This work importantly differs from our model in that the admissions process is decentralized, and thus any student's chance of admission is modeled *independent* of any other students' behavior. As our matching mechanism is RP-DA, we cannot make this independence assumption.

Chade and Smith [3] investigate a similar problem, with similar motivations as those discussed here. They assume that all residents agree perfectly on the hospital ranking, whereas hospitals' preferences are dependent on noisy signals of students' caliber. However, they investigate a decentralized market, instead of one running DA.

Coles et al. [4] discuss signaling in matching markets. Agents' preferences are distributed according to some (restricted) a priori known distributions, and each agent knows their own preferences. Firms can send at most one job offer, and workers can send one *signal* to a firm indicating their interest. Under this setting, firms can do better than simply offering their top candidate a job. In some ways, this achieves more of the overlap desired by the Lee and Schwarz perfect overlap equilibrium [11] without having any firm coordination, but requires extra machinery not explicitly present in matching markets like the NRMP. Furthermore, Kushnir [9] provides an example where signaling is harmful.

The work most closely related to the problem posed in this paper investigates an interviewing game, where firms and workers interview with each other in order to be matched [11]. This is a two-stage game, where firms choose to interview with workers for some fixed cost. These interviews reveal both firms' and workers' preferences. Then, all participants in the market submit the results of their interviews to the matching mechanism, firm-proposing deferred acceptance. Their results show that unless firms coordinate when picking their interviewing set, picking k workers to interview at random is the firm's best strategy. If firms do coordinate, it is best for them to interview with perfect overlap (i.e., the interviews form n/k complete bipartite disconnected components).

Lee and Schwarz make the following assumptions: firms bear all cost of interviewing workers; firms and workers must interview with each other to be matched; firms may find some workers they interview unemployable; workers prefer all firms over being unemployed; and, most importantly, all firms and workers are ex-ante homogeneous, with agents' revealed preferences idiosyncratic and independent [11].

This last assumption is an incredibly strong one; for their 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. (That is, everyone knows and agrees on who the most desirable firms and workers are, but preferences between those top workers are ex-ante homogeneous; note that this is equivalent to all agents having block-correlated preferences, and those blocks are known ex-ante). In this paper, we thus focus on investigating a very different set of assumptions from Lee and Schwarz, but find a similar, naturally arising equilibrium.

2.1 Probabilistic Preference Models

While the payoff function as formulated in Section 3.2 is indifferent with respect to the probability distribution, it is currently formulated as a distribution over rankings. Two commonly studied distributions over rankings are the Plackett-Luce model [18, 13] and the ϕ-Mallows model [14, 15].

When characterizing equilibria for a specific distribution, we focus on the ϕ-Mallows model, as there is an intuitive relationship between the parameters of the model, and how "similar" market participants' preferences are. The Mallows model is characterized by a reference ranking σ, and a dispersion parameter $\phi \in (0, 1]$, [3] which we denote as $\mathscr{D}^{\phi,\sigma}$. Let A denote the set of alternatives that we are ranking, and let $P(A)$ denote the set of all

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

permutations of A. The probability of any given ranking r is:

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

Here d is Kendall's τ distance metric, and Z is a normalizing factor; $Z = \sum_{r' \in P(A)} \phi^{d(r,\sigma)} = (1)(1+\phi)(1+\phi+\phi^2)...(1+...+\phi^{|A|-1})$, as shown in [12].

As $\phi \to 0$, the distribution approaches drawing the reference ranking σ with probability 1; when $\phi = 1$, this is equivalent to drawing from the uniform distribution. The Mallows model (and mixtures of Mallows) have plausible psychometric motivations and are commonly used in machine learning [16, 10, 12]. Mallows models have also been used in previous investigations of preference elicitation schemes for stable matching problems (e.g., [5, 6]).

3 Model

There are n residents and n hospital programs. The set of residents is denoted by $R = \{r_0, r_1, ..., r_{n-1}\}$; the set of hospital programs is denoted by $H = \{h_0, h_1, ..., h_{n-1}\}$. We are interested in *one-to-one matchings* which means that residents can only do their residency at a single hospital, and that hospitals can accept at most one resident. A *matching* is a function $\mu : R \cup H \to R \cup H$, such that $\forall r \in R$, $\mu(r) \in H \cup \{r\}$, and $\forall h \in H$, $\mu(h) \in R \cup \{h\}$. If $\mu(r) = r$ or $\mu(h) = h$ then we say that r or h is unmatched. A matching μ is *stable* if there does not exist some $(r, h) \in R \times H$, such that $h \succ_r \mu(r)$ and $r \succ_h \mu(h)$.

Both hospitals and residents have (strict) preferences over each other, and we let $P(H)$ and $P(R)$ denote the sets of all possible preference rankings over H and R respectively. We assume that hospitals have identical preferences over all residents, which we call the *master list*, $\succ_H$. Without loss of generality, let $\succ_H = r_0 \succ r_1 \succ \ldots \succ r_{n-1}$ where $r_i \succ_H r_j$ means that r_i is preferred to r_j, according to $\succ_H$. We further assume that the master list is common knowledge to all members of H and R. That is, all hospitals agree on the preference ranking over residents and each resident knows where they, and all others, rank in the list. While each resident, r, has idiosyncratic preferences over the hospitals, we assume that these are drawn *i.i.d.* from some common distribution $\mathscr{D}$, and that this is common knowledge. If resident r draws preference ranking η from $\mathscr{D}$, then $h_i \succ_\eta h_j$ means that h_i is preferred to h_j by r under η. We assume there is some common valuation function $v : H \times P(H) \mapsto \mathbb{R}$, applied to rankings η drawn from $\mathscr{D}$ such that, given any $\eta \in P(H)$ with $h_i \succ_\eta h_j$, $v(h_i, \eta) > v(h_j, \eta)$.

Critical to our model is the assumption that residents do not initially know their true preferences, but can refine their knowledge by conducting *interviews*. We let $I(r_j) \subset H$ denote the interview set of resident r_j, and assume that $|I(r_j)| \leq k$ for some fixed $k \leq n$. Once r_j has finished interviewing, r_j knows their preference ranking over $I(r_j)$. It then submits this information to the matching algorithm, resident-preferred deferred acceptance (RP-DA). The matching proceeds in rounds, where in each round unmatched residents propose to their next favorite hospital from their interview set to whom they have not yet proposed. Each hospital chooses its favorite resident from amongst the set of residents who have just proposed and its current match, and the hospital and its choice are then tentatively matched. This process continues until everyone is matched. The resulting matching, μ, is guaranteed to be stable, resident-optimal, and hospital-pessimal [7].

3.1 Description of the Game

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

1. Each resident $r \in R$ simultaneously selects an interviewing set $I(r) \subset H$, based on their knowledge of $\mathscr{D}$ and the hospitals' master list $\succ_H$, where $|I(r)| \leq k$.
2. Each resident, r, interviews with hospitals in $I(r)$ and discovers their preferences over members of $I(r)$.
3. Each resident reports their learned preferences over $I(r)$ and reports all other hospitals as unacceptable, and each hospital reports the master list to a centralized clearing-house, which runs resident proposing deferred acceptance (RP-DA), resulting in the matching μ.

3.2 Payoff function for Interviewing with a Limited Budget

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

$$u_{r_j}(S) = u_{r_j}(S|\mathscr{D}, I(r_0), ..., I(r_{j-1})).$$

Given fixed interviewing sets $I(r_0), I(r_1), ..., I(r_{j-1})$, and some partial match $m = \mu_{|r_0,r_1,...,r_{j-1}}$, we must compute the probability that m happened via RP-DA. Let $m(r_i)$ denote who resident r_i is matched to under m. For any r_i, there is a set of rankings consistent with r_i being matched with $m(r_i)$ under RP-DA (and hospitals' master list $\succ_H$). Denote this set as $T(r_i, m)$. Formally, $T(r_i, m) \subseteq P(H)$ is defined as:

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

Given the interviewing sets of residents $r_0, \ldots, r_{j-1}$, the probability of partial match m is

$$Pr(m|I(r_0), ..., I(r_{j-1})) = \prod_{r_i \in \{r_0,...,r_{j-1}\}} \sum_{\xi \in T(r_i,m)} Pr(\xi|\mathscr{D}). \tag{1}$$

where $Pr(\eta|\mathscr{D})$ is the probability that some resident drew ranking $\eta \in P(H)$ from $\mathscr{D}$.

Using Eq. 1, we can now determine the probability that some hospital h is matched to r_j using RP-DA, when r_j has interviewed with set S, and has preference list η. We simply sum over all possible matches in which this could happen. Because RP-DA is resident optimal, and all hospitals have a master list, any hospital that r_j both interviews with and prefers to h must already be matched. We formally define the set of such matchings, $M^*(S, \eta, I(r_0), ..., I(r_{j-1}))$:

$$\begin{aligned} &M^*(S, \eta, I(r_0), ..., I(r_{j-1}), h) = \\ &\quad \{m \in M | m(r_j) = h; \forall r_i \in \{r_0, ..., r_{j-1}\} m(r_i) \in I(r_i); \\ &\quad \text{ and } \forall x \in S, \text{ if } x \succ_\eta h, \exists r_i \in \{r_0, ..., r_{j-1}\} \text{ s.t. } x \in I(r_i) \text{ and } m(r_i) = x\} \end{aligned}$$

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

$$Pr(\mu(h) = r_j|\eta, S, I(r_0), ..., I(r_{j-1})) = \sum_{m \in M^*(S,\eta,I(r_0),...,I(r_{j-1}),h)} Pr(m|I(r_0), ..., I(r_{j-1})). \quad (2)$$

Finally, we have all of the building blocks to formally define the payoff function. Recall that $v(h, \eta)$ is the imposed utility function, which is deterministic given η. Then, our payoff function is:

$$u_{r_j}(S) = \sum_{h \in S} \sum_{\eta \in P(H)} v(h, \eta) Pr(\eta|\mathscr{D}) Pr(\mu(h) = r_j|\eta, S, I(r_0), ..., I(r_{j-1})) \quad (3)$$

Intuitively, what the payoff function in Eq. 3 does is weight the value for some given alternative by how likely r_j is to be matched to that item, given the interview sets of the "more desirable" residents, $r_0, \ldots, r_{j-1}$.

As an illustrative example, imagine there are two residents, r_0 and r_1, each of whom have interviewed with hospitals h_0 and h_1. Resident r_0 will be matched with whom ever she most prefers, while r_1 will be assigned the other. The probability that r_1 will be assigned h_0 is simply the probability that r_0 drew ranking $h_1 \succ h_0$, while the probability that r_1 is matched to h_1 is the probability that r_0 drew ranking $h_0 \succ h_1$.

4 Equilibrium Analysis

We provide an equilibria analysis for the game as presented in Section 3. We first show that a pure equilibrium for this game always exists, but may be difficult to calculate. We then show that under some additional distributional assumptions, for two interviews, residents' best response is easy: any two residents r_{2j}, r_{2j+1} will both interview with hospitals h_{2j}, h_{2j+1}, forming the perfect overlap interviewing structure found in the Lee and Schwarz paper [11], but under very different modelling assumptions.

Theorem 1 *A pure strategy always exists for the Interviewing with a Limited Budget game.*

Proof: We wish to show that if every resident chooses their expected utility maximizing interviewing set, this forms a pure strategy. Given any resident r_j who is jth in the hospitals' rank ordered list, r_j's expected payoff function only depends on residents $r_0, ..., r_{j-1}$. As r_j knows that each other resident r_i is drawing from distribution $\mathscr{D}$ iid, she can calculate $r_0, ..., r_{j-1}$'s expected utility maximizing interview set. Her payoff function depends only on $\mathscr{D}$ and $I(r_0), ..., I(r_{j-1})$, both of which she now has. She then calculates the expected payoff for each $\binom{n}{k}$ potential interviewing sets, and interviews with the one that maximizes her expected utility. □

Our equilibrium analysis for two interviews when all residents draw preferences from a Mallows model with low dispersion requires some additional results regarding Mallows models, as shown in Lemma 2, Corollary 3, and Corollary 4. To the best of our knowledge, the following results regarding Mallows models have not been stated previously, and may be of more general interest. The proofs are provided in Appendix A.

Lemma 2 *Given some Mallows model $\mathscr{D}^{\phi,\sigma}$ with fixed dispersion parameter ϕ and reference ranking $\sigma = a_i \succ a_j$, then the probability that a ranking η is drawn from $\mathscr{D}^{\phi,\sigma}$ such that $a_i \succ_\eta a_j$ is equal to drawing from some distribution $\mathscr{D}^{\phi,\sigma'}$ where σ is a prefix of σ'. By symmetry, this proof also holds when σ is a suffix of σ'.*

Corollary 3 *Given any reference ranking* σ *and two alternatives* a_i, a_j *such that* $rank(a_j, \sigma) = rank(a_i, \sigma) + 1$*, then* $Pr(a_i \succ a_j | D^{\phi,\sigma}) = \frac{1}{1+\phi}$.

Corollary 4 *Given any reference ranking* σ *and three alternatives* a_w, a_x, a_y *such that* $rank(a_x, \sigma) = rank(a_w, \sigma) + 1$ *and* $rank(a_y, \sigma) = rank(a_x, \sigma) + 1$ *and some* $\eta \in P(\{a_w, a_x, a_y\})$*, then the probability that we draw some ranking* β *consistent with* η *is:* $Pr(\beta | \mathscr{D}^{\phi,\sigma}) = \frac{\phi^{d(\eta, a_w \succ a_x \succ a_y)}}{(1+\phi)(1+\phi+\phi^2)}$.

We now begin the proofs for the main equilibrium characterization. We show that, for two interviews ($k = 2$) with a sufficiently small dispersion parameter, there is a naturally arising equilibrium for all residents to interview assortatively in tiers. As discussed in Section 5, we note that the following result does not hold for all values of ϕ.

Theorem 5 *Given residents' valuation function* $v(h, \eta) = n - rank(h|\eta)$ *(i.e., Borda score) for any ranking* η *and market size* n*, for a Mallows model with reference ranking* $\sigma = h_0, h_1, ..., h_{n-1}$ *with dispersion parameter* ϕ *such that* $0 < \phi \leq 0.265074$*, resident* r_1 *maximizes her expected payoff by interviewing with* $\{h_0, h_1\}$.

Proof Sketch: We provide a proof sketch here, but leave full details to the appendix.

As resident r_0 greedily chooses to interview with $\{h_0, h_1\}$, we note that resident r_1 must calculate a trade-off between a higher expected value for hospitals in a potential interviewing set, and competition for those hospitals. We prove that r_1 maximizes her expected payoff in this interval by bounding the difference in expected payoff between interviewing sets: $u_{r_1}(h_0, h_1 | D^{\phi^*,\sigma}) - u_{r_1}(h_i, h_j | D^{\phi^*,\sigma}) \geq 0$, for all i, j.

Note that r_1 has $\binom{n}{k}$ interviewing sets to choose from, but many are dominated by some other set with equal availability probability, but higher expected value. $\{h_0, h_1\}, \{h_0, h_2\}, \{h_1, h_2\}, \{h_2, h_3\}$ are undominated. Intuitively, the difference between these sets is a trade-off between a higher expected valuation of the hospitals, versus more competition with resident r_0. We compute lower bounds comparing the difference in expected utility between interviewing with $\{h_0, h_1\}$ and each of the potential interview sets independently. We provide a proof sketch for $\{h_1, h_2\}$, but omit all other details to the appendix (as the arguments are analogous).

We prove that choosing $\{h_0, h_1\}$ is better than choosing $\{h_1, h_2\}$, 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_0, h_1, h_2. We then pair these summed permutations in such a manner that makes it easy to find a lower bound for $u_{r_1}(\{h_0, h_1\}) - u_{r_1}(\{h_1, h_2\})$. 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_0, h_1, h_2 as follows:
Case 1: all rankings η consistent with $h_1 \succ h_0 \succ h_2$ or η' consistent with $h_1 \succ h_2 \succ h_0$;
Case 2: all rankings η consistent with $h_0 \succ h_1 \succ h_2$ or η' consistent with $h_2 \succ h_1 \succ h_0$;
Case 3: all rankings η consistent with $h_0 \succ h_2 \succ h_1$ or η' consistent with $h_2 \succ h_0 \succ h_1$.

Note that as we have enumerated all possible permutations of h_0, h_1, h_2, these three cases generate every ranking in $P(H)$. Furthermore, for any one of the three cases, we can iterate over only all possible rankings η that are consistent with the first member of the pair, and generate the ranking η' consistent with the second member of the pair by simply swapping two alternatives in the rank. Moreover, given some η, the number of discordant pairs in η' is simply the number in η, plus the number of additional discordant pairs between h_0, h_1, h_2 caused by swapping the two alternatives.

For clarity, let $u_{r_1}(\{h_0, h_1\}) - u_{r_1}(\{h_1, h_2\}) = U_1 + U_2 + U_3$, where U_1, U_2, U_3 correspond to our three cases. We also introduce the notation $Pr_{\mu(r_i)}(h)$ to denote the probability that r_i is matched to hospital h under matching μ. That is, $Pr_{\mu(r_i)}(h) = Pr(\mu(r_i) = h)$.

Case 1. Because we have fixed $h_1 \succ h_0 \succ h_2$ or $h_1 \succ h_2 \succ h_0$, we know exactly what r_1's match will be, given r_0's match. As we know r_0's interviewing set ($\{h_0, h_1\}$), and the distribution r_0's preferences are drawn from iid, we know the likelihood that either h_0 or h_1 is available; by Lemma 3, $Pr(\mu(r_0) = h_0) = \frac{1}{1+\phi}$. Using this information, the payoff function, and the definition of η, η', we find a lower bound for U_1:

$$U_1 \geq Pr_{\mu(r_0)}(h_1)(1)(1-\phi)Pr(h_1 \succ h_0 \succ h_2) \tag{4}$$
$$= \left(\frac{\phi}{1+\phi}\right)\left(\frac{\phi}{(1+\phi)(1+\phi+\phi^2)}\right)(1-\phi) \tag{5}$$

Case 2. We fix $h_0 \succ h_1 \succ h_2$ or $h_2 \succ h_1 \succ h_0$. This case is analogous to Case 1:

$$U_2 \geq Pr(h_0 \succ h_1 \succ h_2)\frac{2}{1+\phi}(\phi - \phi^3 - \phi^4) \tag{6}$$

Case 3. We fix $h_0 \succ h_2 \succ h_1$ or $h_2 \succ h_0 \succ h_1$. Again, we look at pairs of rankings η, η', where η is consistent with $h_0 \succ h_2 \succ h_1$, and η' is identical to η, except $\text{rank}(h_0, \eta) = \text{rank}(h_2, \eta')$, and $\text{rank}(h_2, \eta) = \text{rank}(h_0, \eta')$.

Then, as before, we sum over all possible rankings consistent with $h_0 \succ h_2 \succ h_1$, but we break this into two subcases, so that $U_3 = U_{3a} + U_{3b}$:

$$U_{3a} = \sum_{\eta \in P(H)^{h_0 \succ h_2 \succ h_1}} Pr_{\mu(r_0)}(h_0)[(v(h_1,\eta) - v(h_2,\eta))Pr(\eta|\mathscr{D}^{\phi,\sigma}) + (v(h_1,\eta') - v(h_2,\eta'))Pr(\eta'|\mathscr{D}^{\phi,\sigma})]$$
$$U_{3b} = \sum_{\eta \in P(H)^{h_0 \succ h_2 \succ h_1}} Pr_{\mu(r_0)}(h_1)[(v(h_0,\eta) - v(h_2,\eta))Pr(\eta|\mathscr{D}^{\phi,\sigma}) + (v(h_0,\eta') - v(h_2,\eta'))Pr(\eta'|\mathscr{D}^{\phi,\sigma})]$$

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

$$U_{3b} \geq \frac{\phi}{\phi+1}(1-\phi)Pr(h_0 \succ h_2 \succ h_1) \tag{7}$$

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 that the following bound holds:

$$U_{3a} \geq Pr_{\mu(r_0)}(h_0)\frac{-\phi}{(1+\phi)(1+\phi+\phi^2)}\left(\frac{\phi}{(1-\phi)^4} + \frac{1}{3(1-\phi)^3} + \frac{2}{3}\right)(1+\phi) \tag{8}$$

We have considered all cases, and can now combine them together. We add the bounds for U_1 (Eq. 5), U_2 (Eq. 6), U_{3a} (Eq. 8), and U_{3b} (Eq. 7). We simplify and substitute using Corollaries 3 and 4, giving us:

$$\begin{aligned}
u_{r_1}(\{h_0,h_1\}) - u_{r_1}(\{h_1,h_2\}) \geq\ & \frac{\phi^2}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}(1-\phi) \\
& + \frac{2}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}(\phi - \phi^3 - \phi^4) \\
& - \frac{\phi}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}\left(\frac{\phi}{(1-\phi)^4} + \frac{1}{3(1-\phi)^3} + \frac{2}{3}\right)(1+\phi) \\
& + \frac{\phi^2}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}(1-\phi)
\end{aligned} \tag{9}$$

Thus, Eq. 9 gives us a lower bound for the difference in expected utility between $\{h_0, h_1\}$ and $\{h_1, h_2\}$ for resident r_1, for all n. Using numerical methods to approximate the roots of Eq. 9, we get that there is a root at 0, and a root at $\phi \approx 0.265074$.

In the appendix, we provide bounds such that $u_{r_1}(\{h_1, h_2\}) - u_{r_1}(\{h_2, h_3\}) \geq 0$ if $0 < \phi < 0.3550107$, and that $u_{r_1}(\{h_0, h_1\}) - u_{r_1}(\{h_0, h_2\}) \geq 0$ if $0 < \phi < 0.413633$ (the proofs are analogous to the one presented here). Thus, for the interval $0 < \phi \leq 0.265074$, we have shown that r_1's best move in this interval is to interview with $\{h_0, h_1\}$. $\square$

Theorem 6 *Given residents' valuation function $v(h, \eta) = n - rank(h|\eta)$ for any ranking η and market size n, for a Mallows model with dispersion parameter ϕ such that $0 < \phi < 0.1707951$, if all residents r_{2f}, r_{2f+1} have interviewed with hospitals h_{2f}, h_{2f+1} for $f < j$, then residents r_{2j}, r_{2j+1} will interview with hospitals $\{h_{2j}, h_{2j+1}\}$.*

Proof Sketch: We again omit a full proof to Appendix D. We first note that for any hospital h_a such that $h_a \succ_\sigma h_{2j}$, interviewing with any other hospital dominates interviewing with h_a, because the probability r_{2j} or r_{2j+1} will be matched with h_a is 0, as h_a is already matched to a more desirable doctor. Likewise, interviewing with any alternative h_b such that $h_{2j+3} \succ_\sigma h_b$ is dominated by interviewing with h_{2j+3}. Resident r_{2j} does best by greedily choosing the top two hospitals left, h_{2j} and h_{2j+1}. Resident r_{2j+1} must again investigate the following undominated interviewing sets: $\{h_{2j}, h_{2j+1}\}, \{h_{2j+1}, h_{2j+2}\}, \{h_{2j+2}, h_{2j+3}\}, \{h_{2j}, h_{2j+2}\}$. We provide a sketch of the comparison between $\{h_{2j}, h_{2j+1}\}$ and $\{h_{2j+1}, h_{2j+2}\}$, leaving the remainder to Appendix E as the argument is similar. For clarity, let $h_{2j} = a_0$; $h_{2j+1} = a_1$; $h_{2j+2} = a_2$. We adapt the proof used in Theorem 5.

As in the proof for Theorem 5, we look at three cases, pairing all possible permutations of a_0, a_1, a_2 in the analogous manner. Cases 1 and 2 are completely analogous to the proof in Theorem 5; the argument in these cases only requires that $\text{rank}(h_0), \sigma) = \text{rank}(h_1, \sigma) - 1 = \text{rank}(h_2), \sigma) - 2$. As this holds for a_0, a_1, a_2, the argument holds.

The argument for Case 3, however, does require that h_0, h_1, h_2 are the first three elements in σ for Case 3a. As the payoff function for 3a is still monotonically decreasing, we provide a similar proof, but give a weaker lower bound (by a multiplicative factor of 2) to ensure we count at least the required discordant pairs appropriately.

Using these bounds, we again are able to derive a function entirely in terms of ϕ, finding a zero at roughly 0.1707951. Thus r_{2j+1} chooses to interview with $\{h_{2j}, h_{2j+1}\}$ over $\{h_{2j+1}, h_{2j+2}\}$ whenever $0 \leq \phi \leq 0.1707961$. As in Theorem 5, the interviewing set that imposes the tightest bound on ϕ is $\{h_{2j}, h_{2j+1}\}$. Therefore, h_{2j+1} chooses to interview with $\{h_{2j}, h_{2j+1}\}$ whenever $0 \leq \phi \leq 0.1707961$, as required $\square$

Corollary 7 *Given residents' valuation function $v(h, \eta) = n - rank(h|\eta)$, and a set of n Mallows models, each with reference ranking $\sigma = h_0 \succ h_1 \succ ... \succ h_{n-1}$, but each with different dispersion parameter ϕ_j such that $0 < \phi_j < 0.1707951$, every pair of residents r_{2j}, r_{2j+1} interviewing with $\{h_{2j}, h_{2j+1}\}$ forms an equilibrium.*

Proof: This is a direct result of combining Theorems 5 and 6. First, resident r_0 must greedily pick interviewing with the two hospitals with the best expected valuation, h_0 and h_1. As there is no competition for r_0, the payoff function becomes:

$$u_{r_0}(S) = \sum_{\eta \in P(H)} Pr(\eta|\mathscr{D}) \max_{a \in S} v(a, \eta)$$

Then, by Theorem 5, resident r_1 interviews with hospitals h_0, h_1 at ϕ_{r_1}. By Theorem 6, r_2, r_3 also maximize their expected utility by interviewing with h_2, h_3. This process iteratively continues, and by Theorem 6, every pair r_{2j}, r_{2j+1} maximizes their expected utility by interviewing with h_{2j}, h_{2j+1} $\square$.

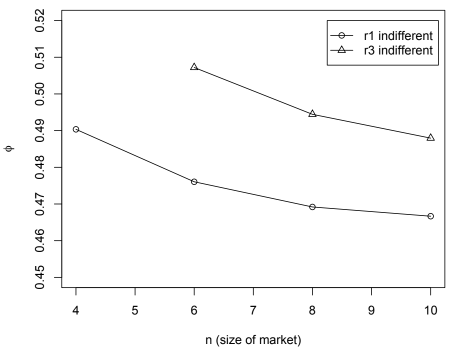

Figure 1: Size of the market versus resident r_{2j+1}'s values of ϕ^* such that for $0 < \phi < \phi^*$, that resident prefers interviewing with $\{h_{2j}, h_{2j+1}\}$ to $\{h_{2j+1}, h_{2j+2}\}$.

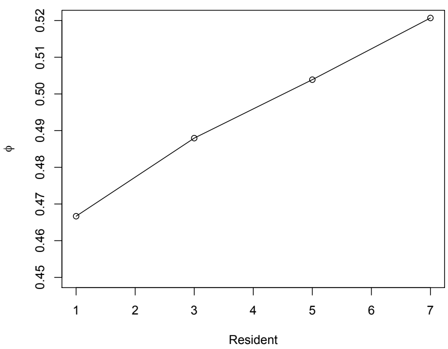

Figure 2: Resident versus their values of ϕ^* such that for $0 < \phi < \phi^*$, that resident prefers interviewing with $\{h_{2j}, h_{2j+1}\}$ to $\{h_{2j+1}, h_{2j+2}\}$, for $n = 10$.

5 Observations and Conjectures

While Section 4 provides theoretical guarantees for equilibria under specific distributions and valuation functions, we hypothesize that this natural equilibrium—residents r_{2j}, r_{2j+1} interview with h_{2j}, h_{2j+1} for all j— is present in a much larger range of distributions and valuation functions.

We note that Figures 1 and 2 provide evidence that the bounds from Theorems 5 and 6 are loose. The markets in Figures 1 and 2 are small ($n \leq 10$), so we can exactly calculate the payoff function. We find that the ϕ such that r_1 is indifferent between interviewing with $\{h_0, h_1\}$ and $\{h_1, h_2\}$ appears to converge quickly to approximately 0.46, as shown in Figure 1. We also find that, contrary to the bound provided in Theorem 6, the value of ϕ such that r_{2j+1} is indifferent between $\{h_{2j}, h_{2j+1}\}$, $\{h_{2j+1}, h_{2j+2}\}$ actually *increases* as j increases. This is shown in Figure 2. This leads us to the following conjecture:

Conjecture 8 *Given residents' valuation function $v(h, \eta) = n - rank(h|\eta)$, and a Mallows model with some reference ranking $\sigma = h_0 \succ h_1 \succ ... \succ h_{n-1}$ and dispersion parameter ϕ^* such that $0 < \phi^* \leq 0.46$, every pair of residents r_{2j}, r_{2j+1} interviewing with h_{2j}, h_{2j+1} forms an equilibrium.*

We find evidence that the perfect interviewing overlap equilibrium noted in the Lee and Schwarz paper extends to models that are close to uniform. We provide evidence of this when $n = 4$, as shown in Figure 3. The expected payoff of the three best interviewing sets are shown as ϕ increases from almost identical preferences to fully uniform preferences. Here, we explicitly see the trade-off between more choice (interviewing with h_2, h_3 for distributions close to uniform) and expected value. Interestingly, when $\phi \in [0.5, 0.6]$, r_1's best option is to split the difference, and interview with one hospital he is guaranteed to get (h_2) and one hospital that will be available with sufficiently high probability, but has a higher expected value (h_1). This choice from r_1 also causes some of the "reach" behavior we see in college admissions markets; r_2's best response now is to interview with h_0, h_3 (a "reach" choice, and a "safe" bet). We hypothesize that there may be many interesting results for preferences located in this parameter range.

We note that the desired n/k complete bipartite interviewing subgraphs equilibrium as described by Lee and Schwarz appears to hold for two large regions of the distribution space: when ϕ is sufficiently close to 0, and when ϕ is sufficiently close to 1. When ϕ is close to

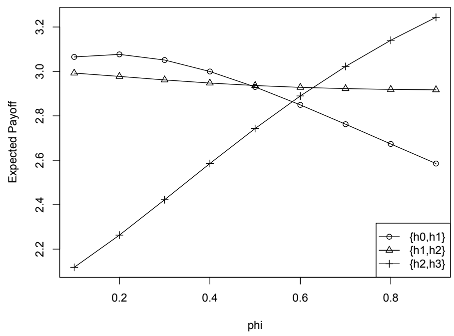

Figure 3: r_1's expected payoff for interviewing with various interviewing sets, as ϕ goes from 0 to 1, $n = 4$.

1, residents choose the set that has the least competition; when ϕ is close to 0, they choose the set that has the best hospitals. We thus conjecture:

Conjecture 9 *Given residents' valuation function $v(h, \eta) = n - rank(h|\eta)$, and residents' preferences distributed according to some ϕ-Mallows model with reference ranking σ, for ϕ sufficiently close to 0 or ϕ sufficiently close to 1, the interview graph forms n/k complete bipartite components.*

We also hypothesize that Theorem 5, Theorem 6, and Corollary 7 can be generalized to an arbitrary number of interviews. While $k = 2$ greatly simplifies the payoff function calculation, the intuition stays the same as the number of interviews increases: when residents are fairly certain their preferences are all very similar, residents choose to apply to the best set of hospitals they have non-zero probability of getting matched to.

Conjecture 10 *Given residents' valuation function $v(h, \eta) = n - rank(h|\eta)$ for any ranking η and market size n, for a fixed interviewing budget of k interviews, there exists some $\phi^{(k)}$ (which may be dependent on k) such that for all ϕ with $0 < \phi \leq \phi^{(k)}$, every block of residents $r_{2j}, ..., r_{2j+k-1}$ interviewing with the set of hospitals $\{h_{2j}, ..., h_{2j+k-1}\}$ forms an equilibrium.*

6 Conclusion

We investigate equilibria for interviewing with a limited budget when master lists are present in the market. We provide a generic payoff function, that is indifferent to both the number of interviews provided and the distribution used, and use this payoff function to show that a pure strategy equilibrium always exists for this game.

We then focus on this game when residents' preferences are drawn from the same distribution and residents are allowed to interview with two hospitals. We show that there is a naturally arising equilibrium where the maximum number of residents are matched: residents assortatively interview in tiers, forming an n/k bipartite interviewing graph structure seen in work by Lee and Schwarz. However, this structure naturally arises in our model, and we characterize a very different preference space than the Lee and Schwarz paper, which investigates the impartial culture model.

This work raises a number of interesting questions. First, we believe that the bounds on the Mallows model parameters used to characterize the equilibria can be improved, and conjecture that the assortive equilibrium exists for $\phi < 0.46$. We also hypothesize that similar results also hold for different valuation functions (e.g. harmonic) and preference distributions (e.g. Plackett-Luce). Perhaps the most important direction for future work is relaxing the master lists assumption; we hypothesize that similar equilibria arise if preferences on both sides of the market are distributed according to a Mallows model with low dispersion.

References

[1] Kimberly D. Anderson, Marsha M. Dorough, Catherine R. Stein, Scott A. Optenberg, Ian M. Thompson, and others. The urology residency matching program in practice. *The Journal of Urology*, 163(6):1878–1887, 2000.

[2] Hector Chade, Gregory Lewis, and Lones Smith. Student portfolios and the college admissions problem. *The Review of Economic Studies*, 81(3):971–1002, 2014.

[3] Hector Chade and Lones Smith. Simultaneous search. *Econometrica*, 74(5):1293–1307, 2006.

[4] Peter Coles, Alexey Kushnir, and Muriel Niederle. Preference signaling in matching markets. Technical report, National Bureau of Economic Research, 2010.

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

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

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

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

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

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

[11] Robin S. Lee and Michael Schwarz. Interviewing in two-sided matching markets. Working Paper 14922, National Bureau of Economic Research, April 2009.

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

[13] R. Duncan Luce. *Individual Choice Behavior: A Theoretical Analysis*. Wiley, 1959.

[14] Colin L. Mallows. Non-null ranking models. *Biometrika*, 44:114–130, 1957.

[15] John Marden. *Analyzing and modeling rank data*. Chapman and Hall, London, 1995.

[16] Thomas Brendan Murphy and Donal Martin. Mixtures of distance-based models for ranking data. *Computational Statistics and Data Analysis*, 41:645–655, January 2003.

[17] Muriel Niederle, Alvin E. Roth, and Tayfun Sonmez. Matching and market design. In Steven N. Durlauf and Lawrence E. Blume, editors, *The New Palgrave Dictionary of Economics (2nd Ed.)*, volume 5, pages 436–445. Palgrave Macmillan, Cambridge, 2008.

[18] R. Plackett. The analysis of permutations. *Applied Statistics*, 24:193–202, 1975.

[19] National Resident Matching Program. National resident matching program, results and data: 2015 main residency match®. 2015.

[20] Baharak Rastegari, Anne Condon, Nicole Immorlica, and Kevin Leyton-Brown. Two-sided matching with partial information. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, pages 733–750. ACM, 2013.

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

[22] Min Zhu. College admissions in china: A mechanism design perspective. *China Economic Review*, 30:618–631, September 2014.

Joanna Drummond
Department of Computer Science
University of Toronto
Toronto, Ontario, Canada
Email: jdrummond@cs.toronto.edu

Allan Borodin
Department of Computer Science
University of Toronto
Toronto, Ontario, Canada
Email: bor@cs.toronto.edu

Kate Larson
Cheriton School of Computer Science
University of Waterloo
Waterloo, Ontario, Canada
Email: klarson@uwaterloo.ca

A Proofs of Mallows Lemmas

Lemma 11 *Given some Mallows model $\mathscr{D}^{\phi,\sigma}$ with fixed dispersion parameter ϕ and reference ranking σ given two alternatives a_i and a_j such that $a_i \succ_\sigma a_j$, then the probability that a ranking η is drawn from $\mathscr{D}^{\phi,\sigma}$ such that $a_i \succ_\eta a_j$ is equal to drawing from some distribution $\mathscr{D}^{\phi,\sigma'}$ where σ is a prefix of σ'. By symmetry, this proof also holds when σ is a suffix of σ'.*

Proof: First, 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 $Pr(a_i \succ a_j|\mathscr{D}^{\phi,\sigma'})$, and using algebraic manipulations to show this is equivalent to the definition of $Pr(a_i \succ a_j|\mathscr{D}^{\phi,\sigma})$.

$$Pr(a_i \succ a_j|\mathscr{D}^{\phi,\sigma'}) = \frac{\sum_{\eta' \in P(\{a_0,\ldots,a_{p-1},a_p\})^{a_i \succ a_j}} \phi^{d(\eta',\sigma')}}{1(1+\phi)...(1+...+\phi^{p-1}+\phi^p)} \tag{10}$$

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

$$
\begin{aligned}
Pr(a_i \succ a_j | \mathscr{D}^{\phi,\sigma'}) &= \frac{\sum_{\eta \in P(\{a_0,\dots,a_{p-1}\})^{a_i \succ a_j}} \sum_{l=0}^{p} \phi^{d(\eta,\sigma)+l}}{1(1+\phi)\dots(1+\dots+\phi^p)} && (11)\\
&= \frac{\left[\sum_{\eta \in P(\{a_0,\dots,a_{p-1}\})^{a_i \succ a_j}} \phi^{d(\eta,\sigma)}\right]\left[\sum_{l=0}^{p} \phi^{l}\right]}{1(1+\phi)\dots(1+\dots+\phi^p)} && (12)\\
&= \frac{\left[\sum_{\eta \in P(\{a_0,\dots,a_{p-1}\})^{a_i \succ a_j}} \phi^{d(\eta,\sigma)}\right](1+\dots+\phi^p)}{1(1+\phi)\dots(1+\dots+\phi^{p-1})(1+\dots+\phi^p)} && (13)\\
&= \frac{\sum_{\eta \in P(\{a_0,\dots,a_{p-1}\})^{a_i \succ a_j}} \phi^{d(\eta,\sigma)}}{1(1+\phi)\dots(1+\dots+\phi^{p-1})} && (14)\\
&= Pr(a_i \succ a_j | \mathscr{D}^{\phi,\sigma}) && (15)
\end{aligned}
$$

□

Corollary 12 *Given any reference ranking σ and two alternatives a_i, a_j such that $rank(a_j,\sigma) = rank(a_i,\sigma)+1$, then $Pr(a_i \succ a_j | D^{\phi,\sigma}) = \frac{1}{1+\phi}$.*

Proof: Consider $\sigma = a_i \succ a_j$, a reference ranking with two elements in it. Then, the set of all potential rankings such that $a_i \succ a_j$ under $\mathscr{D}^{\phi,\sigma}$ 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 2, the probability that some η is drawn from $\mathscr{D}^{\phi,\gamma}$ such that $a_i \succ_\eta a_j$ is $\frac{1}{1+\phi}$ as required. □

Corollary 13 *Given any reference ranking σ and three alternatives a_w, a_x, a_y such that $rank(a_x,\sigma) = rank(a_w,\sigma)+1$ and $rank(a_y,\sigma) = rank(a_x,\sigma)+1$ and some $\eta \in P(\{a_w,a_x,a_y\})$, then the probability that we draw some ranking β consistent with η is: $Pr(\beta|\mathscr{D}^{\phi,\sigma}) = \frac{\phi^{d(\eta, a_w \succ a_x \succ a_y)}}{(1+\phi)(1+\phi+\phi^2)}$.*

Proof: Consider $\sigma^* = a_w \succ a_x \succ a_y$, a reference ranking with three elements in it. The set of all potential rankings under $\mathscr{D}^{\phi,\sigma^*}$ such that $a_w \succ a_x \succ a_y$ is solely that ranking. Using the same argument as in Lemma 2, we note that creating some new reference ranking $\gamma = \sigma' \succ \sigma^* \succ \sigma''$ and drawing from $\mathscr{D}^{\phi,\gamma}$ does not change the likelihood that we draw a ranking consistent with $a_w \succ a_x \succ a_y$.

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

B Convergence Proof

Lemma 14 *Given $0 < \phi < 1$:*

$$\sum_{w=0}^{\infty}\sum_{a=w}^{\infty}\sum_{b=a}^{\infty} \phi^{w+a+b} \le \frac{1}{3(1-\phi)^3} + \frac{2}{3} \tag{16}$$

As noted in the paper, as $n \to \infty$, $\sum_{w=0}^{\infty}\sum_{a=0}^{\infty}\sum_{b=0}^{\infty} \phi^{w+a+b} \to \frac{1}{(1-\phi)^3}$. However, by dropping a and b all the way down to 0, we are counting many items multiple times. We want to show that:

$$\sum_{w=0}^{\infty}\sum_{a=w}^{\infty}\sum_{b=a}^{\infty} \phi^{w+a+b} \le \frac{2}{3} + \frac{1}{3}\sum_{w=0}^{\infty}\sum_{a=0}^{\infty}\sum_{b=0}^{\infty} \phi^{w+a+b} \tag{17}$$

We investigate 3 cases based on the indices of w, a, b: when all three are different, when exactly two are equal, and when all are the same.
Case 1: $w \neq a \neq b$. Let $w = r,\ a = s,\ b = t$. The only valid permutation is $r < s < t$, as in $\sum_{w=0}^{\infty}\sum_{a=w}^{\infty}\sum_{b=a}^{\infty}\phi^{w+a+b}$. However, as $\sum_{w=0}^{\infty}\sum_{a=0}^{\infty}\sum_{b=0}^{\infty}\phi^{w+a+b}$ allows for all permutations of r, s, t, we end up adding ϕ^{r+s+t} 6 times.
Case 2: Exactly 2 of w, a, b are equal. Then, there are two integers r, s such that $r < s$ and the two valid permutations for $\sum_{w=0}^{\infty}\sum_{a=w}^{\infty}\sum_{b=a}^{\infty}\phi^{w+a+b}$ are: $w = r,\ a = r,\ b = s$ and $w = r,\ a = s,\ b = s$. One of these permutations adds up to $2r + s$, and the other adds up to $r + 2s$. However, when summing over all possible permutations, there are three permutations that add up to $2r + s$, and three that add up to $r + 2s$, so we are adding ϕ^{2r+s} three extra times, and ϕ^{r+2s} three extra times.
Case 3: $w = a = b$. Note that in this case there are exactly as many permutations in both. However, note that we have added 6 times as many ϕ^{r+s+t} in Case 1. Then, for all $w > 0$ such that $w = a = b$, there exists some r, s, t such that $r + s + t = 3w$. Looking at all 6 of those permutations such that $r + s + t = 3w$ and adding in $3w$, we now have 7 permutations for $\sum_{w=0}^{\infty}\sum_{a=0}^{\infty}\sum_{b=0}^{\infty}\phi^{w+a+b}$ when we should have 2. As $7/2 > 3$, and Case 2 also has 3 times as many permutations, we have added $3\phi^{w+a+b}$ more than we should have for any w, a, b except for $w = a = b = 0$.

This gives us:

$$3\Big(\sum_{w=0}^{\infty}\sum_{a=w}^{\infty}\sum_{b=a}^{\infty}\phi^{w+a+b}\Big) - 2\phi^{0+0+0} \leq \sum_{w=0}^{\infty}\sum_{a=0}^{\infty}\sum_{b=0}^{\infty}\phi^{w+a+b} \tag{18}$$

$$3\Big(\sum_{w=0}^{\infty}\sum_{a=w}^{\infty}\sum_{b=a}^{\infty}\phi^{w+a+b}\Big) - 2 \leq \sum_{w=0}^{\infty}\sum_{a=0}^{\infty}\sum_{b=0}^{\infty}\phi^{w+a+b} \tag{19}$$

$$\sum_{w=0}^{\infty}\sum_{a=w}^{\infty}\sum_{b=a}^{\infty}\phi^{w+a+b} \leq \frac{2}{3} + \frac{1}{3}\sum_{w=0}^{\infty}\sum_{a=0}^{\infty}\sum_{b=0}^{\infty}\phi^{w+a+b} \tag{20}$$

$$\sum_{w=0}^{\infty}\sum_{a=w}^{\infty}\sum_{b=a}^{\infty}\phi^{w+a+b} \leq \frac{2}{3} + \frac{1}{3(1-\phi)^3} \tag{21}$$

C Full proof of Theorem 5

Theorem 15 *Given residents' valuation function $v(h, \eta) = n - rank(h|\eta)$ (i.e., Borda score) for any ranking η and market size n, for a Mallows model with reference ranking $\sigma = h_0, h_1, ..., h_{n-1}$ with dispersion parameter ϕ such that $0 < \phi \leq 0.265074$, resident r_1 maximizes her expected payoff by interviewing with $\{h_0, h_1\}$.*

Proof: As resident r_0 greedily chooses to interview with $\{h_0, h_1\}$, we note that resident r_1 must calculate a trade-off between a higher expected value for hospitals in a potential interviewing set, and competition for those hospitals. We prove that r_1 maximizes her expected payoff in this interval by bounding the difference in expected payoff between interviewing sets: $u_{r_1}(h_0, h_1|D^{\phi^*,\sigma}) - u_{r_1}(h_i, h_j|D^{\phi^*,\sigma}) \geq 0$, for all i, j.

Note that for any pair of alternatives h_i, h_j such that h_i and h_j have the same probability of being available, if $h_i \succ_\sigma h_j$, interviewing with h_i dominates interviewing with h_j. When applied to elements of an interviewing set, these swaps create a chain of dominated sets. Instead of comparing all $\binom{n}{k}$ sets to find the best interviewing set, we can just compare the undominated ones. Thus, for r_1, the set of undominated interviewing sets is: $\{h_0, h_1\}, \{h_0, h_2\}, \{h_1, h_2\}, \{h_2, h_3\}$. Intuitively, the difference between these sets is a trade-off between a higher expected valuation of the hospitals, versus more competition with resident r_0. We compute lower bounds comparing the difference in expected utility between interviewing with $\{h_0, h_1\}$ and each of the potential interview sets independently. We present the case for $\{h_1, h_2\}$ (as we find it imposes the tightest bound), but leave the remainder of the set to Section E for clarity, as the arguments are analogous.

We prove that choosing $\{h_0, h_1\}$ is better than choosing $\{h_1, h_2\}$, 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_0, h_1, h_2. We then pair these summed permutations in such a manner that makes it easy to find a lower bound for $u_{r_1}(\{h_0, h_1\}) - u_{r_1}(\{h_1, h_2\})$. 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_0, h_1, h_2 as follows:
Case 1: all rankings η consistent with $h_1 \succ h_0 \succ h_2$ or η' consistent with $h_1 \succ h_2 \succ h_0$;
Case 2: all rankings η consistent with $h_0 \succ h_1 \succ h_2$ or η' consistent with $h_2 \succ h_1 \succ h_0$;
Case 3: all rankings η consistent with $h_0 \succ h_2 \succ h_1$ or η' consistent with $h_2 \succ h_0 \succ h_1$.

Note that as we have enumerated all possible permutations of h_0, h_1, h_2, these three cases generate every ranking in $P(H)$. Furthermore, for any one of the three cases, we can iterate over only all possible rankings η that are consistent with the first member of the pair, and generate the ranking η' consistent with the second member of the pair by simply swapping two alternatives in the rank. Moreover, given some η, the number of discordant pairs in η' is simply the number in η, plus the number of additional discordant pairs between h_0, h_1, h_2 caused by swapping the two alternatives.

For clarity, let $u_{r_1}(\{h_0, h_1\}) - u_{r_1}(\{h_1, h_2\}) = U_1 + U_2 + U_3$, where U_1, U_2, U_3 correspond to our three cases. We also introduce the notation $Pr_{\mu(r_i)}(h)$ to denote the probability that r_i is matched to hospital h under matching μ. That is, $Pr_{\mu(r_i)}(h) = Pr(\mu(r_i) = h)$.

Case 1. Because we have fixed $h_1 \succ h_0 \succ h_2$ or $h_1 \succ h_2 \succ h_0$, we know exactly what r_1's match will be, given h_0's match. Additionally, given any η consistent with $h_1 \succ h_0 \succ h_2$, we generate η' consistent with $h_1 \succ h_2 \succ h_0$ by letting $\text{rank}(h_0, \eta) = \text{rank}(h_2, \eta')$ and $\text{rank}(h_2, \eta) = \text{rank}(h_0, \eta')$. This adds one additional discordant pair, so $d(\eta', \sigma) = d(\eta, \sigma) + 1$.

We first note that if r_0 is assigned h_0, r_1 gets h_1 in either interviewing set, and thus the difference in expected utility is 0. Focusing on when r_0 is assigned h_1, we calculate the difference in r_1's expected utility between the two interviewing sets:

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

By construction, we can rewrite all η' in terms of η, switching the rank as constructed, and adding 1 to all distances (as $d(h_1 \succ h_0 \succ h_2, h_1 \succ h_2 \succ h_0) = 1$):

$$U_1 = \sum_{\eta \in P(H)^{h_1 \succ h_0 \succ h_2}} Pr_{\mu(r_0)}(h_1) \Big[(v(h_0, \eta) - v(h_2, \eta)) \frac{\phi^{d(\eta, \sigma)}}{Z} + (v(h_2, \eta) - v(h_0, \eta)) \frac{\phi^{d(\eta, \sigma)+1}}{Z} \Big] \quad (22)$$

$$= \sum_{\eta \in P(H)^{h_1 \succ h_0 \succ h_2}} Pr_{\mu(r_0)}(h_1) (v(h_0, \eta) - v(h_2, \eta)) \frac{\phi^{d(\eta, \sigma)}}{Z} (1 - \phi) \quad (23)$$

Note that by Corollaries 3 and 4, $\sum_{\eta \in P(H)^{h_1 \succ h_0 \succ h_2}} \frac{\phi^{d(\eta, \sigma)}}{Z} = Pr(h_1 \succ h_0 \succ h_2)$, and $Pr_{\mu(r_0)}(h_1) = \frac{\phi}{1+\phi}$. Also note that by construction, everything in Eq. 23 is positive. Therefore, $v(h_0, \eta) - v(h_2, \eta) \geq 1$, which implies that:

$$U_1 \geq Pr_{\mu(r_0)}(h_1)(1)(1 - \phi) Pr(h_1 \succ h_0 \succ h_2) \quad (24)$$

$$= \left(\frac{\phi}{1 + \phi} \right) \left(\frac{\phi}{(1 + \phi)(1 + \phi + \phi^2)} \right) (1 - \phi) \quad (25)$$

Case 2. For this case we fix $h_0 \succ h_1 \succ h_2$ or $h_2 \succ h_1 \succ h_0$. Again, we list all η such that $h_0 \succ h_1 \succ h_2$, and transform that into an η' that is identical to η, except except $\text{rank}(h_0, \eta') = \text{rank}(h_2, \eta)$ and $\text{rank}(h_2, \eta') = \text{rank}(h_0, \eta)$. Note that $d(\eta', \sigma) = d(\eta, \sigma) + 3$.

$$\begin{aligned}
U_2 = &\sum_{\eta \in P(H)^{h_0 \succ h_1 \succ h_2}} Pr_{\mu(r_0)}(h_0)\left[(0)Pr(\eta|\mathscr{D}^{\phi,\sigma}) + (v(h_1,\eta') - v(h_2,\eta'))Pr(\eta'|\mathscr{D}^{\phi,\sigma})\right] \\
&+ Pr_{\mu(r_0)}(h_1)\left[(v(h_0,\eta) - v(h_2,\eta))Pr(\eta|\mathscr{D}^{\phi,\sigma}) + (v(h_0,\eta') - v(h_2,\eta'))Pr(\eta'|\mathscr{D}^{\phi,\sigma})\right] \\
= &\sum_{\eta \in P(H)^{h_0 \succ h_1 \succ h_2}} Pr_{\mu(r_0)}(h_0)(v(h_1,\eta) - v(h_0,\eta))\left(\frac{\phi^{d(\eta,\sigma)+3}}{Z}\right) \\
&+ Pr_{\mu(r_0)}(h_1)\left[(v(h_0,\eta) - v(h_2,\eta))\left(\frac{\phi^{d(\eta,\sigma)}}{Z}\right) + (v(h_2,\eta) - v(h_0,\eta))\left(\frac{\phi^{d(\eta,\sigma)+3}}{Z}\right)\right]
\end{aligned}$$

By the definition of the valuation function v, $v(h_1, \eta) - v(h_0, \eta) \geq v(h_2, \eta) - v(h_0, \eta)$. Then:

$$\begin{aligned}
U_2 \geq &\sum_{\eta \in P(H)^{h_0 \succ h_1 \succ h_2}} Pr_{\mu(r_0)}(h_0)(v(h_2,\eta) - v(h_0,\eta))\frac{\phi^{d(\eta,\sigma)+3}}{Z} \\
&+ Pr_{\mu(r_0)}(h_1)\,(v(h_0,\eta) - v(h_2,\eta))\left(\frac{\phi^{d(\eta,\sigma)} - \phi^{d(\eta,\sigma)+3}}{Z}\right) &(26) \\
= &\sum_{\eta \in P(H)^{h_0 \succ h_1 \succ h_2}} \frac{\phi^{d(\eta,\sigma)}}{Z}\Big[\frac{1}{1+\phi}(-\phi^3) + \frac{\phi}{1+\phi}(1-\phi^3)\Big](v(h_0,\eta) - v(h_2,\eta)) &(27) \\
= &\sum_{\eta \in P(H)^{h_0 \succ h_1 \succ h_2}} \frac{\phi^{d(\eta,\sigma)}}{Z(1+\phi)}(v(h_0,\eta) - v(h_2,\eta))(\phi - \phi^3 - \phi^4) &(28)
\end{aligned}$$

Note that if $0 < \phi \leq 0.7548$, $\phi - \phi^3 - \phi^4 \geq 0$, all terms in Eq 28 are positive, as $v(h_0, \eta) - v(h_2, \eta) \geq 2$. We thus impose our first restriction on ϕ; we now only look at the range where $0 \leq \phi \leq 0.7548$. Making the substitution that $v(h_0, \eta) - v(h_2, \eta) \geq 2$ by construction, we get the following:

$$U_2 \geq Pr(h_0 \succ h_1 \succ h_2)\frac{2}{1+\phi}(\phi - \phi^3 - \phi^4) \tag{29}$$

Case 3. We fix $h_0 \succ h_2 \succ h_1$ or $h_2 \succ h_0 \succ h_1$. Again, we look at pairs of rankings η, η', where η is consistent with $h_0 \succ h_2 \succ h_1$, and η' is identical to η, except $\text{rank}(h_0, \eta) = \text{rank}(h_2, \eta')$, and $\text{rank}(h_2, \eta) = \text{rank}(h_0, \eta')$.

Then, as before, we sum over all possible rankings consistent with $h_0 \succ h_2 \succ h_1$:

$$\begin{aligned}
U_3 = &\sum_{\eta \in P(H)^{h_0 \succ h_2 \succ h_1}} Pr_{\mu(r_0)}(h_0)[(v(h_1,\eta) - v(h_2,\eta))Pr(\eta|\mathscr{D}^{\phi,\sigma}) + (v(h_1,\eta') - v(h_2,\eta'))Pr(\eta'|\mathscr{D}^{\phi,\sigma})] \\
&+ Pr_{\mu(r_0)}(h_1)[(v(h_0,\eta) - v(h_2,\eta))Pr(\eta|\mathscr{D}^{\phi,\sigma}) + (v(h_0,\eta') - v(h_2,\eta')Pr(\eta'|\mathscr{D}^{\phi,\sigma})]
\end{aligned}$$

We break this equation into two subcases, so that $U_3 = U_{3a} + U_{3b}$:

$$U_{3a} = \sum_{\eta \in P(H)^{h_0 \succ h_2 \succ h_1}} Pr_{\mu(r_0)}(h_0)[(v(h_1,\eta) - v(h_2,\eta))Pr(\eta|\mathscr{D}^{\phi,\sigma}) + (v(h_1,\eta') - v(h_2,\eta'))Pr(\eta'|\mathscr{D}^{\phi,\sigma})]$$

$$U_{3b} = \sum_{\eta \in P(H)^{h_0 \succ h_2 \succ h_1}} Pr_{\mu(r_0)}(h_1)[(v(h_0,\eta) - v(h_2,\eta))Pr(\eta|\mathscr{D}^{\phi,\sigma}) + (v(h_0,\eta') - v(h_2,\eta'))Pr(\eta'|\mathscr{D}^{\phi,\sigma})]$$

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

$$U_{3b} = \sum_{\eta \in P(H)^{h_0 \succ h_2 \succ h_1}} Pr_{\mu(r_0)}(h_1)[(v(h_0,\eta) - v(h_2,\eta))\frac{\phi^{d(\eta,\sigma)}}{Z} + (v(h_2,\eta) - v(h_0,\eta))\frac{\phi^{d(\eta,\sigma)+1}}{Z} \quad (30)$$

$$= \sum_{\eta \in P(H)^{h_0 \succ h_2 \succ h_1}} Pr_{\mu(r_0)}(h_1)(v(h_0,\eta) - v(h_2,\eta))[\frac{\phi^{d(\eta,\sigma)}}{Z} - \frac{\phi^{d(\eta,\sigma)+1}}{Z}] \quad (31)$$

$$\geq \frac{\phi}{\phi+1}(1-\phi)Pr(h_0 \succ h_2 \succ h_1) \quad (32)$$

Case U_{3a} is different from all other cases, in that *all* terms are negative. Furthermore, we note that as a function of n (keeping ϕ constant), the equation in U_{3a} is monotonically decreasing. Thus, if this function converges as $n \to \infty$, we have found a lower bound for U_{3a} for all n.

We analyze $v(h_1,\eta) - v(h_2,\eta)$ independently of $v(h_1,\eta') - v(h_2,\eta')$, though the analysis is symmetrical. Take $v(h_1,\eta) - v(h_2,\eta)$. Intuitively, we sum over all permutations of $h_3, ..., h_{n-1}$, and place h_0, h_1, h_2 in all possible indices consistent with the ranking $h_0 \succ h_2 \succ h_1$. This generates all rankings in the set $P(H)^{h_0 \succ h_1 \succ h_2}$. Let $\sigma^* = h_3 \succ h_4 \succ ... \succ h_{n-1}$ (σ with elements h_0, h_1, h_2 removed).

We note that we can calculate the number of discordant pairs via the indices of h_0, h_1, h_2. Given some permutation $\gamma \in P(H \setminus \{h_0, h_1, h_2\})$, suppose we insert h_0, h_1, h_2 into γ such that rank(h_0) $= w$; rank(h_2) $= x$; rank(h_1) $= y$. Then, there are w alternatives from γ before h_0, $x-1$ alternatives from γ before h_2, and $y-2$ alternatives from γ before h_1. As every item in γ before h_0, h_1, or h_2 causes a discordant pair, we get $w + x - 1 + y - 2$ discordant pairs due to our indices, and 1 discordant pair from $h_2 \succ h_1$, giving us a total of $w + x + y - 2$ discordant pairs. Counting the discordant pairs in this manner we get:

$$\sum_{\eta \in P(H)^{h_0 \succ h_2 \succ h_1}} (v(h_1,\eta) - v(h_2,\eta))\frac{\phi^{d(\eta,\sigma)}}{Z} = \sum_{\gamma}\sum_{w=0}^{n-2}\sum_{x=w+1}^{n-1}\sum_{y=x+1}^{n}\frac{1}{Z}(x-y)\phi^{d(\gamma,\sigma*)+w+x+y-2} \quad (33)$$

To simplify the exponent, we make a substitution. Let $x = a+1$, and let $y = b+2$. Then, for $x \in \{1, ..., n-1\}$, $a \in \{0, ..., n-2\}$, and for $y \in \{2, ..., n\}$, $b \in \{0, ..., n-2\}$. We use this substitution, and then consider the worst-case of $(a-b-1)$ by setting $a = 0$:

$$\frac{1}{Z}\sum_{w=0}^{n-2}\sum_{x=w+1}^{n-1}\sum_{y=x+1}^{n}(x-y)\phi^{d(\gamma,\sigma*)+w+x+y-2} = \frac{1}{Z}\sum_{w=0}^{n-2}\sum_{a=w}^{n-2}\sum_{b=a}^{n-2}(a-b-1)\phi^{d(\gamma,\sigma*)+w+a+b+1} \quad (34)$$

$$\geq \frac{1}{Z}\sum_{w=0}^{n-2}\sum_{a=w}^{n-2}\sum_{b=a}^{n-2}(-1-b)\phi^{d(\gamma,\sigma*)+w+a+b+1} \quad (35)$$

Note that $\sum_{\gamma \in P(H \setminus \{h_0,h_1,h_2\})} \phi^{d(\gamma,\sigma*)} = (1+\phi)(1+\phi+\phi^2)...(1+\phi+...+\phi^{n-4})$, and $Z = (1+\phi)...(1+...+\phi^{n-4})(1+...+\phi^{n-3})(1+....+\phi^{n-2})(1+...+\phi^{n-1})$. Then, since $(1+\phi)(1+\phi+\phi^2) \leq (1+...+\phi^{n-3})(1+...+\phi^{n-2})(1+...+\phi^{n-1})$, for $n \geq 3$, $\frac{\sum_{\gamma \in P(H \setminus \{h_0,h_1,h_2\})} \phi^{d(\gamma,\sigma*)}}{Z} \leq \frac{1}{(1+\phi)(1+\phi+\phi^2)}$, allowing us to further simplify the bound.

We also note that $-1-b$ is always negative for all values of $w, a, b > 0$. Thus summing from $a = 0$ (resp. $b = 0$) to $n-2$ is a lower bound for summing from $a = w$ (resp. $b = a$)

to $n-2$. We then simplify Eq 33 by using the substitution for $\sum_{\gamma \in P(H\setminus\{h_0,h_1,h_2\})} \frac{\phi^{d(\gamma,\sigma^*)}}{Z}$, and summing from 0:

$$\frac{1}{(1+\phi)(1+\phi+\phi^2)} \sum_{w=0}^{n-2}\sum_{a=w}^{n-2}\sum_{b=a}^{n-2}(-b-1)\phi^{w+a+b+1}$$
$$\geq \frac{-\phi}{(1+\phi)(1+\phi+\phi^2)}\big[\big(\sum_{w=0}^{n-2}\sum_{a=0}^{n-2}\sum_{b=0}^{n-2} b\phi^{w+a+b}\big) + \big(\sum_{w=0}^{n-2}\sum_{a=w}^{n-2}\sum_{b=a}^{n-2}\phi^{w+a+b}\big)\big] \tag{36}$$

Further note that:

$$\sum_{w=0}^{n-2}\sum_{a=0}^{n-2}\sum_{b=0}^{n-2} b\phi^{w+a+b} = \big(\sum_{w=0}^{n-2}\phi^w\big)\big(\sum_{a=0}^{n-2}\phi^a\big)\big(\sum_{b=0}^{n-2} b\phi^b\big) \tag{37}$$

As $n \to \infty$, we know that $\sum_{a=0}^{n-2}\phi^a$ converges to $\frac{1}{1-\phi}$, as it is simply a geometric series. It is also well known that $\sum_{b=0}^{n-2} b\phi^b$ converges to $\frac{\phi}{(1-\phi)^2}$.

To get a tighter bound on $\sum_{w=0}^{n-2}\sum_{a=w}^{n-2}\sum_{b=a}^{n-2}\phi^{w+a+b}$, we do a bit more analysis. First, we note the following:

$$\sum_{w=0}^{n-2}\sum_{a=w}^{n-2}\sum_{b=a}^{n-2}\phi^{w+a+b} \leq \sum_{w=0}^{n-2}\sum_{a=0}^{n-2}\sum_{b=0}^{n-2} \to \frac{1}{(1-\phi)^3} \tag{38}$$

When summing from $a=0$ (resp. $b=0$) instead of $a=w$ (resp. $b=a$), we count any given ϕ^{w+a+b} thrice (except for $w=a=b=0$). Intuitively, this is because we look at all permutations of the values w, b, a could take, instead of only those such that $w \leq a \leq b$ as required. The full proof is in Appendix B, where we show:

$$\sum_{w=0}^{\infty}\sum_{a=w}^{\infty}\sum_{b=a}^{\infty}\phi^{w+a+b} \leq \frac{1}{3(1-\phi)^3} + \frac{2}{3} \tag{39}$$

Thus, as $n \to \infty$, the RHS of Eq 36 converges, giving us:

$$\sum_{\eta \in P(H)^{h_0 \succ h_2 \succ h_1}} \frac{1}{Z}(v(h_1,\eta) - v(h_2,\eta))\phi^{d(\eta,\sigma)} \geq \frac{-\phi}{(1+\phi)(1+\phi+\phi^2)}\big(\frac{\phi}{(1-\phi)^4} + \frac{1}{3(1-\phi)^3} + \frac{2}{3}\big)$$

A lower bound on $v(h_1,\eta) - v(h_0,\eta)$ can be found identically, though switching from η to η' incurs an additional discordant pair, giving us:

$$\sum_{\eta \in P(H)^{h_0 \succ h_2 \succ h_1}} \frac{1}{Z}(v(h_1,\eta) - v(h_0,\eta))\phi^{d(\eta,\sigma)} \geq \frac{-\phi^2}{(1+\phi)(1+\phi+\phi^2)}\big(\frac{\phi}{(1-\phi)^4} + \frac{1}{3(1-\phi)^3} + \frac{2}{3}\big)$$

This then gives us the final bound for U_{3a}:

$$U_{3a} \geq Pr_{\mu(r_0)}(h_0)\frac{-\phi}{(1+\phi)(1+\phi+\phi^2)}\big(\frac{\phi}{(1-\phi)^4} + \frac{1}{3(1-\phi)^3} + \frac{2}{3}\big)(1+\phi) \tag{40}$$

We have considered all cases, and can now combine them together. We add the bounds for U_1 (Eq. 25), U_2 (Eq. 29), U_{3a} (Eq. 40), and U_{3b} (Eq. 32) giving us:

$$\begin{aligned} U_1 + U_2 + U_3 \geq\ & Pr_{\mu(r_0)}(h_1)(1-\phi)Pr(h_1 \succ h_0 \succ h_2) \\ & + Pr(h_0 \succ h_1 \succ h_2)\frac{2}{1+\phi}(\phi - \phi^3 - \phi^4) \\ & - Pr_{\mu(r_0)}(h_0)Pr(h_0 \succ h_2 \succ h_1)\left(\frac{\phi}{(1-\phi)^4} + \frac{1}{3(1-\phi)^3} + \frac{2}{3}\right)(1+\phi) \\ & + Pr_{\mu(r_0)}(h_1)(1-\phi)Pr(h_0 \succ h_2 \succ h_1) \end{aligned} \tag{41}$$

From Corollary 3, we know that $Pr_{\mu(r_0)}(h_0) = \frac{1}{1+\phi}$ and $Pr_{\mu(r_0)}(h_1) = \frac{\phi}{1+\phi}$. From Corollary 4, we know that $Pr(h_1 \succ h_0 \succ h_2) = Pr(h_0 \succ h_2 \succ h_1) = \frac{\phi}{(1+\phi)(1+\phi+\phi^2)}$ and $Pr(h_0 \succ h_1 \succ h_2) = \frac{1}{(1+\phi)(1+\phi+\phi^2)}$. Substituting this into Eq. 41, we get the final, full bound for $u_{r_1}(\{h_0, h_1\}) - u_{r_1}(\{h_1, h_2\})$:

$$\begin{aligned} u_{r_1}(\{h_0, h_1\}) - u_{r_1}(\{h_1, h_2\}) \geq\ & \frac{\phi^2}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}(1-\phi) \\ & + \frac{2}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}(\phi - \phi^3 - \phi^4) \\ & - \frac{\phi}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}\left(\frac{\phi}{(1-\phi)^4} + \frac{1}{3(1-\phi)^3} + \frac{2}{3}\right)(1+\phi) \\ & + \frac{\phi^2}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}(1-\phi) \end{aligned} \tag{42}$$

Thus, Eq. 42 gives us a lower bound for the difference in expected utility between $\{h_0, h_1\}$ and $\{h_1, h_2\}$ for resident r_1, for all n. Using numerical methods to approximate the roots of Eq. 42, we get that there is a root at 0, and a root at $\phi \approx 0.265074$.

We have now proven the bound showing that when $0 < \phi \leq 0.265074$, r_1 choosing the interview set $\{h_0, h_1\}$ dominates choosing the interview set $\{h_1, h_2\}$. In Section E, we provide bounds such that $u_{r_1}(\{h_1, h_2\}) - u_{r_1}(\{h_2, h_3\}) \geq 0$ if $0 < \phi < 0.3550107$, and that $u_{r_1}(\{h_0, h_1\}) - u_{r_1}(\{h_0, h_2\}) \geq 0$ if $0 < \phi < 0.413633$ (the proofs are analogous to the one presented here). Thus, for the interval $0 < \phi \leq 0.265074$, we have successfully shown that r_1's best move in this interval is to interview with $\{h_0, h_1\}$ as required. □

D Full Proof of Theorem 6

Theorem 16 *Given residents' valuation function $v(h, \eta) = n - rank(h|\eta)$ for any ranking η and market size n, for a Mallows model with dispersion parameter ϕ such that $0 < \phi < 0.1707951$, if all residents r_{2f}, r_{2f+1} have interviewed with hospitals h_{2f}, h_{2f+1} for $f < j$, then residents r_{2j}, r_{2j+1} will interview with hospitals $\{h_{2j}, h_{2j+1}\}$.*

Proof: We first note that for any hospital h_a such that $h_a \succ_\sigma h_{2j}$, interviewing with any other hospital dominates interviewing with h_a, because the probability r_{2j} or r_{2j+1} will be matched with h_a is 0, as h_a is already matched to a more desirable doctor. Likewise, interviewing with any alternative h_b such that $h_{2j+3} \succ_\sigma h_b$ is dominated by interviewing with h_{2j+3}.

Resident r_{2j} does best by greedily choosing the top two hospitals left, h_{2j} and h_{2j+1}. Resident r_{2j+1} must again investigate the following interviewing sets: $\{h_{2j}, h_{2j+1}\}, \{h_{2j+1}, h_{2j+2}\}, \{h_{2j+2}, h_{2j+3}\}, \{h_{2j}, h_{2j+2}\}$. We provide a proof of the comparison between $\{h_{2j}, h_{2j+1}\}$ and $\{h_{2j+1}, h_{2j+2}\}$, leaving the remainder to Appendix E for clarity.

We adapt the proof used in Theorem 5. We again break the expected payoff function into three subcases; $u^n_{r_{2j+1}}(\{h_{2j}, h_{2j+1}\}) - u^n_{r_{2j+1}}(\{h_{2j+1}, h_{2j+2}\}) = U^*_1 + U^*_2 + U^*_3$. For clarity, let $h_{2j} = a_0$; $h_{2j+1} = a_1$; $h_{2j+2} = a_2$. As in the proof for Theorem 5, we look at three cases, pairing all possible permutations of a_0, a_1, a_2 in the following manner:
Case 1: all rankings consistent with $a_1 \succ a_0 \succ a_2$ or $a_1 \succ a_2 \succ a_0$;
Case 2: all rankings consistent with $a_0 \succ a_1 \succ a_2$ or $a_2 \succ a_1 \succ a_0$;
Case 3: all rankings consistent with $a_0 \succ a_2 \succ a_1$ or $a_2 \succ a_0 \succ a_1$.
We again, for some fixed ranking η derive a ranking η' by substituting a_i for a_j in the ranking, to switch between the paired rankings.

Case 1. This case is completely analogous to Case 1 presented in Theorem 5. The keystone of the argument is that h_0, h_1, h_2 are all adjacent in the reference ranking σ, which is again the case with a_0, a_1, a_2. The minimum distance is again the same, and swapping a_0 and a_2 again gives us only one additional discordant pair (as a_0, a_1, a_2 are all adjacent). Therefore, $U^*_1 \geq \left(\frac{\phi}{1+\phi}\right)\left(\frac{\phi}{(1+\phi)(1+\phi+\phi^2)}\right)(1-\phi)$

Case 2. This case is likewise completely analogous to Case 2 presented in Theorem 5, for the same reason as above. Thus, $U^*_2 \geq Pr(a_0 \succ a_1 \succ a_2)\frac{2}{1+\phi}(\phi - \phi^3 - \phi^4)$.

Case 3. We again break this case up in to U^*_{3a} and U^*_{3b}. Again, for the same reasons as in Cases 1 and 2 in this proof, case U^*_{3b} is identical to the one provided in U_{3b} in Theorem 5. However, the bound calculated in U_{3a} requires that h_0, h_1, h_2 are in the first three indices in the reference ranking to accurately calculate the number of discordant pairs. We modify the bound shown in U_{3a} to be for a_0, a_1, a_2, but in doing so significantly loosen it. Given empirical findings (described in the next section), we believe it is likely that $U^* \geq U$, and thus our final bound could be tightened significantly.

To begin, we note that U^*_{3a} is quite similar to U_{3a}:

$$U^*_{3a} = \sum_{\eta \in P(H)^{a_0 \succ a_2 \succ a_1}} Pr_{\mu(r_{2j})}(a_0)[(v(a_1,\eta) - v(a_2,\eta))Pr(\eta|\mathscr{D}^{\phi,\sigma}) + (v(a_1,\eta') - v(a_2,\eta'))Pr(\eta'|\mathscr{D}^{\phi,\sigma})]$$

Again, U^*_{3a} is a monotonically decreasing function in n. To get a lower bound we again analyze convergence as $n \to \infty$. We likewise analyze $v(a_1,\eta) - v(a_2,\eta)$ independently of $v(a_1,\eta) - v(a_0,\eta)$, though the analysis is again symmetrical. We start with $v(a_1,\eta) - v(a_2,\eta)$. Note that by construction, we have $2j$ alternatives that are more desirable than a_0, a_1, a_2 in reference ranking σ. Fix some ranking $\gamma \in P(H \setminus \{a_0, a_1, a_2\})$.

In the proof for U_{3a}, we noted that h_0, h_1, h_2 were better than all other alternatives, and so we knew exactly how many additional discordant pairs we were adding. We do a similar argument here, but must be more careful, as there are $2j$ elements that are better than alternatives a_0, a_1, a_2. We again start by summing over all potential rankings $\gamma \in P(H \setminus \{a_0, a_1, a_2\})$. Again, let $\sigma^* = \sigma \setminus \{a_0, a_1, a_2\}$. We note that the order of γ is important for two reasons: first, calculating the number of discordant pairs within γ, and secondly, counting the number of discordant pairs between γ and a_2.

Let γ' be the set of all alternatives $\gamma_i \in \gamma$ such that under the reference ranking σ, $\gamma_i \succ_\sigma a_0$. Let γ'' be the set of all alternatives $\gamma_q \in \gamma$ such that under the reference ranking σ, $a_2 \succ_\sigma \gamma_q$. By construction, $|\gamma'| = 2j$. While the number of discordant pairs between $\gamma' \succ \gamma''$ and γ have changed, this does not affect the analysis for counting the discordant pairs when placing a_0, a_1, a_2 in different indices of γ. We now sum over all indices that we can place a_0, a_1, a_2 in under this new ranking $\gamma' \succ a_0 \succ a_2 \succ a_1 \succ \gamma''$:

$$\sum_{\eta \in P(H)^{a_0 \succ a_2 \succ a_1}} (v(a_1,\eta) - v(a_2,\eta) Pr(\eta|\mathscr{D}^{\phi,\sigma}) \geq \tag{43}$$

$$\sum_{w=-2j}^{n-2j-2} \sum_{x=w+1}^{n-2j-1} \sum_{y=x+1}^{n-2j} [(2j+x) - (2j+y)] \phi^{d(\gamma,\sigma*)+|w|+|y|+|x|-2} \tag{44}$$

$$= \sum_{w=-2j}^{n-2j-2} \sum_{x=w+1}^{n-2j-1} \sum_{y=x+1}^{n-2j} (x-y) \phi^{d(\gamma,\sigma*)+|w|+|y|+|x|-2} \tag{45}$$

Breaking this into two parts, one consisting of when w, x, y are positive, and one consisting of when w, x, y are negative, we get that:

$$\sum_{\eta \in P(H)^{a_0 \succ a_2 \succ a_1}} (v(a_1,\eta) - v(a_2,\eta)) Pr(\eta|\mathscr{D}^{\phi,\sigma}) \geq \frac{-2\phi}{(1+\phi)(1+\phi+\phi^2)} \left(\frac{\phi}{(1-\phi)^4} + \frac{\phi}{3(1-\phi)^3} + \frac{2}{3}\right) \tag{46}$$

Likewise:

$$\sum_{\eta \in P(H)^{a_0 \succ a_2 \succ a_1}} (v(a_1,\eta) - v(a_0,\eta)) Pr(\eta|\mathscr{D}^{\phi,\sigma}) \geq \frac{-2\phi^2}{(1+\phi)(1+\phi+\phi^2)} \left(\frac{\phi}{(1-\phi)^4} + \frac{\phi}{3(1-\phi)^3} + \frac{2}{3}\right) \tag{47}$$

Combining this all together, we find that this has a zero at roughly 0.1707951. Thus r_{2j+1} chooses to interview with $\{h_{2j}, h_{2j+1}\}$ over $\{h_{2j+1}, h_{2j+2}\}$ whenever $0 \leq \phi \leq 0.1707961$. As in Theorem 5, the interviewing set that imposes the tightest bound on ϕ is $\{h_{2j}, h_{2j+1}\}$; we leave the remainder of the calculations to the appendix, as they are symmetrical to those in Theorem 5 with the additional factor of 2 as added here. Therefore, h_{2j+1} chooses to interview with $\{h_{2j}, h_{2j+1}\}$ whenever $0 \leq \phi \leq 0.1707961$, as required □

E Symmetric Interview Sets Proofs

Lemma 17 *For resident r_1, interviewing with $\{h_1, h_2\}$ dominates interviewing with $\{h_2, h_3\}$ when $0 < \phi < 0.3550107$.*

We again show this by breaking the utility function into three cases:
Case 1: all rankings consistent with $h_2 \succ h_1 \succ h_3$ or $h_2 \succ h_3 \succ h_1$;
Case 2: all rankings consistent with $h_1 \succ h_2 \succ h_3$ or $h_3 \succ h_2 \succ h_1$;
Case 3: all rankings consistent with $h_1 \succ h_3 \succ h_2$ or $h_3 \succ h_1 \succ h_2$.

Note, for Case 1, there is no difference between choosing h_1, h_2, or h_2, h_3, because h_2 will always be chosen no matter what resident r_0 does.

For Case 2, if $\eta = h_1 \succ h_2 \succ h_3$ and $\eta' = h_3 \succ h_2 \succ h_1$, when $\mu(r_0) = h_0$, under η, we are looking at the difference between $v(h_1,\eta) - v(h_2,\eta)$, and under η', we look at the difference between $v(h_3,\eta) - v(h_2,\eta)$. Adding η and η' together gives us a total of: $Pr_{\mu(r_0)}(h_0)[v(h_1,\eta) - v(h_2,\eta)](1-\phi)$. When $\mu(r_0) = h_0$, there is no difference under η, but η' contributes $v(h_2,\eta) - v(h_1,\eta)$. Thus, η' contributes $-Pr_{\mu(r_0)}(h_1)[v(h_1,\eta) - v(h_2,\eta)]$. Combining this with when $\mu(r_0) = h_0$, we get the following contribution from Case 2:

$$Pr(h_1 \succ h_2 \succ h_3) \frac{1}{1+\phi} (1 - \phi^3 - \phi^4) \tag{48}$$

For Case 3, again when $\mu(r_0) = h_0$, with η we get $v(h_1, \eta) - v(h_3, \eta)$ and with η' we get $v(h_3, \eta) - v(h_1, \eta)$, giving us $Pr_{\mu(r_0)}(h_0)[v(h_1, \eta) - v(h_3, \eta)](1 - \phi)$. When $\mu(r_0) = h_1$, however, all terms are negative again. Under η, we get $v(h_3, \eta) - v(h_2, \eta)$ and under η' we get $v(h_1, \eta) - v(h_2, \eta)$. As in Case 3 in the main body of the paper, we simply let $n \to \infty$, and bound using the bound proved there. This means Case 3 contributes:

$$Pr(h_1 \succ h_2 \succ h_3)\big(Pr_{\mu(r_0)}(h_0)(1 - \phi) - Pr_{\mu(r_0)}(h_1)\big[\frac{\phi}{(1 - \phi)^4} + \frac{1}{3(1 - \phi)^3} + \frac{2}{3}(1 + \phi)\big] \tag{49}$$

This means that our bound is:

$$u_{r_1}(\{h_1, h_2\}) - u_{r_1}(\{h_2, h_3\}) \geq \frac{1}{(1 + \phi)^2(1 + \phi + \phi^2)}\big[(2 - \phi - \phi^3 - \phi^4) - (\phi + \phi^2)\big(\frac{\phi}{(1 - \phi)^4} + \frac{1}{3(1 - \phi)^3} + \frac{2}{3}\big)\big] \tag{50}$$

as required □.

Lemma 18 *For resident* r_1*, interviewing with* $\{h_0, h_1\}$ *dominates interviewing with* $\{h_0, h_2\}$ *when* $0 < \phi < 0.413633$.

We break the utility function into three cases:
Case 1: all rankings consistent with $h_0 \succ h_1 \succ h_2$ or $h_0 \succ h_2 \succ h_1$;
Case 2: all rankings consistent with $h_1 \succ h_0 \succ h_2$ or $h_2 \succ h_0 \succ h_1$;
Case 3: all rankings consistent with $h_1 \succ h_2 \succ h_0$ or $h_2 \succ h_1 \succ h_0$.

For Case 1, when $\mu(r_1) = h_1$, r_1 is indifferent between the two interviewing sets, so the difference is:

$$Pr(h_0 \succ h_1 \succ h_2)Pr_{\mu(r_0)}(h_0)(1 - \phi) \tag{51}$$

For Case 2, when $\mu(r_0) = h_0$, we again compare h_1 and h_2 standardly. When $\mu(r_0) = h_1$, r_1 is indifferent for η, and η' contributes $\phi[v(h_0, \eta) - v(h_1, \eta)] > \phi[v(h_2, \eta) - v(h_1, \eta)]$:

$$Pr(h_1 \succ h_0 \succ h_2)[2(1 - \phi) - 2\phi^2] \tag{52}$$

For Case 3, when $\mu(r_0) = h_0$, we have the standard case again. When $\mu(r_0) = h_1$, we again need to bound as $n \to \infty$. Also note, $d(\eta, \sigma) = d(\eta', \sigma)$ in this case.

$$Pr(h_1 \succ h_2 \succ h_0)\big[-2Pr_{\mu(r_0)}(h_1)\big(\frac{\phi}{(1 - \phi)^4} + \frac{1}{3(1 - \phi)^2} + \frac{2}{3}\big)\big] \tag{53}$$

We combine this all to get:

$$\frac{1}{(1 + \phi)(1 + \phi + \phi^2)}\big[1 + \phi - 2\phi^2 - 2\phi^3 - 2\phi^3\big(\frac{\phi}{(1 - \phi)^4} + \frac{1}{3(1 - \phi)^3} + \frac{2}{3}\big)\big] \tag{54}$$

As required □

Lemma 19 *When* r_{2f}, r_{2f+1} *all interview with* h_{2f}, h_{2f+1}*, both the interview sets* a_0, a_2 *and* a_2, a_3 *are dominated by* a_0, a_1*.*

This proof is identical to the one presented in Theorem 6. Using the two functions presented in Lemmas 17 and 18, we modify them in the same way as we did in Theorem 6: we keep all positive terms identical, and double any terms that we let go to infinity.

For a_2, a_3, this means that the difference between a_1, a_2 and a_2, a_3 is:

$$u_{r_{2j+1}}(\{a_1, a_2\}) - u_{r_{2j+1}}(\{a_2, a_3\}) \geq \frac{1}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}[(1-\phi^3-\phi^4)+(1-\phi)-2\phi(\frac{\phi}{(1-\phi)^4}+\frac{1}{3(1-\phi)^3}+\frac{2}{3})(1+\phi)] \tag{55}$$

This difference is greater than 0 whenever $0 < \phi < 0.296649$.

Likewise, for a_0, a_2, the difference between a_0, a_1 and a_0, a_2 is:

$$u_{r_{2j+1}}(\{a_0, a_1\}) - u_{r_{2j+1}}(\{a_0, a_2\}) \geq \frac{1}{(1+\phi)(1+\phi)(1+\phi+\phi^2)}[2+\phi-2\phi^2(\frac{\phi}{(1-\phi)^4}+\frac{1}{3(1-\phi)^3}+\frac{2}{3})] \tag{56}$$

This difference is greater than 0 whenever $0 < \phi < 0.439098$.