Stable Matching: Choosing Which Proposals to Make

Ishan Agarwal
New York University

Richard Cole
New York University

Abstract

To guarantee all agents are matched in general, the classic Deferred Acceptance algorithm needs complete preference lists. In practice, preference lists are short, yet stable matching still works well. This raises two questions:

- Why does it work well?
- Which proposals should agents include in their preference lists?

We study these questions in a model, introduced by Lee [17], with preferences based on correlated cardinal utilities: these utilities are based on common public ratings of each agent together with individual private adjustments. Lee showed that for suitable utility functions, in large markets, with high probability, for most agents, all stable matchings yield similar valued utilities. By means of a new analysis, we strengthen Lee's result, showing that in large markets, with high probability, for *all* but the agents with the lowest public ratings, all stable matchings yield similar valued utilities. We can then deduce that for *all* but the agents with the lowest public ratings, each agent has an easily identified length $O(\log n)$ preference list that includes all of its stable matches, addressing the second question above. We note that this identification uses an initial communication phase.

We extend these results to settings where the two sides have unequal numbers of agents, to many-to-one settings, e.g. employers and workers, and we also show the existence of an ϵ-Bayes-Nash equilibrium in which every agent makes relatively few proposals. These results all rely on a new technique for sidestepping the conditioning between the tentative matching events that occur over the course of a run of the Deferred Acceptance algorithm. We complement these theoretical results with an experimental study.

2012 ACM Subject Classification Mathematics of computing → Graph algorithms; Mathematics of computing → Matchings and factors; Theory of computation → Random network models

Keywords and phrases Stable matching, randomized analysis

Digital Object Identifier 10.4230/LIPIcs.ICALP.2023.37

Related Version Full version hosted on arXiv
Full Version: `link to the full paper`

Funding This work was supported in part by NSF Grant CCF-1909538.

Acknowledgements We thank the reviewers for their helpful feedback.

1 Introduction

Consider a doctor applying for residency positions. Where should she apply? To the very top programs for her specialty? Or to those where she believes she has a reasonable chance of success (if these differ)? And if the latter, how does she identify them? We study these questions in the context of Gale and Shapley's deferred acceptance (DA) algorithm [5]. It is well-known that in DA the optimal strategy for the proposing side is to list their choices in order of preference. However, this does not address which choices to list.

The DA algorithm is widely used to compute matchings in real-world applications: the National Residency Matching Program (NRMP), which matches future residents to hospital programs [25]; university admissions programs which match students to programs, e.g. in Chile [24], school choice programs, e.g. for placement in New York City's high schools [1], the Israeli psychology Masters match [9], and no doubt many others (e.g. [7]).

Recall that each agent provides the mechanism a list of its possible matches in preference order, including the possibility of "no match" as one of its preferences. These mechanisms promise that the output will be a stable matching with respect to the submitted preference lists. In practice, preference lists are relatively short. This may be directly imposed by the mechanism or could be a reflection of the costs—for example, in time or money—of determining these preferences. Note that a short preference list is implicitly stating that the next preference after the listed ones is "no match".

Thus it is important to understand the impact of short preference lists. Roth and Peranson observed that the NRMP data showed that preference lists were short compared to the number of programs and that these preferences yielded a single stable partner for most participants; we note that this single stable partner could be the "no match" choice, and in fact this is the outcome for a constant fraction of the participants. They also confirmed this theoretically for the simplest model of uncorrelated random preferences; namely that with the preference lists truncated to the top O(1) preferences, almost all agents have a unique stable partner. Subsequently, in [10] the same result was obtained in the more general popularity model which allows for correlations among different agents' preferences; in their model, the first side—men—can have arbitrary preferences; on the second side—women—preferences are selected by weighted random choices, the weights representing the "popularity" of the different choices. These results were further extended by Kojima and Parthak in [15].

The popularity model does not capture behavior in settings where bounds on the number of proposals lead to proposals being made to plausible partners, i.e. partners with whom one has a realistic chance of matching. One way to capture such settings is by way of tiers [2], also known as block correlation [4]. Here agents on each side are partitioned into tiers, with all agents in a higher tier preferred to agents in a lower tier, and with uniformly random preferences within a tier. Tiers on the two sides may have different sizes. If we assign tiers successive intervals of ranks equal to their size, then, in any stable matching, the only matches will be between agents in tiers whose rank intervals overlap.

A more nuanced way of achieving these types of preferences bases agent preferences on cardinal utilities; for each side, these utilities are functions of an underlying common assessment of the other side, together with idiosyncratic individual adjustments for the agents on the other side. These include the separable utilities defined by Ashlagi, Braverman, Kanoria and Shi in [2], and another class of utilities introduced by Lee in [17]. This last model will be the focus of our study.

To make this more concrete, we review a simple special case of Lee's model, the *linear separable model.* Suppose that there are n men and n women seeking to match with each other. Each man m has a public rating r_m, a uniform random draw from $[0,1]$. These ratings can be viewed as the women's joint common assessment of the men. In addition, each woman w has an individual adjustment, which we call a score, $s_w(m)$ for man m, again a uniform random draw from $[0,1]$. All the draws are independent. Woman w's utility for man m is given by $\frac{1}{2}[r_m + s_w(m)]$; her full preference list has the men in decreasing utility order. The men's utilities are defined similarly.

Lee stated that rather than being assumed, short preference lists should arise from the model; this appears to have been a motivation for the model he introduced. A natural first step would be to show that for some or all stable matchings, the utility of each agent can be well-predicted, for this would then allow the agents to limit themselves to the proposals achieving such a utility. Lee proved an approximate version of this statement, namely that with high probability (w.h.p., for short) most agents obtain utility within a small ϵ of an easily-computed individual benchmark. However, this does not imply that agents can restrict their proposals to a reduced utility range. (See the paragraph preceding Definition 4 for the specification of the benchmarks.)

Our work seeks to resolve this issue. We obtain the following results. Note that in these results, when we refer to the bottommost agents, we mean when ordered by decreasing public rating. Also, we let the term loss mean the difference between an agent's benchmark utility and their achieved utility.

1. We show that in the linearly separable model, for any constant $c > 0$, with probability $1 - 1/n^c$, in every stable matching, apart from a sub-constant σ fraction of the bottommost agents, *all* the other agents obtain utility equal to an easily-computed individual benchmark $\pm\epsilon$, where ϵ is also sub-constant.
   We show that both $\sigma, \epsilon = \widetilde{\Theta}(n^{-1/3})$.[1] As we will see, this implies, w.h.p., that for all the agents other than the bottommost σ fraction, each agent has $\Theta(\ln n)$ possible edges (proposals) that could be in any stable matching, namely the proposals that provide both agents utility within ϵ of their benchmark. Furthermore, we show our bound is tight: with fairly high probability, there is no matching, let alone stable matching, providing every agent a partner if the values of ϵ and σ are reduced by a suitable constant factor.
   An interesting consequence of this lower bound on the agents' utilities is that the agents can readily identify a moderate sized subset of the edge set to which they can safely restrict their applications. More precisely, any woman w outside the bottommost σ fraction, knowing only her own public rating, the public ratings of the men, and her own private score for each man, can determine a preference list of length $\widetilde{\Theta}(n^{1/3})$ which, w.h.p, will yield the same result as her true full-length list. Our analysis also shows that if w obtained the men's private scores for these proposals, then w.h.p. she could safely limit herself to a length $O(\ln n)$ preference list.
2. The above bounds apply not only to the linearly separable model, but to a significantly more general bounded derivative model (in which derivatives of the utility functions are bounded).
3. The result also immediately extends to settings with unequal numbers of men and women. Essentially, our analysis shows that the loss for an agent is small if there is a σ fraction of agents of lower rank on the opposite side. Thus even on the longer side, w.h.p., the topmost $n(1 - \sigma)$ agents all obtain utility close to their benchmark, where n is the size of the shorter side. This limits the "stark effect of competition" [3]—namely that the agents on the longer side are significantly worse off—to a lower portion of the agents on the longer side.
4. The result extends to the many-to-one setting, in which agents on one side seek multiple matches. Our results are given w.r.t. a parameter d, the number of matches that each agent on the "many" side desires. For simplicity, we assume this parameter is the same for all these agents. In fact, we analyze a more general many-to-many setting.
5. A weaker result with arbitrarily small $\sigma, \epsilon = \Theta(1)$ holds when there is no restriction on the derivatives of the utility functions, which we call the general values model. Again, we show this bound cannot be improved in general. This setting is essentially the general

[1] The $\widetilde{\Theta}(\cdot)$ notation means up to a poly-logarithmic term; here $\sigma, \epsilon = \Theta((n/\ln n)^{-1/3})$.

setting considered by Lee [17]. He had shown there was a σ fraction of agents who might suffer larger losses; our bound identifies this σ fraction of agents as the bottommost agents.

6. In the bounded derivative model, with slightly stronger constraints on the derivatives, we also show the existence of an ϵ-Bayes-Nash equilibrium in which no agent proposes more than $O(\ln^2 n)$ times and all but the bottommost $O((\ln n/n)^{1/3})$ fraction of the agents make only the $O(\ln n)$ proposals identified in (1) above. Here $\epsilon = \Theta(\ln n/n^{1/3})$.

These results all follow from a lemma showing that, w.h.p., each non-bottommost agent has at most a small loss. In turn, the proof of this lemma relies on a new technique which sidesteps the conditioning inherent to runs of DA in these settings.

Experimental results

Much prior work has been concerned with preference lists that have a constant bound on their length. For moderate values of n, say $n \in [10^3, 10^6]$, $\ln n$ is quite small, so our $\Theta(\ln n)$ bound may or may not be sufficiently small in practice for this range of n. What matters are the actual constants hidden by the Θ notation, which our analysis does not fully determine. To help resolve this, we conducted a variety of simulation experiments.

We have also considered how to select the agents to include in the preference lists, when seeking to maintain a constant bound on their lengths, namely a bound that, for the values of n we considered, was smaller than the $\Theta(\ln n)$ bound determined by the above simulations; again, our investigation was experimental.

Other Related work

The random preference model was introduced by Knuth [12] (for a version in English see [13]), and subsequently extensively analyzed [20, 14, 21, 18, 23, 22, 16]. In this model, each agent's preferences are an independent uniform random permutation of the agents on the other side. An important observation was that when running the DA algorithm, the proposing side obtained a match of rank $\Theta(\ln n)$ on the average, while on the other side the matches had rank $\Theta(n/\ln n)$.

A recent and unexpected observation in [3] was the "stark effect of competition": that in the random preferences model the short side, whether it was the proposing side or not, was the one to enjoy the $\Theta(\ln n)$ rank matches. Subsequent work showed that this effect disappeared with short preference lists in a natural modification of the random preferences model [11]. Our work suggests yet another explanation for why this effect may not be present: it does not require that short preference lists be imposed as an external constraint, but rather that the preference model generates few edges that might ever be in a stable matching.

The number of edges present in any stable matching has also been examined for a variety of settings. When preference lists are uniform the expected number of stable pairs is $\Theta(n \ln n)$ [21]; when they are arbitrary on one side and uniform on the other side, the expected number is $O(n \ln n)$ [14]. This result continues to hold when preference lists are arbitrary on the men's side and are generated from general popularities on the women's side [6]. Our analysis shows that in the linear separable model (and more generally in the bounded derivative setting) the expected number of stable pairs is also $O(n \ln n)$.

Another important issue is the amount of communication needed to identify who to place on one's preference lists when they have bounded length. In general, the cost is $\Omega(n)$ per agent (in an n agent market) [8], but in the already-mentioned separable model of Ashlagi et al. [2] this improves to $\widetilde{O}(\sqrt{n})$ given some additional constraints, and further improves to

$O(\ln^4 n)$ in a tiered separable market [2]. We note that for the bounded derivatives setting, with high probability, the communication cost will be $O(n^{1/3} \ln^{2/3} n)$ for all agents except the bottommost $\Theta(n^{2/3} \ln^{1/3} n)$, for whom the cost can reach $O(n^{2/3} \ln^{1/3} n)$.

Another approach to selecting which universities to apply to was considered by Shorrer who devised a dynamic program to compute the optimal choices for students assuming universities had a common ranking of students [26].

Roadmap

In Section 2 we review some standard material. In Section 3 we state our main result in two parts: Theorem 1, which bounds the losses in the setting of the linear model, and Theorem 2, which shows it suffices to limit preference lists to a small set of edges. We prove these theorems in Sections 4 and 5, respectively. We also present some numerical simulations for the linear separable model in Section 6 We conclude with a brief discussion of open problems in Section 7.

In the appendices of the full version of the paper, we formally state and prove all the other results alluded to in the introduction and we also present further numerical simulations for the linear separable model. For the reader's convenience, in the text that follows, we provide pointers to these appendices, as appropriate. We note that Appendix A provides a complete summary of the content in these appendices.

2 Preliminaries

2.1 Stable Matching and the Deferred Acceptance (DA) Algorithm

Let M be a set of n men and W a set of n women. Each man m has an ordered list of women that represents his preferences, i.e. if a woman w comes before a woman w' in m's list, then m would prefer matching with w rather than w'. The position of a woman w in this list is called m's ranking of w. Similarly each woman w has a ranking of her preferred men[2]. The stable matching task is to pair (match) the men and women in such a way that no two people prefer each other to their assigned partners. More formally:

▶ **Definition 1** (Matching). *A matching is a pairing of the agents in M with the agents in W. It comprises a bijective function μ from M to W, and its inverse $\nu = \mu^{-1}$, which is a bijective function from W to M.*

▶ **Definition 2** (Blocking pair). *A matching μ has a blocking pair (m, w) if and only if:*

1. *m and w are not matched: $\mu(m) \neq w$.*
2. *m prefers w to his current match $\mu(m)$.*
3. *w prefers m to her current match $\nu(w)$.*

▶ **Definition 3** (Stable matching). *A matching μ is stable if it has no blocking pair.*

Gale and Shapley [5] proposed the seminal deferred acceptance (DA) algorithm for the stable matching problem. We present the woman-proposing DA algorithm (Algorithm 1); the man-proposing DA is symmetric. The following facts about the DA algorithm are well known. We state them here without proof and we shall use them freely in our analysis.

[2] Throughout this paper, we assume that each man m (woman w) ranks all the possible women (men), i.e. m's (w's) preference list is complete.

Algorithm 1 Woman Proposing Deferred Acceptance (DA) Algorithm.

```
Initially, all the men and women are unmatched.
while some woman w with a non-empty preference list is unmatched do
    let m be the first man on her preference list;
    if m is currently unmatched then
        tentatively match w to m.
    end
    if m is currently matched to w', and m prefers w to w' then
        make w' unmatched and tentatively match w to m.
    else
        remove m from w's preference list.
    end
end
```

▶ **Observation 1.**

1. *DA terminates and outputs a stable matching.*
2. *The stable matching generated by DA is independent of the order in which the unmatched agents on the proposing side are processed.*
3. *Woman-proposing DA is woman-optimal, i.e. each woman is matched with the best partner she could be matched with in any stable matching.*
4. *Woman-proposing DA is man-pessimal, i.e. each man is matched with the worst partner he could be matched with in any stable matching.*

2.2 Useful notation and definitions

There are n men and n women. In all of our models, each man m has a utility $U_{m,w}$ for the woman w, and each woman w has a utility $V_{m,w}$ for the man m. These utilities are defined as

$$U_{m,w} = U(r_w, s_m(w)), \text{ and}$$
$$V_{m,w} = V(r_m, s_w(m)),$$

where r_m and r_w are common public ratings, $s_m(w)$ and $s_w(m)$ are private scores specific to the pair (m,w), and $U(\cdot,\cdot)$ and $V(\cdot,\cdot)$ are continuous and strictly increasing functions from $\mathbb{R}_+^2$ to $\mathbb{R}_+$. The ratings are independent uniform draws from $[0,1]$ as are the scores.

In the *Linear Separable Model*, each man m assigns each woman w a utility of $U_{m,w} = \lambda \cdot r_w + (1-\lambda)\cdot s_m(w)$, where $0 < \lambda < 1$ is a constant. The women's utilities for the men are defined analogously as $V_{m,w} = \lambda \cdot r_m + (1-\lambda)\cdot s_w(m)$. All our experiments are for this model.

We let $\{m_1, m_2, \ldots, m_n\}$ be the men in descending order of their public ratings and $\{w_1, w_2, \ldots, w_n\}$ be a similar ordering of the women. We say that m_i has public rank i, or rank i for short, and similarly for w_i. We also say that m_i and w_i are *aligned*. In addition, we often want to identify the men or women in an interval of public ratings. Accordingly, we define $M(r,r')$ to be the set of men with public ratings in the range (r,r'), and $M[r,r']$ to be the set with public ratings in the range $[r,r']$; we also use the notation $M(r,r']$ and $M[r,r')$ to identify the men with ratings in the corresponding semi-open intervals. We use an analogous notation, with W replacing M, to refer to the corresponding sets of women.

We will be comparing the achieved utilities in stable matchings to the following benchmarks: the rank i man has as benchmark $U(r_{w_i}, 1)$, the utility he would obtain from the combination of the rank i woman's public rating and the highest possible private score; and similarly for the women. Based on this we define the loss an agent faces as follows.

▶ **Definition 4** (Loss). *Suppose man m and woman w both have rank i. The* loss *m sustains from a match of utility u is defined to be $U(r_w, 1) - u$. The loss for women is defined analogously.*

In our analysis we will consider a complete bipartite graph whose two sets of vertices correspond to the men and women, respectively. For each man m and woman w, we view the possible matched pair (m, w) as an edge in this graph. Thus, throughout this work, we will often refer to edges being proposed, as well as edges satisfying various conditions.

3 Upper Bound in The Linear Separable Model

To illustrate our proof technique for deriving upper bounds, we begin by stating and proving our upper bound result for the special case of the linear separable model with $\lambda = \frac{1}{2}$.

▶ **Theorem 1.** *In the linear separable model with $\lambda = 1/2$, when there are n men and n women, for any given constant $c > 0$, for large enough n, with probability at least $1 - n^{-c}$, in every stable matching, for every i, with $r_{w_i} \geq \overline{\sigma} \triangleq 3\overline{L}/2$, agent m_i suffers a loss of at most $\overline{L}$, where $\overline{L} = (16(c+2)\ln n/n)^{1/3}$, and similarly for the agents w_i.*

In words, w.h.p., all but the bottommost agents (those whose aligned agents have public rating less than $\overline{\sigma}$) suffer a loss of no more than $\overline{L}$. This is a special case of our basic upper bound for the bounded utilities model (Theorem 12).

One of our goals is to be able to limit the number of proposals the proposing side needs to make. We identify the edges that could be in some stable matching, calling them acceptable edges. Our definition is stated generally so that it covers all our results; accordingly we replace the terms $\overline{L}$ and $\overline{\sigma}$ in Theorem 1 with parameters L and σ.

▶ **Definition 5** (Acceptable edges). *Let $0 < \sigma < 1$ and $0 < L < 1$ be two parameters. An edge (m_i, w_j) is* (L, σ)-man-acceptable *either if it provides m_i utility at least $U(r_{w_i}, 1) - L$, or if $m_i \in M[0, \sigma)$. The definition of (L, σ)-woman-acceptable is symmetric. Finally, (m_i, w_j) is* (L, σ)-acceptable *if it is both (L, σ)-man and (L, σ)-woman-acceptable.*

To prove our various results, we choose L and σ so that w.h.p. the edges in every stable matching are (L, σ)-*acceptable*. We call this high probability event $\mathcal{E}$. We will show that if $\mathcal{E}$ occurs, then running DA on the set of acceptable edges, or any superset of the acceptable edges obtained via loss thresholds, produces the same stable matching as running DA on the full set of edges.

▶ **Theorem 2.** *If $\mathcal{E}$ occurs, then running woman-proposing DA with the edge set restricted to the acceptable edges or to any superset of the acceptable edges obtained via loss thresholds (including the full edge set) result in the same stable matching.*

The implication is that w.h.p. a woman can safely restrict her proposals to her acceptable edges, or to any overestimate of this set of edges obtained by her setting an upper bound on the loss she is willing to accept. There is a small probability— at most n^{-c}—that this may result in a less good outcome, which can happen only if $\mathcal{E}$ does not occur. Note that Theorem 2 applies to every utility model we consider. Then, w.h.p., every stable matching gives each woman w, whose aligned agent m has public rating $r_m \geq \overline{\sigma} = \Omega((\ln n/n)^{1/3})$, a partner with public rating in the range $[r_m - 2\overline{L}, r_m + \frac{5}{2}\overline{L}]$ (see Theorem 25 in Appendix F.1).

The bound $r_m - 2\overline{L}$ is a consequence of the bound on the woman's loss; the bound $r_m + \frac{5}{2}\overline{L}$ is a consequence of the bound on the men's losses. An analogous statement applies to the men.

This means that if we are running woman-proposing DA, each of these women might as well limit her proposals to her woman-acceptable edges, which is at most the men with public ratings in the range $r_m \pm \Theta(\overline{L})$ for whom she has private scores of at least $1 - \Theta(\overline{L})$. In expectation, this yields $\Theta(n^{1/3}(\ln n)^{2/3})$ men to whom it might be worth proposing. It also implies that a woman can have a gain of at most $\Theta(\overline{L})$ compared to her target utility.

If, in addition, each man can inexpensively signal the women who are man-acceptable to him, then the women can further limit their proposals to just those men providing them with a signal; this reduces the expected number of proposals these women can usefully make to just $\Theta(\ln n)$.

4 Sketch of the Proof of Theorem 1

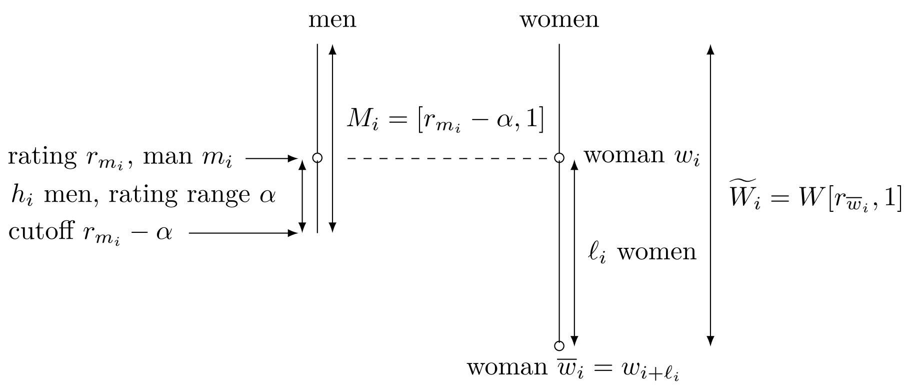

Figure 1

We begin by outlining the main ideas used in our analysis. Our goal is to show that when we run woman proposing DA, w.h.p. each man receives a proposal that gives him a loss of at most L (except possibly for men among the bottommost $\Theta(nL)$). As the outcome is the man-pessimal stable matching, this means that w.h.p., in all stable matchings, these men have a loss of at most L. By symmetry, the same bound holds for the women.

Next, we provide some intuition for the proof of this result. See Fig. 1. Our analysis uses 3 parameters $\alpha, \beta, \gamma = \Theta(L)$. Let m_i be a non-bottommost man. We consider the set of men with public rank at least $r_{m_i} - \alpha$: $M_i = M[r_{m_i} - \alpha, 1]$. We consider a similar, slightly larger set of women: $\widetilde{W}_i = W[r_{w_i} - 3\alpha, 1]$. Now we look at the best proposals by the women in $\widetilde{W}_i$, i.e. the ones they make first. Specifically, we consider the proposals that give these women utility at least $V(r_{m_i} - \alpha, 1)$, proposals that are therefore guaranteed to be to the men in M_i. Let $|M_i| = i + h_i$ and $|W_i| = i + \ell_i$. In expectation, $\ell_i - h_i = 2\alpha n$. Necessarily, at least $\ell_i - h_i + 1$ women in M_i cannot match with men in $M_i \setminus \{m_i\}$. But, as we will see, these women all have probability at least β of having a proposal to m_i which gives them utility at least $V(r_{m_i} - \alpha, 1)$. These are proposals these women must make before they make any proposals to men with public rating less than $r_{m_i} - \alpha$. Furthermore, for each of these proposals, m_i has probability at least γ of having a loss of L or less. Thus, in expectation, m_i receives at least $2\alpha\beta\gamma n$ proposals which give him a loss of L or less.

We actually want a high-probability bound. So we choose α, β, γ so that $\alpha\beta\gamma n \geq c \log n$ for a suitable constant $c > 0$, and then apply a series of Chernoff bounds. There is one difficulty. The Chernoff bounds requires the various proposals to be independent. Unfortunately, in general, this does not appear to be the case. However, we are able to show that the failure probability for our setting is at most the failure probability in an artificial setting in which the events are independent, which yields the desired bound.

We now embark on the actual proof.

We formalize the men's rating cutoff with the notion of DA stopping at public rating r.

▶ **Definition 6** (DA stops). *The women* stop at public rating r *if, in each woman's preference list, all the edges with utility less than $V(r, 1)$ are removed. The women* stop at man m *if, in each woman's preference list, all the edges following her edge to m are removed. The women* double cut at man m and public rating r, *if they each stop at m or r, whichever comes first. Men stopping and double cutting are defined similarly. Finally, an edge is said to* survive *the cutoff if it is not removed by the stopping.*

To obtain our bounds for man m_i, we will have the women double cut at rating $r_{m_i} - \alpha$ and at man m_i, where $\alpha > 0$ is a parameter we will specify later.

Our upper bounds in all of the utility models depend on a parameterized key lemma (Lemma 3) stated shortly. This lemma concerns the losses the men face in the woman-proposing DA; a symmetric result applies to the women. The individual theorems follow by setting the parameters appropriately. Our key lemma uses three parameters: $\alpha, \beta, \gamma > 0$. To avoid rounding issues, we will choose α so that αn is an integer. The other parameters need to satisfy the following constraints.

$$\text{for } r \geq \alpha: \qquad V(r - \alpha, 1) \leq V(r, 1 - \beta) \tag{1}$$

$$\text{for } r \geq 3\alpha: \qquad U(r, 1) - U(r - 3\alpha, 1 - \gamma) \leq L \tag{2}$$

Equation (1) relates the range of private values that will yield a woman an edge to m_i that survives the cut at $r_{m_i} - \alpha$, or equivalently the probability of having such an edge. Observation 2 below, shows that Equation (2) identifies the range of m_i's private values for proposals from $\widetilde{W}_i$ that yield him a loss of at most L (for we will ensure the women in $\widetilde{W}_i$ have public rating at least $r_{w_i} - 3\alpha$).

▶ **Observation 2.** *Consider the proposal from woman w to the rank i man m_i. Suppose the rank i woman w_i has rating $r_{w_i} \geq 3\alpha$. If w has public rating $r \geq r_{w_i} - 3\alpha$ and m_i's private score for w is at least $1-\gamma$, then m_i's utility for w is at least $U(r_{w_i} - 3\alpha, 1-\gamma) \geq U(r_{w_i}, 1) - L$.*

In the linear separable model with $\lambda = \frac{1}{2}$, we set $\alpha = \beta = \gamma$ and $L = 2\alpha$.

The next lemma determines the probability that man m_i receives a proposal causing him a loss of at most L. The lemma calculates this probability in terms of the parameters we just defined. Note that the result does not depend on the utility functions $U(\cdot, \cdot)$ and $V(\cdot, \cdot)$ being linear. In fact, the same lemma applies to much more general utility models which we also study (see Appendix C) and it is the crucial tool we use in all our upper bound proofs.

In what follows, to avoid heavy-handed notation, by $r_{m_i} - \alpha$ we will mean $\max\{0, r_{m_i} - \alpha\}$.

In order to state our next result crisply, we define the following Event $\mathcal{E}_i$. It concerns a run of woman-proposing DA with double cut at the rank i man m_i and at public rating $r_{m_i} - \alpha$. Let $h_i = \left|M[r_{m_i} - \alpha, r_{m_i})\right|$, $\ell_i = \left|W[r_{w_i} - 3\alpha, r_{w_i})\right|$, and $\overline{w}_i$ be the woman with rank $i + \ell_i$. See Figure 1 for an illustration of these definitions. Event $\mathcal{E}_i$ occurs if $r_{w_i} \geq 3\alpha$ and between them the $i + \ell_i$ women in $W[r_{w_i} - 3\alpha, 1]$ make at least one proposal to m_i that causes him a loss of at most L.

Finally we define Event $\mathcal{E}$: it happens if $\mathcal{E}_i$ occurs for all i such that $r_{w_i} \geq 3\alpha$.

▶ **Lemma 3.** *Let $\alpha > 0$ and $L > 0$ be given, and suppose that β and γ satisfy* (1) *and* (2)*, respectively. Then, Event $\mathcal{E}$ occurs with probability at least $1 - p_f$, where the failure probability*

$$p_f = n \cdot \exp(-\alpha(n-1)/12) + n \cdot \exp(-\alpha(n-1)/24) + n \exp(-\alpha\beta n/8) + n \cdot \exp(-\alpha\beta\gamma n/2).$$

The following simple claim notes that the men's loss when running the full DA is no larger than when running double-cut DA.

▷ Claim 4. Suppose a woman-proposing double-cut DA at man m_i and rating $r_{m_i} - \alpha$ is run, and suppose m_i incurs a loss of L. Then in the full run of woman-proposing DA, m_i will incur a loss of at most L.

Proof. Recall that when running the women-proposing DA the order in which unmatched women are processed does not affect the outcome. Also note that as the run proceeds, whenever a man's match is updated, the man obtains an improved utility. Thus, in the run with the full edge set we can first use the edges used in the double-cut DA and then proceed with the remaining edges. Therefore if in the double-cut DA m_i has a loss of L, in the full run m_i will also have a loss of at most L. ◀

To illustrate how this lemma is applied, we now prove Theorem 1. Note that $\overline{L}$ is the value of L used in this theorem. Our other results use other values of L.

Proof. (Of Theorem 1) By Lemma 3, in the double-cut DA, for all i with $r_{w_i} \geq 3\alpha$, m_i obtains a match giving him loss at most $\overline{L}$, with probability at least $1 - n \cdot \exp(-\alpha(n-1)/12) - n \cdot \exp(-\alpha n/24) - n \exp\big(-\alpha^2 n/8\big) - n \cdot \exp\big(-\alpha^3 n/2\big)$.

By Claim 4, m_i will incur a loss of at most $\overline{L}$ in the full run of woman-proposing DA with at least as large a probability. But this is the man-pessimal match. Consequently, in every stable match, m_i has a loss of at most $\overline{L}$. By symmetry, the same bound applies to each woman w_i such that $r_{m_i} \geq 3\alpha$.

We choose $\overline{L} = [16(c+2)\ln n/n]^{1/3}$. Recalling that $\alpha = \overline{L}/2$, we see that for large enough n the probability bound, over all the men and women, is at most $1 - n^{-c}$. The bounds $r_{w_i} \geq 3\alpha$ and $r_{m_i} \geq 3\alpha$ imply we can set $\overline{\sigma} = 3\alpha = \frac{3}{2}\overline{L}$. ◀

Proof. (Of Lemma 3.) We run the double-cut DA in two phases, defined as follows. Recall that $h_i = \big|M[r_{m_i} - \alpha, r_{m_i})\big|$ and $\ell_i = \big|W[r_{w_i} - 3\alpha, r_{w_i})\big|$. Note that women with rank at most $i + \ell_i$ have public rating at least $r_{w_i} - 3\alpha$.
Phase 1. Every unmatched woman with rank at most $i + \ell_i$ keeps proposing until her next proposal is to m_i, or she runs out of proposals.
Phase 2. Each unmatched women makes her next proposal, if any, which will be a proposal to m_i.

Our analysis is based on the following four claims. The first two are simply observations that w.h.p. the number of agents with public ratings in a given interval is close to the expected number. We defer the proofs to the appendix.

A critical issue in this analysis is to make sure the conditioning induced by the successive steps of the analysis does not affect the independence needed for subsequent steps. To achieve this, we use the Principle of Deferred Decisions, only instantiating random values as they are used. Since each successive bound uses a different collection of random variables this does not present a problem.

▷ Claim 5. Let $\mathcal{B}_1$ be the event that for some i, $h_i \geq \frac{3}{2}\alpha(n-1)$. $\mathcal{B}_1$ occurs with probability at most $n \cdot \exp(-\alpha(n-1)/12)$. The only randomness used in the proof are the choices of the men's public ratings. The same bound applies to the women.

Proof. (Sketch.) As $\mathrm{E}[h_i] = \alpha(n-1)$, w.h.p., $h_i < \frac{3}{2}\alpha(n-1)$. This claim uses a Chernoff bound with the randomness coming from the public ratings of the men. ◀

▷ Claim 6. Let $\mathcal{B}_2$ be the event that for some i, $\ell_i \le \frac{5}{2}\alpha(n-1)$. $\mathcal{B}_2$ occurs with probability at most $n \cdot \exp(-\alpha(n-1)/24)$. The only randomness used in the proof are the choices of the women's public ratings. The same bound applies to the men.

Proof. This is very similar to the proof of Claim 5. ◀

▷ Claim 7. Let $\mathcal{B}_3$ be the event that between them, the women with rank at most $i + \ell_i$ make fewer than $\frac{1}{2}\alpha\beta n$ Step 2 proposals to m_i. If events $\mathcal{B}_1$ and $\mathcal{B}_2$ do not occur, then $\mathcal{B}_3$ occurs with probability at most $\exp(-\alpha\beta n/8)$. The only randomness used in the proof are the choices of the women's private scores.

This bound uses the private scores of the women and employs a novel argument given below to sidestep the conditioning among these proposals.

▷ Claim 8. If none of the events $\mathcal{B}_1$, $\mathcal{B}_2$, or $\mathcal{B}_3$ occur, then at least one of the Step 2 proposals to m_i will cause him a loss of at most L with probability at least $1-(1-\gamma)^{\alpha\beta n/2} \ge 1 - \exp(-\alpha\beta\gamma n/2)$. The only randomness used in the proof are the choices of the men's private scores.

Proof. Note that each Phase 2 proposal is from a woman w with rank at most $i + \ell_i$. As already observed, her public rating is at least $r_{w_i} - 3\alpha$. Recall that man m_i's utility for w equals $U(r_w, s_{m_i}(w)) \ge U(r_{w_i} - 3\alpha, s_{m_i}(w))$. To achieve utility at least $U(r_{w_i}, 1) - L \le U(r_{w_i} - 3\alpha, 1-\gamma)$ (using (2)) it suffices to have $s_{m_i}(w) \ge 1-\gamma$, which happens with probability γ. Consequently, utility at least $U(r_{w_i}, 1) - L$ is achieved with probability at least γ.

For each Phase 2 proposal these probabilities are independent as they reflect m_i's private scores for each of these proposals. Therefore the probability that there is no proposal providing m_i a loss of at most L is at most

$$(1-\gamma)^{\alpha\beta n/2} \le \exp(\alpha\beta\gamma n/2).$$

◀

Concluding the proof of Lemma 3: The overall failure probability summed over all n choices of i is

$$n \cdot \exp(-\alpha(n-1)/12) + n \cdot \exp(-\alpha(n-1)/24) + n\exp(-\alpha\beta n/8) + n \cdot \exp(-\alpha\beta\gamma n/2).$$

◀

Proof. (Of Claim 7.) First, we simplify the action space by viewing the decisions as being made on a discrete utility space, as specified in the next claim, proved in the appendix.

▷ Claim 9. For any $\delta > 0$, there is a discrete utility space in which for each woman the probability of selecting m_i is only increased, and the probability of having any differences in the sequence of actions in the original continuous setting and the discrete setting is at most δ.

We represent the possible computations of the double-cut DA in this discrete setting using a tree T. Each woman will be going through her possible utility values in decreasing order, with the possible actions of the various women being interleaved in the order given by the DA processing. Each node u corresponds to a woman w processing her next utility value. The possible choices at this utility are each represented by an edge descending from u. These choices are:

i. Proposing to some man (among those men w has not yet proposed to); or

ii. "no action". This corresponds to w making no proposal achieving the current utility.

We observe the following important structural feature of tree T. Let S be the subtree descending from the edge corresponding to woman w proposing to m_i; in S there are no further actions of w, i.e. no nodes at which w makes a choice, because the double cut DA cuts at the proposal to m_i.

The assumption that $\mathcal{B}_1$ and $\mathcal{B}_2$ do not occur means that for all i, $h_i < \frac{3}{2}\alpha(n-1)$ and $\ell_i > \frac{5}{2}\alpha(n-1)$, and therefore $\ell_i - h_i > \alpha(n-1)$.

At each leaf of T, up to $i + h_i - 1$ women will have been matched with someone other than m_i. The other women either finished with a proposal to m_i or both failed to match and did not propose to m_i. Let w be a woman in the latter category. Then, on the path to this leaf, w will have traversed edges corresponding to a choice at each discrete utility in the range $[V(r_{m_i} - \alpha, 1), V(1, 1)]$.

We now create an extended tree, T_x, by adding a subtree at each leaf; this subtree will correspond to pretending there were no matches; the effect is that each women will take an action at all their remaining utility values in the range $[V(r_{m_i} - \alpha, 1), V(1, 1)]$, except that in the sub-subtrees descending from edges that correspond to some woman w selecting m_i, w has no further actions. For each leaf in the unextended tree, the probability of the path to that leaf is left unchanged. The probabilities of the paths in the extended tree are then calculated by multiplying the path probability in the unextended tree with the probabilities of each woman's choices in the extended portion of the tree.

Next, we create an artificial mechanism $\mathcal{M}$ that acts on tree T_x. The mechanism $\mathcal{M}$ is allowed to put $i + h_i - 1$ "blocks" on each path; blocks can be placed at internal nodes. A block names a woman w and corresponds to her matching (but we no longer think of the matches as corresponding to the outcome of the edge selection; they have no meaning beyond making all subsequent choices by this woman be the "no action" choice).

DA can be seen as choosing to place up to $i + h_i - 1$ blocks at each of the nodes corresponding to a leaf of T. $\mathcal{M}$ will place its blocks so as to minimize the probability p of paths with at least $\frac{1}{2}\alpha\beta n$ women choosing edges to m_i. Clearly p is a lower bound on the probability that the double-cut DA makes at least $\frac{1}{2}\alpha\beta n$ proposals in Step 2. Given a choice of blocks we call the resulting probability of having fewer than $\frac{1}{2}\alpha\beta n$ women choosing edges to m_i the *blocking probability*.

▷ Claim 10. The probability that $\mathcal{M}$ makes at least $\frac{1}{2}\alpha\beta n$ proposals to m_i is at least $1 - \exp(-\alpha\beta n/8)$.

▶ **Corollary 1.** *The probability that the double-cut DA makes at least* $\frac{1}{2}\alpha\beta n$ *proposals to* m_i *is at least* $1 - \exp(-\alpha\beta n/8)$.

Proof. For any fixed δ, by Claim 10, the probability that $\mathcal{M}$ makes at least $\frac{1}{2}\alpha\beta n$ proposals to m_i is at least $1 - \exp(-\alpha\beta n/8)$. By construction, the probability is only larger for the double-cut DA in the discrete space.

Therefore, by Claim 4, the probability that the double-cut DA makes at least $\frac{1}{2}\alpha\beta n$ proposals to m_i in the actual continuous space is at least $1-\exp(-\alpha\beta n/8)-\delta$, and this holds for any $\delta > 0$, however small. Consequently, this probability is at least $1-\exp(-\alpha\beta n/8)$. ◀

Proof. (Of Claim 10.) We will show that the most effective blocking strategy is to block as many women as possible before they have made any choices. This leaves at least $(i + \ell_i) - (i - 1 + h_i) \geq 1 + \alpha(n-1) \geq \alpha n$ women unmatched. Then, as we argue next, each of these remaining at least αn women w has independent probability at least β that their proposal to m_i is cutoff-surviving. To be cutoff-surviving, it suffices that $V(r_{m_i}, s_w(m_i)) \geq V(r_{m_i} - \alpha, 1)$. But we know by (1) that $V(r_{m_i} - \alpha, 1) \leq V(r_{m_i}, 1-\beta)$, and therefore it suffices that $s_w(m_i) \geq 1-\beta$, which occurs with probability β.

Consequently, in expectation, there are at least $\alpha\beta n$ proposals to m_i, and therefore, by a Chernoff bound, at least $\frac{1}{2}\alpha\beta n$ proposals with probability at least $\exp(-\alpha\beta n/8)$.

We consider the actual blocking choices made by $\mathcal{M}$ and modify them bottom-up in a way that only reduces the probability of there being $\frac{1}{2}\alpha\beta n$ or more proposals to m_i.

Clearly, $\mathcal{M}$ can choose to block the same maximum number of women on every path as it never hurts to block more women (we allow the blocking of women who have already proposed to m_i even though it does not affect the number of proposals to m_i).

Consider a deepest block at some node u in the tree, and suppose b women are blocked at u. Let v be a sibling of u. As this is a deepest block, there will be no blocks at proper descendants of u, and furthermore as there are the same number of blocks on every path, v will also have b blocked women.

Observe that if there is no blocking in a subtree, then the probability that a woman makes a proposal to m_i is independent of the outcomes for the other women. Therefore the correct blocking decision at node u is to block the b women with the highest probabilities of otherwise making a proposal to m_i, which we call their *proposing probabilities*; the same is true at each of its siblings v.

Let x be u's parent. Suppose the action at node x concerns woman $\widetilde{w}_x$. Note that the proposing probability for any woman $w \neq \widetilde{w}_x$ is the same at u and v because the remaining sequence of actions for woman w is the same at nodes u and v, and as they are independent of the actions of the other women, they yield the same probability of selecting m_i at some point.

We need to consider a number of cases.

Case 1. w is blocked at every child of x.
Then we could equally well block w at node x.

Case 2. At least one woman other than $\widetilde{w}_x$ is blocked at some child of x.
Each such blocked woman w has the same proposing probability at each child of x. Therefore by choosing to block the women with the highest proposing probabilities, we can ensure that at each node either $\widetilde{w}_x$ plus the same $b-1$ other women are blocked, or these $b-1$ woman plus the same additional woman $w' \neq \widetilde{w}_x$ are blocked. In any event, the blocking of the first $b-1$ women can be moved to x.

Case 2.1. $\widetilde{w}_x$ is not blocked at any child of x.
Then the remaining identical blocked woman at each child of x can be moved to x.

Case 2.2. $\widetilde{w}_x$ is blocked at some child of x but not at all the children of x.
Notice that we can avoid blocking $\widetilde{w}_x$ at the child u of x corresponding to selecting m_i, as the proposing probability for $\widetilde{w}_x$ after it has selected m_i is 0, so blocking any other women would be at least as good. Suppose that $w \neq \widetilde{w}_x$ is blocked at node u.

Let v be another child of x at which $\widetilde{w}_x$ is blocked. Necessarily, $p_{v,\widetilde{w}_x}$, the proposing probability for $\widetilde{w}_x$ at node v, is at least the proposing probability $p_{v,w}$ for w at node v (for otherwise w would be blocked at node v); also, $p_{v,w}$ equals the proposing probability for w at every child of x including u; in addition, $p_{v,\widetilde{w}_x}$ equals the proposing probability for $\widetilde{w}_x$ at every child of x other than u. It follows that w is blocked at u and $\widetilde{w}_x$ can be blocked at every other child of x. But then blocking $\widetilde{w}_x$ at x only reduces the proposing probability.

Thus in every case one should move the bottommost blocking decisions at a collection of sibling nodes to a single blocking decision at their parent. ◀

◀

5 Making Fewer Proposals

We identify a sufficient set of edges that contains all stable matchings, and on which the DA algorithm produces the same outcome as when it runs on the full edge set.

▶ **Definition 7** (Viable edges). *An edge (m, w) is* man-viable *if, according to m's preferences, w is at least as good as the woman he is matched to in the man-pessimal stable match. Woman-viable is defined symmetrically. An edge is* viable *if it is both man and woman-viable. E_v is the set of all viable edges.*

▶ **Lemma 11.** *Running woman-proposing DA with the edge set restricted to E_v and with any superset obtained via loss thresholds, including the full edge set, results in the same stable matching.*

Proof. Suppose a new stable matching, S, now exists in the restricted edge set: it could not be present when using the full edge set, therefore there must be a blocking edge (m, w) in the full edge set. But neither m nor w would have removed this edge when forming their restricted edge set since for both of them it is better than an edge they did not remove (the edge they are matched with in S).

It follows that w.h.p. the set of stable matchings is the same when using E_v (or any super set of it generated by truncation with larger loss thresholds) and the whole set. Thus woman-proposing DA run on the restricted edge set will yield the same stable matching as on the full edge set.

◀

Proof. (Of Theorem 2.) If $\mathcal{E}$ occurs, the set of acceptable edges contains all the viable edges. Furthermore, the acceptable edges are defined by means of loss thresholds. The result now follows from Lemma 11. ◀

For some of the very bottommost agents, almost all edges may be acceptable. However, in the bounded derivatives model, with slightly stronger constraints on the derivatives, we also show (see Appendix H) the existence of an ϵ-Bayes-Nash equilibrium in which all but a bottom $\Theta((\ln n/n)^{1/3})$ fraction of agents use only $\Theta(\ln n)$ edges, and all agents propose using at most $\Theta(\ln^2 n)$ edges, with $\epsilon = O(\ln n/n^{1/3})$.

6 Numerical Simulations

We present several simulation results which are complementary to our theoretical results. Throughout this section, we focus on the linear separable model.

6.1 NRMP Data

We used NRMP data to motivate some of our choices of parameters for our simulations. The NRMP provides extensive summary data [19]. We begin by discussing this data.

Over time, the number of positions and applicants has been growing. We mention some numbers for 2021. There were over 38,000 positions available and a little over 42,000 applicants. The main match using the DA algorithm (modified to allow for couples, who comprise a little over 5% of the applicants) filled about 95% of the available positions. The NRMP also ran an aftermarket, called SOAP, after which about 0.5% of the positions remained unfilled.

The positions cover many different specialities. These specialities vary hugely in the number of positions available, with the top 11, all of size at least 1,000, accounting for 75% of the positions. In addition, about 75% of the doctors apply to only one speciality. We think that as a first approximation, w.r.t. the model we are using, it is reasonable to view each speciality as a separate market. Accordingly, we have focused our simulations on markets with 1,000–2,000 positions (though the largest speciality in the NRMP data had over 9,000 positions).

On average, doctors listed 12.5 programs in their preference lists, hospital programs listed 88 doctors, and the average program size was 6.5 (all numbers are approximate). While there is no detailed breakdown of the first two numbers, it is clear they vary considerably over the individual doctors and hospitals. For our many-to-one simulations we chose to use a fixed size for the hospital programs. Our simulations cause the other two numbers to vary over the individual doctors and programs because the public ratings and private scores are chosen by a random process.

6.2 Numbers of Available Edges

The first question we want to answer is how long do the preference lists need to be in order to have a high probability of including all acceptable edges, for all but the bottommost agents?

We chose bottommost to mean the bottom 20% of the agents, based on where the needed length of the preference lists started to increase in our experiments for $n = 1{,}000$–$2{,}000$.

We ran experiments with $\lambda = 0.5, 0.67, 0.8$, corresponding to the public rating having respectively equal, twice, and four times the weight of the private scores in their contribution to the utility. We report the results for $\lambda = 0.8$. The edge sets were larger for smaller values of λ, but the results were qualitatively the same. We generated 100 random markets and determined the smallest value of L that ensured all agents were matched in all 100 markets. $L = 0.12$ sufficed. In Figure 2, we show results by decile of women's rank (top 10%, second 10%, etc.), specifically the average length of the preference list and the average number of edges proposed by a woman in woman-proposing DA, over these 100 randomly generated markets. We also show the max and min values over the 100 runs; these can be quite far from the average value. Note that the min values in Figure 2(a) are close to the max values in Figure 2(b), which suggests that being on the proposing side does not significantly reduce the value of L that the women could use compared to the value the men use. We also show data for a typical single run in Figure 3.

We repeated the simulation for the many-to-one setting. In Figure 4, we show the results for 2000 workers and 250 companies, each with 8 positions. Now, on average, a typical worker (i.e. among the top 80%) has an average preference list length of 55 and makes 7 proposals.

The one-to-one results show that for non-bottommost agents, the preference lists have length 150 on the average, while women make 30 proposals on the average (these numbers

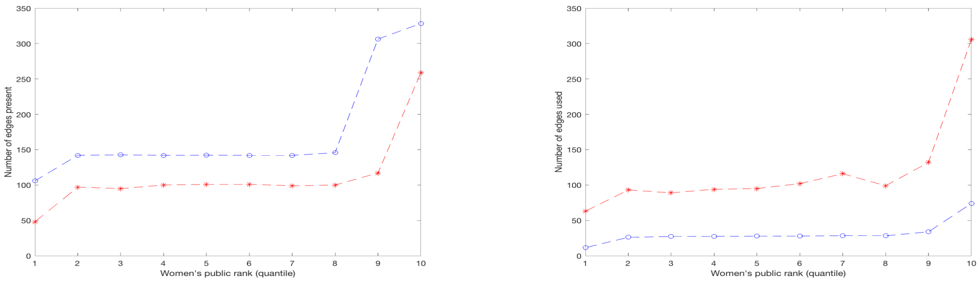

(a) Number of edges in the acceptable edge set, per woman, by decile; average in blue with circles, minimum in red with stars. ($n = 2{,}000$, $\lambda = 0.8$, $L = 0.12$.)

(b) Number of edges in the acceptable edge set proposed during the run of DA, per women, by decile; average in blue with circles, maximum in red with stars.

Figure 2 One-to-one case: summary statistics.

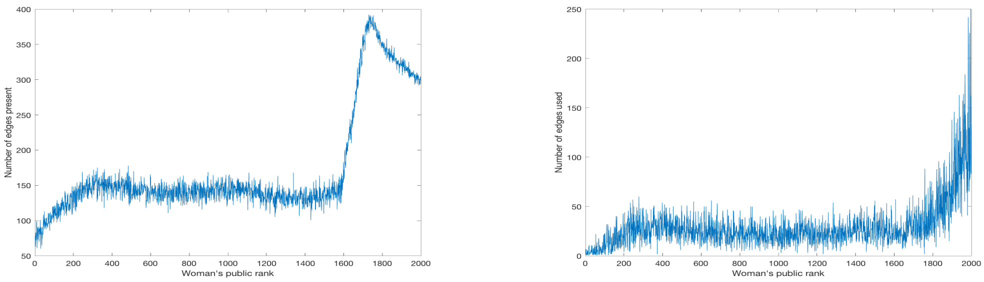

(a) Number of edges in the acceptable edge set for each woman.

(b) Number of edges in the acceptable edge set proposed by each woman.

Figure 3 One-to-one case: a typical run.

are slightly approximate). What is going on? We believe that the most common matches provide a small loss or gain ($\Theta(n^{-1/3})$ in our theoretical bounds) as opposed to the maximum loss possible ($\Theta(n^{-1/3}\ln^{1/3} n)$ in our theoretical bounds), as is indicated by our distribution bound on the losses (see item 4 in Appendix E.1). The question then is where do these edges occur in the preference list, and the answer is about one fifth of the way through (for one first has the edges providing a gain, which only go to higher up agents on the opposite side, and then one has the edges providing a loss, and these go both up and down). However, a few of the women will need to go through most of their list, as indicated by the fact that the max and min lines (for example in Figure 4) roughly coincide.

This effect can also be seen in the many-to-one experiment but it is even more stark on the worker's side. The reason is that the number of companies with whom a worker w might match which are above w, based on their public ratings alone, is $\Theta(L_c n_c)$, while the number below w is $\Theta(L_w n_c)$, a noticeably larger number. (See Appendix F.1 for a proof of these bounds.) The net effect is that there are few edges that provide w a gain, and so the low-loss edges, which are the typical matches, are reached even sooner in this setting.

Now we turn to why the number of edges in the available edge set per woman changes at the ends of the range. There are two factors at work. The first factor is due to an increasing loss bound as we move toward the bottommost women, which increases the sizes of their available edge sets. The second factor is due to public ratings. For a woman w the range of men's public ratings for its acceptable edges is $[r_m - \Theta(\overline{L}), r_m + \Theta(\overline{L})]$, where m is aligned

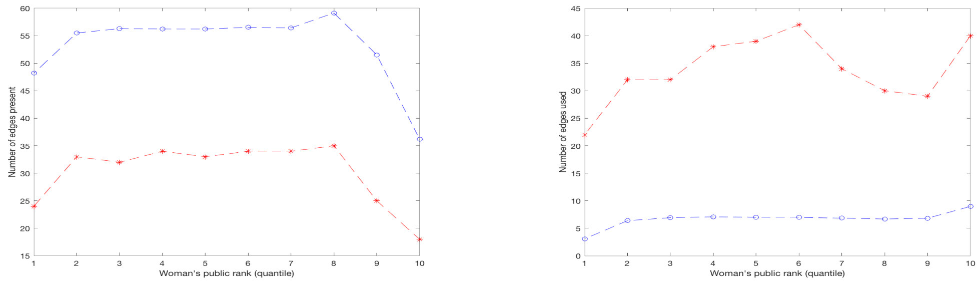

(a) Many to One Setting: Number of edges in the acceptable edge set per worker, by decile; average in blue with circles, minimum in red with stars. (n_w = 2,000, $d = 8$, $\lambda = 0.8$, $L_c = 0.14$, $L_w =$ 0.24.)

(b) Number of edges in the acceptable edge set proposed during the run of DA, per worker, by decile; average in blue with circles, maximum in red with stars.

Figure 4 Many to One Setting.

with w. But at the ends a portion of this range will be cut off, reducing the number of acceptable edges, with the effect more pronounced for low public ratings. Because $\lambda = 0.8$, initially, as we move to lower ranked women, the gain due to increasing the loss bound dominates the loss due to a reduced public rating range, but eventually this reverses. Both effects can be clearly seen in Figure 3(a), for example.

6.3 Unique Stable Partners

Another interesting aspect of our simulations is that they showed that most agents have a unique stable partner. This is similar to the situation in the popularity model when there are short preference lists, but here this result appears to hold with full length preference lists. In Figure 5, we show the outcome on a typical run and averaged over 100 runs, for $n =$ 2,000 in the one-to-one setting. We report the results for the men, but as the setting is symmetric they will be similar for the women. On the average, among the top 90% of agents by rank, 0.5% (10 of 1,800) had more than one stable partner, and among the remainder another 2% had multiple stable partners (40 of 200).

Also, as suggested by the single run illustrated in Figure 5(a), the pair around public rank 1,600 and the triple between 1,200 and 1,400 have multiple stable partners which they can swap (or exchange via a small cycle of swaps) to switch between different stable matchings. This pattern is typical for the very few men with multiple stable partners outside the bottommost region.

6.4 Constant Number of Proposals

Our many-to-one experiments suggest that the length of the preference lists needed by our model are larger than those observed in the NRMP data. In addition, even though there is a simple rule for identifying these edges, in practice the communication that would be needed to identify these edges may well be excessive. In light of this it is interesting to investigate what can be done when the agents have shorter preference lists.

We simulated a strategy where the workers' preference lists contain only a constant number of edges. We construct an *Interview Edge Set* which contains the edges (w, c) satisfying the following conditions:

1. Let r_w and r_c be the public ratings of w and c respectively. Then $|r_w - r_c| \le p$.

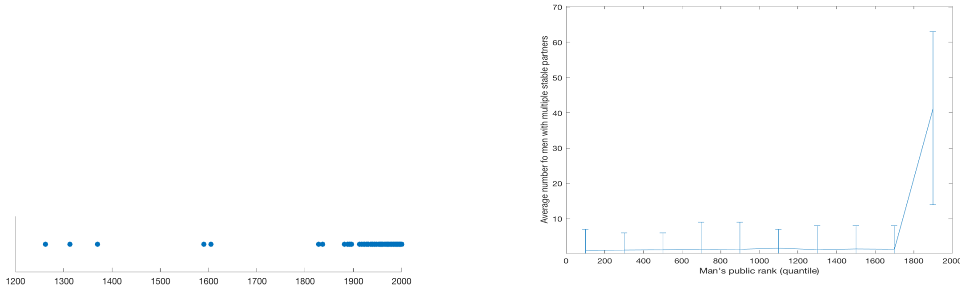

(a) Public ranks of men with multiple stable partners in a typical run.

(b) Average numbers of men with multiple stable partners, by decile.

Figure 5 Unique stable partners, one-to-one setting.

2. The private score w has for c as well as the private score of c for w are both greater than q.

We choose the parameters p and q so as to have 15 edges per agent on average. Many combinations of p and q would work. We chose a pair that caused relatively few mismatches. We then ran worker proposing DA on the Interview Edge Set.

One way of identifying these edges is with the following communication protocol: the workers signal the companies which meet their criteria (the workers' criteria); the companies then reply to those workers who meet their criteria. In practice this would be a lot of communication on the workers's side, and therefore it may be that an unbalanced protocol where the workers use a larger q_w as their private score cutoff and the companies a correspondingly smaller q_c is more plausible. Clearly this will affect the losses each side incurs when there is a match, but we think it will have no effect on the non-match probability, and as non-matches are the main source of losses, we believe our simulation is indicative. We ran the above experiment with $p = 0.19$ and $q = 0.60$, with the company capacity being 8. Figure 6(a) shows the locations of unmatched workers in a typical run of this experiment while 6(b) shows the average numbers of unmatched workers per quantile (of public ratings) over 100 runs. We observe that the number of unmatched workers is very low (about 1.5% of the workers) and most of these are at the bottom of the public rating range.

Figure 6(c) compares the utility obtained by the workers in the match obtained by running worker-proposing DA on the Interview Edge Set to the utility they obtain in the worker-optimal stable match. We observe that only a small number of workers have a significantly worse outcome when restricted to the Interview Edge Set.

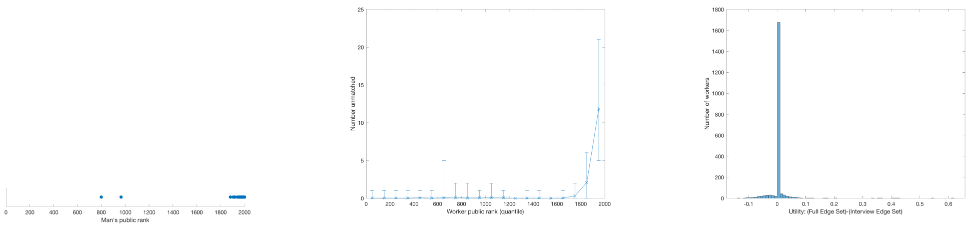

(a) Public ranks of unmatched workers in a typical run.

(b) Average numbers of unmatched workers by public rating decile.

(c) Distribution of workers' utilities with worker-proposing DA: (full edge set result) −(Interview edge set result)

Figure 6 Constant number of proposals.

7 Discussion and Open Problems

Our work shows that in the bounded derivatives model, apart from a sub-constant fraction of the agents, each of the other agents has $O(\ln n)$ easily identified edges on their preference list which cover all their stable matches w.h.p.

As described in Section 6, our experiments for the one-to-one setting yield a need for what appear to be impractically large preference lists. While the results in the many-to-one setting are more promising, even here the preference lists appear to be on the large side. Also, while our rule for identifying the edges to include is simple, in practice it may well require too much communication to identify these edges. At the same time, our outcome is better than what is achieved in practice: we obtain a complete match with high probability, whereas in the NRMP setting a small but significant percentage of positions are left unfilled. Our conclusion is that it remains important to understand how to effectively select smaller sets of edges.

In the popularity model, it is reasonable for each agent to simply select their favorite partners. But in the current setting, which we consider to be more realistic, it would be an ineffective strategy, as it would result in most agents remaining unmatched. Consequently, we believe the main open issue is to characterize what happens when the number of edges k that an agent can list is smaller than the size of the allowable edge set. We conjecture that following a simple protocol for selecting edges to list, such as the one we use in our experiments (see Section 6.4), will lead to an ϵ-Bayes-Nash equilibrium, where ϵ is a decreasing function of k. Strictly speaking, as the identification of allowable edges requires communication, we need to consider the possibility of strategic communication, and so one would need to define a notion of ϵ-equilibrium akin to a Subgame Perfect equilibrium. We conjecture that even with this, it would still be an ϵ-equilibrium.

Finally, it would be interesting to resolve whether the experimentally observed near uniqueness of the stable matching for non-bottom agents is a property of the linear separable model. We conjecture that in fact it also holds in the bounded derivatives model.

References

1. Atila Abdulkadiroğlu, Parag A. Pathak, and Alvin E. Roth. The new york city high school match. *American Economic Review*, 95(2):364–367, May 2005. URL: https://www.aeaweb.org/articles?id=10.1257/000282805774670167, doi:10.1257/000282805774670167.
2. Itai Ashlagi, Mark Braverman, Yash Kanoria, and Peng Shi. Clearing matching markets efficiently: Informative signals and match recommendations. *Management Science*, 66(5):2163–2193, 2019. doi:10.1287/mnsc.2018.3265.
3. Itai Ashlagi, Yash Kanoria, and Jacob D. Leshno. Unbalanced random matching markets: The stark effect of competition. *Journal of Political Economy*, 125(1), 2017. doi:10.1086/689869.
4. Peter Coles, Alexey Kushnir, and Muriel Niederle. Preference signaling in matching markets. *American Economic Journal: Microeconomics*, 5(2):99–134, May 2013. URL: https://www.aeaweb.org/articles?id=10.1257/mic.5.2.99, doi:10.1257/mic.5.2.99.
5. D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962. URL: http://www.jstor.org/stable/2312726.
6. Hugo Gimbert, Claire Mathieu, and Simon Mauras. Incentives in popularity-based random matching markets, 2019. arXiv:1904.03890v1.
7. Yannai A. Gonczarowski, Noam Nisan, Lior Kovalio, and Assaf Romm. Matching for the israeli "mechinot" gap-year programs: Handling rich diversity requirements. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19, page 321, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3328526.3329620.

8 Yannai A. Gonczarowski, Noam Nisan, Rafail Ostrovsky, and Will Rosenbaum. A stable marriage requires communication. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, page 1003–1017, USA, 2015. Society for Industrial and Applied Mathematics.

9 Avinatan Hassidim, Assaf Romm, and Ran I. Shorrer. Redesigning the israeli psychology master's match. *American Economic Review*, 107(5):205–09, May 2017. URL: `https://www.aeaweb.org/articles?id=10.1257/aer.p20171048`, `doi:10.1257/aer.p20171048`.

10 Nicole Immorlica and Mohammad Mahdian. Incentives in large random two-sided markets. *ACM Trans. Econ. Comput.*, 3(3), jun 2015. `doi:10.1145/2656202`.

11 Yash Kanoria, Seungki Min, and Pengyu Qian. In which matching markets does the short side enjoy an advantage? In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1374–1386. SIAM, 2021.

12 Donald E. Knuth. *Mariages stables et leurs relations avec d'autres problèmes combinatoires : introduction à l'analyse mathémathique des algorithmes.* Les Presses de l'Université de Montréal, 1976.

13 Donald E. Knuth. *Stable Marriage and Its Relation to Other Combinatorial Problems: An Introduction to the Mathematical Analysis of Algorithms*, volume 10. CRM Proceedings & Lecture Notes, 1996.

14 Donald E Knuth, Rajeev Motwani, and Boris Pittel. Stable husbands. *Random Structures & Algorithms*, 1(1):1–14, 1990.

15 Fuhito Kojima and Parag A. Pathak. Incentives and stability in large two-sided matching markets. *American Economic Review*, 99(3):608–27, June 2009. URL: `https://www.aeaweb.org/articles?id=10.1257/aer.99.3.608`, `doi:10.1257/aer.99.3.608`.

16 Ron Kupfer. The influence of one strategic agent on the core of stable matchings. In *WINE*, 2020.

17 SangMok Lee. Incentive compatibility of large centralized matching markets. *The Review of Economic Studies*, 84(1):444–463, 2016.

18 Stephan Mertens. Random stable matchings. *Journal of Statistical Mechanics: Theory and Experiment*, 2005, 10 2005. `doi:10.1088/1742-5468/2005/10/P10008`.

19 nrmp.org. Results and data, 2021 main residency match, May 2021. URL: `https://www.nrmp.org/match-data-analytics/residency-data-reports/`.

20 Boris Pittel. The average number of stable matchings. *SIAM Journal on Discrete Mathematics*, 2(4):530–549, 1989.

21 Boris Pittel. On Likely Solutions of a Stable Marriage Problem. *The Annals of Applied Probability*, 2(2):358 – 401, 1992. `doi:10.1214/aoap/1177005708`.

22 Boris Pittel. On likely solutions of the stable matching problem with unequal numbers of men and women. *Mathematics of Operations Research*, 44(1):122–146, 2019.

23 Boris Pittel, Larry Shepp, and Eugene Veklerov. On the number of fixed pairs in a random instance of the stable marriage problem. *SIAM Journal on Discrete Mathematics*, 21(4):947–958, 2008.

24 Ignacio Rios, Tomás Larroucau, Giorgiogiulio Parra, and Roberto Cominetti. Improving the chilean college admissions system. *Operations Research*, 69(4):1186–1205, 2021. `doi:10.1287/opre.2021.2116`.

25 Alvin E. Roth and Elliott Peranson. The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American Economic Review*, 89(4):748–780, September 1999. URL: `https://www.aeaweb.org/articles?id=10.1257/aer.89.4.748`, `doi:10.1257/aer.89.4.748`.

26 Ran I. Shorrer. Simultaneous search: Beyond independent successes. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19, page 347–348, New York, NY, USA, 2019. Association for Computing Machinery. `doi:10.1145/3328526.3329599`.