Optimal truncation in matching markets [☆]

Peter Coles ^{a,*}, Ran Shorrer ^{b,c}

^a eBay Data Labs, United States

^b Harvard University, United States

^c Harvard Business School, United States

ARTICLE INFO

Article history:

Received 23 June 2013

Available online 24 February 2014

JEL classification:

C78

D47

Keywords:

Matching

Deferred acceptance

Truncation

Market design

Manipulation

ABSTRACT

Although no stable matching mechanism can induce truth-telling as a dominant strategy for all participants (Roth, 1982), recent studies have presented conditions under which truthful reporting by all agents is close to optimal (Immerlica and Mahdian, 2005; Kojima and Pathak, 2009; Lee, 2011). Our results demonstrate that in large, uniform markets using the Men-Proposing Deferred Acceptance Algorithm, each woman's best response to truthful behavior by all other agents is to truncate her list substantially. In fact, the optimal degree of truncation for such a woman goes to 100% of her list as the market size grows large. In general one-to-one markets we provide comparative statics for optimal truncation strategies: reduction in risk aversion and reduced correlation across preferences each lead agents to truncate more. So while several recent papers focused on the limits of strategic manipulation, our results serve as a reminder that without pre-conditions ensuring truthful reporting, there exists a potential for significant manipulation even in settings where agents have little information.

© 2014 Elsevier Inc. All rights reserved.

1. Introduction

One of the great success stories in economic theory is the application of matching theory to two-sided markets. A classic example is the National Resident Matching Program (NRMP), in which medical school students are matched to residency positions in hospitals. Rather than hospitals pursuing students via a decentralized series of offers, refusals and acceptances, matching occurs via a centralized mechanism. In this mechanism, each student ranks the hospital programs, and each hospital ranks the students. They submit these lists to an algorithm, which determines which students will be matched to which programs.

Such a centralized process has a number of advantages. First and foremost, the algorithm on which this and many similar centralized processes are based produces an outcome that is *stable* with respect to reported preferences.¹ In a stable matching, no two agents mutually prefer each other to their assigned match, nor does any matched participant prefer to be unmatched. A second advantage is that eliminating a decentralized offer process may save time and other resources.

[☆] We thank Yakov Babichenko, B. Douglas Bernheim, William Diamond, Michael Grubb, Avinatan Hassidim, John Hatfield, Guido Imbens, Jacob Leshno, Jonathan Levin, Paul Milgrom, Muriel Niederle, Mikko Packalen, and Al Roth for helpful discussions, and two anonymous referees for providing detailed comments. We thank the Stanford Institute for Economic Policy Research for their generous support.

* Corresponding author.

E-mail addresses: pacoles@gmail.com (P. Coles), rshorrer@hbs.edu (R. Shorrer).

¹ In 1998, the algorithm used in the NRMP was altered to accommodate student couples and allow for specialized hospital positions, so that the outcome is “close to” a stable matching (see Roth and Peranson, 1999).

Finally, as Roth and Xing (1994) have shown, a centralized mechanism can successfully halt the unraveling of a market.² Centralized matching mechanisms also power a variety of other markets, including the public school systems in New York, Boston, Singapore and other cities, as well as numerous specialized medical fellowships.

These centralized markets all employ versions of an algorithm proposed by Gale and Shapley (1962). The algorithm, which in one-to-one markets is often referred to as the *Men-Proposing Deferred Acceptance Algorithm* (MP-DA), takes as its inputs preference lists reported by agents, and outputs a stable matching. When agents are asked to report preference lists for submission to MP-DA, this begs the question: Do all agents have an incentive to report truthfully? Dubins and Freedman (1981) and Roth (1982) provide the answer: they do not. In fact, Roth showed that no mechanism that produces stable matchings will induce truth-telling as a dominant strategy for all agents. However, in the preference list submission game induced by MP-DA, for all participants on *one side* of the market, the “men,” truth-telling is a dominant strategy.³ But this leaves open the question of how participants on the other side of the market, the “women,” might benefit by strategically misrepresenting their preferences.

Recent work has examined conditions under which gains to strategic manipulation are limited for all participants in the market, not just those on one side. One approach in the literature concerns *large markets*. Roth and Peranson (1999) observe that in the data from the NRMP, very few participants could have improved their outcomes by reporting different preferences. They show via simulations that when the length of preference lists is held fixed and the number of participants grows, the size of the set of stable matching shrinks (a property they term “core convergence”), so that opportunities for manipulation are reduced. Immorlica and Mahdian (2005) demonstrate this result theoretically, finding that in large marriage markets where preference list length is bounded, nearly all players have an incentive to truthfully report preferences. Kojima and Pathak (2009) generalize this result, showing that in many-to-one markets, preference list manipulation, as well as other modes of strategic manipulation such as non-truthful reporting of capacities (see also Sönmez, 1997), are again limited. Lee (2011) considers one-to-one matching markets where agent utilities are drawn from distributions with bounded support that have both a common and an independent component. He shows that when all agents report truthfully, the proportion of participants who can achieve a significant utility gain from manipulation vanishes as the market grows large.⁴

Our approach takes a different tack. We do not require preference lists to be short, and ask: how should players optimally misrepresent preferences in markets that do not satisfy non-manipulability conditions? How “far” could optimal behavior be from truthfulness? We wish to study optimal manipulation, along with payoffs and market-wide welfare effects, and ask how strategic behavior and outcomes change as we vary market conditions.

The particular form of strategic misrepresentation we focus on is preference list *truncation*; that is, listing in order the first several partners from one's true preference list, and identifying all other partners as unacceptable. Truncation has an intuitive logic: by listing less-preferred partners as unacceptable, the probability of being matched with these partners drops to zero. Agents using this strategy might hope that correspondingly, the likelihood of being matched to a partner who remains on the truncated list will go up. In the context of MP-DA, this intuition is confirmed: submitting a truncated preference list weakly increases the likelihood of being matched to some agent on the truncated list, regardless of beliefs about the lists other agents submit. But submitting a truncated preference list is a risky strategy. Limiting acceptable partners also increases the likelihood of ending up with no match. Analysis of this tradeoff is the crux of the results in this paper.

While always a method for weakly increasing the likelihood of matching with better-ranked opponents, in some uncertain settings, truncation is optimal: Roth and Rothblum (1999) show that when agents' beliefs satisfy a form of symmetry termed “ $\mathcal{M}$ -symmetry,” they can do no better than to truncate. Ehlers (2004) demonstrated that this result holds under somewhat more general conditions.

Whether optimal or not, we analyze truncation, both in symmetric and general settings. We ask: to what degree should players truncate, if at all? (Note that submitting one's true preference list is also a form of truncation.) Can a participant realistically gain from truncation when she is extremely uncertain about what opponents might report? If players anticipate that others may be truncating, how does this affect their behavior? Do participants truncate in equilibrium? What are the welfare implications in a truncation equilibrium?

To evaluate the consequences of truncation, we first characterize the payoff from truncation for a woman with general beliefs over the preference lists other agents will submit in terms of the distributions of her most and least preferred achievable mates. In a market with N men and N women, when a woman believes submitted preferences of others are uniformly chosen from the set of all full-length preference lists, she may safely truncate a large fraction of her list with low risk of becoming unmatched. Further, as there is a large gap between the expected rank of the mate she receives from truthful revelation and her most preferred achievable mate (Pittel (1989) shows these asymptote to $N/\log N$ and $\log N$

² Before the NRMP was introduced in the 1950s, offers and interviews were made as early as the fall of students' third year in medical school, which was undesirable for a number of reasons. The willingness of both hospitals and students to participate argues strongly in favor of the program's effectiveness. The NRMP enjoys participation rates of close to 100% of eligible students, with over 38 000 students participating in the March 2012 match.

³ This is true in one-to-one, or “marriage” markets, where each agent has the capacity to match with at most one other agent. In many-to-one settings, e.g. students matching to hospitals, truth-telling is no longer a dominant strategy when, in the Deferred Acceptance Algorithm, the “hospitals” side makes the offers (see Roth, 1985).

⁴ In a very different approach, Featherstone and Mayefsky (2010) run lab experiments in 5×5 marriage markets, and find that participants have trouble learning to find beneficial deviations under MP-DA, even if there are potential gains (though participants have more success in finding successful manipulation when facing “priority” mechanisms).

respectively), truncation can lead to gains. The optimal degree of truncation can be significant. We demonstrate that in large, balanced, uniform markets the optimal level of truncation approaches 100%.⁵ That is, when there are many agents in the market, a woman optimally submits an extremely short list relative to her full preference list – in the limit, the fraction of men that she leaves on her list goes to zero.

The two sides of the market diverge in their tastes for truncation equilibria. Compared to the outcome from truthful reporting, women prefer any equilibrium in which they all use truncation strategies; for men the opposite is true. Furthermore, if there are two truncation equilibria that can be compared in the degree of truncation, women prefer the equilibrium in which they truncate more, while men prefer the equilibrium in which women truncate less. Under uniform preferences, we demonstrate the existence of a symmetric equilibrium where all women use the same truncation strategy. However, even under uniform preferences, asymmetric equilibria, where women vary in their degrees of truncation, may also exist. In such equilibria, and in contrast to the across-equilibria results, the women who truncate *least* are best off. Intuitively, while women benefit from truncation by other women, they would prefer not to bear the risk of truncation themselves.

Relaxing the uniform preferences assumption and returning to the environment where players have arbitrary beliefs, we examine comparative statics. We find that optimal truncation levels vary with risk preference: regardless of beliefs over reported opponent preferences, the less risk-averse a player, the more she should truncate.

We then turn to correlation in players' preferences. The correlated preferences we consider are meant to capture the notion that in many settings, agents largely agree in their preferences over partners on the other side of the market, but that an individual's preferences may idiosyncratically depart from common opinion. We find that the higher the likelihood a participant places on opponents having preferences similar to her own, the less she should truncate. Our findings largely corroborate the simulation results of Roth and Peranson (1999), who find that when preferences are correlated, the set of stable matchings is small, and therefore the set of submitted preference lists that could lead to gains is minimal. An important difference between our correlation setting and Roth and Peranson's is that we consider incomplete information, where realized matchings may be unstable with respect to true preferences (even while stable with respect to reported preferences).

To place this analysis in context, several comments are in order. While, for the reasons stated earlier, we believe truncating the *bottom* of one's list is an intuitive manipulation in the preference list submission problem, in different environments eliminating *better-ranked* members from one's preference list might be a reasonable strategy. For example, in the job market for economists, departments may choose not to interview certain highly-accomplished candidates, reasoning that these candidates will receive offers they prefer more (see Coles et al., 2010). In efforts to best use costly and scarce interview slots, departments may effectively “top-truncate” their preferences lists, focusing instead on candidates more likely to ultimately accept an offer. In general, it is when market frictions generate costs that this behavior arises. Lee and Schwarz (2012) consider a setting where information acquisition is costly, so that firms prefer to interview workers who have a high likelihood of accepting (and likelihood is based on the number and identity of other firms interviewing a worker). Coles et al. (2013) consider a setting where workers can signal their preferences to firms, so that firms may choose not to make offers to better-ranked candidates, and instead make offers to candidates who have indicated likeliness to accept. In our paper, the analysis is performed after any costly information gathering has taken place, so these considerations do not arise.

The rest of the paper is organized as follows: Section 2 lays out the stable marriage setting and illustrates the fundamental tradeoff associated with truncation. In Section 3 we characterize the return to truncation, first for general beliefs, then in a uniform setting. In Section 4, we prove the existence of a truncation equilibrium in symmetric settings, and explore equilibrium welfare implications. Based on this existence result, we provide simulation evidence for a significant degree of truncation in equilibrium. Section 5 and Section 6 examine how truncation behavior relates to risk preferences and correlation of agent preferences, respectively. Section 7 concludes.

2. Matching markets background

We begin by setting out the basic model of matching. In contrast to some of the well-known matching papers, we approach the notion of preferences of participants from a cardinal rather than an ordinal perspective, which allows us to discuss choice under uncertainty.⁶ Note, however, that standard matching results involving ordinal preferences also apply as we may infer preference orderings from cardinal utilities.

2.1. Marriage markets and stability

A *marriage market* of size N consists of a triplet $(\mathcal{M}, \mathcal{W}, u)$, where $\mathcal{M}$ is the set of men, $\mathcal{W}$ is the set of women, $|\mathcal{M}| = |\mathcal{W}| = N$, and $u = \prod_{i \in \mathcal{M} \cup \mathcal{W}} u_i$ is the profile of preferences for men and women.⁷

⁵ Recently, Ashlagi et al. (2013) have shown that the requirement that the number of men and women is balanced is crucial for the existence of the large gap in expected ranks. We discuss the robustness of our results to imbalances in the number of men and women in Appendix A.

⁶ Some matching papers do manage to study choice under uncertainty even when agents have only ordinal preferences. For example, Ehlers and Masso (2007) and Roth and Rothblum (1999) use the related concepts of “Ordinal Bayesian Nash Equilibrium” and “ P_w -stochastic dominance,” respectively.

⁷ The assumption $|\mathcal{M}| = |\mathcal{W}|$ was made for technical and notational convenience. However, it plays an important role in Theorem 1, as discussed in Appendix A.

Preferences $u_m : \mathcal{W} \cup \{m\} \rightarrow \mathbb{R}$ for man $m \in \mathcal{M}$ are given by a von Neumann–Morgenstern utility function in which m derives utility $u_m(w)$ from matching with woman w and $u_m(m)$ from remaining single. For simplicity, we assume that u_m is one-to-one, so that m is never indifferent between any two certain options. Preferences u_w for woman $w \in \mathcal{W}$ are defined similarly on $\mathcal{M} \cup \{w\}$.

As u_m is one-to-one, m 's preferences u_m induce a strict preference ordering P_m over $\mathcal{W} \cup \{m\}$. We refer to P_m as m 's *preference list*. For example, if $N = 3$, $u_m(w_1) > u_m(w_3) > u_m(w_2) > u_m(m)$ yields preference list (w_1, w_3, w_2, m) , meaning m prefers woman w_1 to w_3 to w_2 to being single. Note that men may prefer bachelorhood over some of the women. For example, (w_1, w_3, m, w_2) indicates that m prefers w_1 to w_3 to remaining single to w_2 . We say that man m finds w *acceptable* if m prefers w to remaining single. When convenient, we list only a man's acceptable women. Preference lists for women are defined similarly, and we let P denote the profile of preference lists.

A *matching* is a pairing of men and women, so that each woman is assigned at most one man and each man at most one woman. Formally, a matching μ is a mapping from $\mathcal{M} \cup \mathcal{W}$ to $\mathcal{M} \cup \mathcal{W}$ such that for every $m \in \mathcal{M}$, $\mu(m) \in \mathcal{W} \cup \{m\}$, and for every $w \in \mathcal{W}$, $\mu(w) \in \mathcal{M} \cup \{w\}$, and also for every $m, w \in \mathcal{M} \cup \mathcal{W}$, $\mu(m) = w$ if and only if $\mu(w) = m$. When $\mu(x) = x$, agent x is *single* or *unmatched* under matching μ . For agents that are not single, we refer to $\mu(m)$ as m 's wife and $\mu(w)$ as w 's husband. The terms partner and mate are also used. In a matching, each agent cares only about his or her partner, and not about the partners of other agents, so that we may discuss agent preferences over matchings.

Given preferences, a matching is *stable* if no agent desires to leave his or her mate to remain single, and no pair of agents mutually desire to leave their mates and pair with each other. Formally, given a matching μ , we say that it is *blocked* by (m, w) if m prefers w to $\mu(m)$ and w prefers m to $\mu(w)$. A matching μ is *individually rational* if for each $x \in \mathcal{M} \cup \mathcal{W}$ with $\mu(x) \neq x$, x finds $\mu(x)$ acceptable. A matching μ is *stable* if it is individually rational and is not blocked. In general, more than one stable matching may exist for given preferences.

Given preferences, a woman w is *achievable* for m if there is *some* stable matching μ in which $w = \mu(m)$. Achievable mates of women are defined similarly.

2.2. The Men-Proposing Deferred Acceptance Algorithm

In their seminal 1962 paper, Gale and Shapley prove that in any marriage market there exists a stable matching. To demonstrate this result, they propose an algorithm, the Men-Proposing Deferred Acceptance Algorithm (MP-DA), to generate a stable matching given any profile of preferences lists.

MP-DA takes as its input a preference list profile P for agents $\mathcal{M} \cup \mathcal{W}$, and the output is a matching $\mu^M[P]$. When P is clear from the context, we write μ^M to denote $\mu^M[P]$.⁸ The algorithm works iteratively as follows:

- **Step 1.** Each man proposes to the first woman on his preference list. Each woman then considers her offers, rejects all men deemed unacceptable, and if any others remain, rejects all but her most preferred mate.
- **Step k .** Each man who was rejected in step $k - 1$ makes an offer to the next woman on his preference list. If his preference list is exhausted, or if he prefers bachelorhood to the next woman on his list, he makes no offer. Each woman behaves as in step 1, considering offers in hand (including any man she has retained from the previous step) and rejects all but her most preferred acceptable suitor.
- **Termination.** If in any step k , no man makes an offer, the algorithm terminates. Each woman is paired with her current mate and this matching is final.

Gale and Shapley show that this algorithm must terminate in finite time, and they provide a remarkable characteristic of the resulting outcome.

Theorem (Gale–Shapley). *The matching μ^M resulting from MP-DA is stable. Furthermore, for any other stable matching μ , every man weakly prefers μ^M to μ and every woman weakly prefers μ to μ^M .*

The stability of the matching produced by MP-DA offers numerous advantages, as outlined in the introduction. But men are particularly satisfied with this outcome. For the men, the algorithm produces the *optimal* stable matching, based on reported preferences. For the women, however, this is not the case. As we will see, this feature may mean some women prefer to strategically misreport preferences, causing the algorithm to produce a different matching.

2.3. The Preference List Submission Problem for men

We now turn to the incentive properties of MP-DA. That is, in a setting where agents are asked to submit preference lists to the algorithm, we ask if they have an incentive to report something other than the truth. We will see that women may, while men do not.

⁸ In one-to-one markets, the women-proposing version of the algorithm (WP-DA) has identical but reversed properties, with output denoted by $\mu^W[P]$.

Consider a set of agents $\mathcal{M} \cup \mathcal{W}$. Agent $i \in \mathcal{M} \cup \mathcal{W}$ with preferences u_i must submit a preference list $\hat{P}_i$ to MP-DA, where $\hat{P}_i$ is chosen from the set of i 's possible preference lists $\mathcal{P}_i$. The agent's beliefs about what preference lists others will report are represented by the random variable $\tilde{P}_{-i}$, which takes as its range $\mathcal{P}_{-i}$, the set of all possible preference list profiles for others. Note that since u_i is a von Neumann–Morgenstern utility function, agent i may compare outcomes in this incomplete information setting.

Agent i solves the *Preference List Submission Problem*:

$$\max_{\hat{P}_i \in \mathcal{P}_i} \mathbb{E}[u_i(\mu^M[\hat{P}_i, \tilde{P}_{-i}](i))].$$

Dubins and Freedman (1981) and Roth (1982) have shown that for any man m with preferences u_m and beliefs $\tilde{P}_{-m}$, it is optimal for m to submit his true preference list P_m (which corresponds to u_m).

Theorem (Dubins and Freedman; Roth). *In the Preference List Submission Problem,*

$$P_m \in \arg \max_{\hat{P}_m \in \mathcal{P}_m} \mathbb{E}[u_m(\mu^M[\hat{P}_m, \tilde{P}_{-m}](m))].$$

2.4. The Preference List Submission Problem for women

For women submitting preference lists to MP-DA, truth-telling may not be optimal. One way a woman w might misrepresent preferences is by submitting a *truncation* of her true preference list; that is, listing in order the first several men from her true preference list and declaring all other men unacceptable. Truncation generates a tradeoff: it may cause a woman to match with the better-ranked men she leaves on her list, but may also cause her to be left unmatched. In this section we demonstrate this tradeoff, pose the problem of optimal truncation, and describe conditions so that in the Preference List Submission Problem, among all possible preference list submission strategies, truncation is optimal.

The following example demonstrates the tradeoff at hand.

Example 1 (*Strategic truncation*). Suppose men and women have the following preference lists:

m_1	m_2	m_3	w_1	w_2	w_3
w₃	w₂	w₁	m_1	m_1	m₁
w_1	w_1	w_2	m_2	m_3	m_2
w_2	w_3	w_3	m₃	m₂	m_3

We consider the strategic incentives of woman 1, assuming all other agents report truthfully. First, suppose w_1 submits her true preference list. In this case, MP-DA stops after one step and w_1 is matched to m_3 , her least preferred mate. The stable matching is indicated above in bold.

Now suppose that w_1 misrepresents her preferences and submits the truncated list (m_1, m_2) . In this case, she will reject man m_3 's first round offer in the MP-DA. Man m_3 must then make an offer to w_2 in the next round. Woman w_2 will accept m_3 over m_2 , who made her an offer in the previous round. Man m_2 then finds himself single, and must make an offer to w_1 . Woman w_1 accepts m_2 's offer and MP-DA terminates, yielding the matching in bold below:

m_1	m_2	m_3	w_1	w_2	w_3
w₃	w_2	w_1	m_1	m_1	m₁
w_1	w₁	w₂	m₂	m₃	m_2
w_2	w_3	w_3		m_2	m_3

Therefore, by truncating her list, w_1 improves her outcome.

To see how truncation can be dangerous, suppose w_1 truncates her list even more and submits (m_1) only. In this case, the algorithm will leave her unmatched, as shown in bold below:

m_1	m_2	m_3	w_1	w_2	w_3
w₃	w_2	w_1	m_1	m_1	m₁
w_1	w_1	w₂		m₃	m_2
w_2	w_3	w_3		m_2	m_3

Remark 1 characterizes woman w 's match when she submits a truncated version of her preference list, demonstrating generally how truncation can lead to the three outcomes in our example. For $k \in \{0, \dots, N\}$, we denote by P_w^k the preference list which includes in order only w 's k most preferred men, and call this the k -truncation of her true preference list P_w . If fewer than k men are acceptable to w , then $P_w^k \equiv P_w$.

Remark 1. Let P be the preference list profile of all agents in $\mathcal{M} \cup \mathcal{W}$. Then $\mu^M[P_w^k, P_{-w}](w)$ is w 's least preferred achievable mate under P with rank $\leq k$. Should no such mate exist, $\mu^M[P_w^k, P_{-w}](w) = w$.

The example illustrates a general principle; given the preference lists submitted by others, truncation by woman w can have one of three consequences:

1. No effect. Woman w has truncated below her least preferred achievable mate.
2. Improvement. Woman w truncates above her least preferred mate, and is matched with her least preferred achievable mate above the point of truncation.
3. Unmatched. Woman w has over-truncated, truncating above her most preferred achievable mate.

If woman w is certain of the preference lists P_{-w} others are submitting, her truncation decision is simple: she calculates her most preferred achievable mate under $P = (P_w, P_{-w})$ and truncates her list to just include him. If instead w believes her opponents will submit preference lists according to some probability distribution, then truncating her list at k generates a lottery over outcomes in which either her partner will be among her k most preferred men, or else she will be unmatched. This tradeoff, between improvement and becoming unmatched, will guide our analysis in this paper.

2.4.1. Optimality of truncation

Truncation is not the only possible misrepresentation of preferences. A woman could reverse two men in her preference list, list men as acceptable who are in fact unacceptable, drop men from the middle of her list, or use some combination of these. However, under some conditions, truncation is optimal.

The next proposition states that under certainty, women can do no better than to truncate (Roth and Vande Vate, 1991).

Proposition (Roth and Vande Vate). Suppose woman w has preferences u_w and knows others will report preference lists P_{-w} to MP-DA. Then truncating such that $\mu^W(w)$ is the last acceptable partner on her list is an optimal strategy for w .

Perhaps surprisingly, when a woman has very little information about the preference lists others might report, she again can do no better than to truncate. In order to gain from non-truncation misrepresentations, such as swapping the positions of two men in her reported preference list, a woman must have very specific information about the preference lists others report. Without such information, it is best to leave the men in their correct order. Roth and Rothblum (1999) demonstrate this principle using the following framework.

Let woman w 's beliefs about reported preference lists of others be represented by $\tilde{P}_{-w}$, a random variable taking on values in $\mathcal{P}_{-w}$. If P_{-w} is a preference list profile for agents $-w$, define $P_{-w}^{m \leftrightarrow m'}$ to be the preference list profile in which m and m' swap preference lists, and all women swap the positions of m and m' in their lists. We say that woman w 's beliefs are (m, m') -symmetric if $\Pr(\tilde{P}_{-w} = P_{-w}) = \Pr(\tilde{P}_{-w} = P_{-w}^{m \leftrightarrow m'})$ for all $P_{-w} \in \mathcal{P}_{-w}$. For a subset $\mathcal{M}' \subseteq \mathcal{M}$, beliefs $\tilde{P}_{-w}$ are $\mathcal{M}'$ -symmetric if they are (m, m') -symmetric for all $m, m' \in \mathcal{M}'$.

Theorem (Roth and Rothblum). Suppose w 's beliefs about reported preference lists of others are $\mathcal{M}$ -symmetric. Then any preference list $\hat{P}_w$ she might submit to MP-DA is weakly P_w -stochastically dominated by some truncation of her true preference list.⁹

Hence, when w is certain about reported preference lists of her opponents, or when she has extreme, symmetric uncertainty, truncation is optimal.

2.4.2. The Truncation Problem

Even when truncation is not optimal, we may sometimes wish to restrict the choice set for women to truncations of her true preference list. We define the *Truncation Problem* for woman w with preferences u_w and beliefs $\tilde{P}_{-w}$ on others' submitted preference lists as

$$\max_{k \in \{0, \dots, N\}} \mathbb{E}[u_w(\mu^M[P_w^k, \tilde{P}_{-w}](w))].$$

For convenience, whenever we consider the Truncation Problem for a woman w , we will relabel men so that w has $u_w(m_1) > u_w(m_2) > \dots > u_w(m_N)$.

3. Characterizing truncation payoffs

In this section we explore a woman's payoff from submitting a truncation of her true preference list. We first derive a formula for her payoff from truncating at any point in her list in terms of the distribution of her most and least preferred

⁹ $\hat{P}_w$ is P_w -stochastically dominated by $\hat{P}_w'$ iff for any vNM utility function that corresponds to P_w , the expected utility from submitting $\hat{P}_w'$ is at least as great as the expected utility from submitting $\hat{P}_w$.

achievable mates. When a woman believes that reported preference lists of her opponents are distributed symmetrically over the set of all preference lists, we can say more about her gains from truncation: conditional on truncation yielding an improvement, w 's expected partner rank will be exactly half of $(1 + \text{her point of truncation})$. Further, in uniform markets, it is highly likely that for an individual woman, some degree of truncation will yield an improvement, and that, in fact, she may safely and beneficially truncate a large fraction of her list. We demonstrate that as the market grows large, the length of a woman's optimal reported list, as a fraction of her full preference list length, goes to zero. Hence, even in a setting where agents possess very little information about opponent preferences, there is room for significant strategic misrepresentation.

3.1. Truncation in two stages: Match, then divorce

To aid us in our analysis, we show that when woman w submits a k -truncation of her preference list to *MP-DA*, the outcome is identical to that from a two stage *Divorcing Algorithm*. In the first stage of the algorithm, agents submit preference lists to *MP-DA*. In the second stage, w 'divorces' her mate and declares all men with rank $\geq k$ unacceptable. This sets off a chain of new offers and proposals, ending in a new match.¹⁰

The Divorcing Algorithm takes as its input a set P of preference lists, a woman w , and a truncation point $k \in \{0, \dots, N\}$, and generates a matching $\mu^{\text{DIV}}[P, k, w]$.

The Divorcing Algorithm.

- **Step 0.** Initialization. Run *MP-DA* to find the men-optimal matching $\mu^M[P]$. If w is single or if w 's mate has rank $\leq k$ in P_w , terminate. Otherwise, divorce w from her mate. Declare candidates with rank $\geq k$ unacceptable for w . Iteration over steps 1 and 2:
- **Step 1.** Pick an arbitrary single man who has not exhausted his preference list. If no such man exists, terminate.
- **Step 2.** The man chosen in the previous step makes an offer to the most preferred woman on his preference list who has not already rejected him. If this woman finds the man acceptable and she prefers him to her current mate (or is single), she divorces (if necessary) and holds the new man's offer. Return to step 1.

The Divorcing Algorithm yields a matching identical to the output from w 's submission of a k -truncated list to *MP-DA*:

Proposition 1. For all $k \in \{0, \dots, N\}$, $P \in \mathcal{P}$ and $w \in \mathcal{W}$, we have $\mu^{\text{DIV}}[P, k, w] = \mu^M[P_w^k, P_{-w}]$.

With this equivalence in hand, when we consider the submission of a truncated preference list, we can think of it as a two stage process, focusing on the chain of offers (if there is one) in the Divorcing Algorithm. We will be interested in whether a chain will end with *i*) a new acceptable man proposing to woman w , or with *ii*) a single man making an acceptable offer to a single woman in $\mathcal{W} \setminus w$, or else exhausting his list. These outcomes correspond to truncation yielding an improvement over truthful reporting, and truncation leaving w unmatched, respectively. Knowing that following a "divorce", w will receive at most one more offer will enable us to calculate the returns to truncation, conditional on truncation yielding an improvement.

3.2. Truncation under general beliefs

In this section, we characterize woman w 's payoff from submitting a truncated version of her true preference list in terms of the distributions of her most and least preferred achievable mates. The results build on [Remark 1](#), which illustrates how in settings of certainty, a woman may gain, lose or see no change from truncation.

We consider the Truncation Problem for woman w with preferences u_w and beliefs $\tilde{P}_{-w}$ about reports of other agents. Throughout the section, u_w (and hence, P_w) is fixed, so we can denote w 's payoff from k -truncation when others submit preference lists P_{-w} as

$$v(k, P_{-w}) \equiv u_w(\mu^M[P_w^k, P_{-w}](w)).$$

Note that $v(N, P_{-w})$ gives w 's payoff if she reports truthfully, and $v(k, P_{-w}) = u_w(w)$ if k -truncation leaves w unmatched. The Truncation Problem then becomes

$$\max_{k \in \{0, \dots, N\}} \mathbb{E}[v(k, \tilde{P}_{-w})].$$

To evaluate $\mathbb{E}[v(k, \tilde{P}_{-w})]$, we condition on the three possible effects of truncation: no effect, improvement, and causing w to become unmatched.

¹⁰ Our Divorcing Algorithm is closely related to the techniques used in [McVitie and Wilson \(1971\)](#), where a "breakmarriage" operation is used to generate all the stable matchings for a given marriage market.

Define $k_l(P_{-w})$ and $k_h(P_{-w})$ to be w 's rank of her mate under $\mu^M[P]$ and $\mu^W[P]$, respectively. That is, $k_l(P_{-w})$ ($k_h(P_{-w})$) gives the rank of w 's least (most) preferred achievable mate when $-w$ report preference lists P_{-w} . Set $k_l(P_{-w}) = k_h(P_{-w}) = N + 1$ when w has no achievable mate. Let $f(\cdot)$ be the probability mass function of the random variable $k_l(\tilde{P}_{-w})$ so that

$$f(x) = \Pr(k_l(\tilde{P}_{-w}) = x)$$

for $x \in \{1, \dots, N + 1\}$. Let $F(\cdot)$ be the associated cumulative distribution function. Similarly, let $g(\cdot)$ be the probability mass function and $G(\cdot)$ be the cumulative distribution function of the random variable $k_h(\tilde{P}_{-w})$.

Using $F(\cdot)$, $G(\cdot)$, and Remark 1, we can express w 's expected payoff from k -truncating her list by using the law of conditional expectations:

$$\begin{aligned} \mathbb{E}[v(k, \tilde{P}_{-w})] &= F(k) \cdot \sum_{i=1}^k \frac{f(i)}{F(k)} u_w(m_i) \\ &\quad + [G(k) - F(k)] \cdot \mathbb{E}[v(k, \tilde{P}_{-w}) \mid \tilde{P}_{-w} \in \mathcal{P}_2(k)] \\ &\quad + [1 - G(k)] \cdot u_w(w), \end{aligned} \quad (3.2.1)$$

where the set $\mathcal{P}_2(k) \equiv \{P_{-w} \mid v(k, P_{-w}) > v(N, P_{-w})\}$ gives the cases when truncation yields an improvement, compared to truthful reporting. When truncation causes w to be unmatched, her payoff is clearly $u_w(w)$, and when truncation has no effect, the likelihood of being matched with x is $f(x)/F(x)$.¹¹

In the next two sections, we will focus on the middle term of the sum in (3.2.1); that is, the cases where truncation yields improvement. We will first see that when there are gains, the improvement can be significant. We will see later that in large markets, these gains may outweigh the risk of being left unmatched.

3.3. Truncation under $\mathcal{M}$ -symmetric beliefs

In this section, we examine the Truncation Problem when woman w has $\mathcal{M}$ -symmetric beliefs. We show that conditional on w 's truncation yielding an improvement compared to truthful reporting, her mate is equally likely to be any man she lists as acceptable. This is somewhat surprising, because when w has unconditional $\mathcal{M}$ -symmetric beliefs and submits preferences in the MP-DA, we would certainly not expect w 's mate to be uniformly distributed; because of the stability requirement, she is far more likely to be matched with her more preferred mates.

Lemma (Truncation under $\mathcal{M}$ -symmetric beliefs). Suppose woman w 's beliefs $\tilde{P}_{-w}$ about the reported preference lists of her opponents are $\mathcal{M}$ -symmetric. Then according to her beliefs,

$$\Pr\{\mu^M[P_w^k, \tilde{P}_{-w}](w) = m_i \mid \tilde{P}_{-w} \in \mathcal{P}_2(k)\} = \Pr\{\mu^M[P_w^k, \tilde{P}_{-w}](w) = m_j \mid \tilde{P}_{-w} \in \mathcal{P}_2(k)\}$$

for all $k \in \{1, \dots, N\}$, $i, j \in \{1, \dots, k\}$. Hence,

$$\mathbb{E}[v(k, \tilde{P}_{-w}) \mid \tilde{P}_{-w} \in \mathcal{P}_2(k)] = \frac{\sum_{i=1}^k u_w(m_i)}{k}.$$

The intuition in this result comes from the Divorcing Algorithm. Consider the settings where k -truncation will yield an improvement for w ($\tilde{P}_{-w} \in \mathcal{P}_2(k)$). By first reporting her true preferences and then divorcing her mate, we know that there will ensue a chain of offers. This chain ends when exactly one man – ranked better than her former mate – will make an offer to w . By the symmetry of w 's beliefs, this is equally likely to be any of these men.

Crucial to the reasoning is that since we know truncation will yield an improvement, this corresponds to an algorithmic outcome where exactly one new superior offer is made to w . In MP-DA generally, multiple offers may be made to w , making it difficult to pinpoint the distribution of her mate's rank.¹²

With our lemma in hand, we can now express w 's expected payoff from truncation at k as

$$\begin{aligned} \mathbb{E}[v(k, \tilde{P}_{-w})] &= F(k) \cdot \sum_{i=1}^k \frac{f(i)}{F(k)} u_w(m_i) \\ &\quad + [G(k) - F(k)] \cdot \frac{\sum_{i=1}^k u_w(m_i)}{k} \\ &\quad + [1 - G(k)] \cdot u_w(w). \end{aligned} \quad (3.3.1)$$

¹¹ If $F(k) = 0$, the first term in the sum is zero.

¹² This result is related to the Principle of Deferred Decisions (“don't do today what you can put off until tomorrow”), which was applied to the stable marriage problem in Knuth (1976). We may think of woman w as deferring her views on her preferences over men $\{1, \dots, k\}$ until she actually receives an offer from one of them. When the first one arrives, only then does she assign the man a rank, which in expectation will be $\frac{1+k}{2}$.

3.4. Optimal truncation in large markets

We now investigate optimal truncation for women when the market size grows large. We will focus on the special case of *uniform beliefs*. That is, when facing the Truncation Problem, w believes reported preference lists P_{-w} of her opponents to be chosen uniformly and randomly from the set of all possible full preference list profiles $\mathcal{P}_{-w}$ (where a full preference list profile is a profile in which each agent prefers any possible mate to being unmatched). Uniform beliefs are a special case of $\mathcal{M}$ -symmetric beliefs.¹³ Hence, under uniform beliefs, we can be sure that truncation is optimal.

The stable marriage problem under uniform beliefs has received attention, especially in the mathematics and computer science literature, in large part because this setting facilitates average and worst case analyses (see [Dzierzawa and Oméro, 2000](#); [Knuth, 1976](#), and [Pittel, 1989](#)). But for our purposes, uniform beliefs offer a tractable incomplete information setting where agents know little about the preferences of others.

Suppose that woman w has preferences $u_w(\cdot)$ linear in the rank of her match (where being unmatched is treated as rank $N + 1$), or else has any strictly increasing, convex transformation of such preferences. Suppose further that w has uniform beliefs. Define $k^*(N) \equiv \max(\arg \max_{k \in \{0, \dots, N\}} \mathbb{E}[u_w(\mu^M[P_w^k, \tilde{P}_{-w}](w))])$. For a market of size N , $k^*(N)$ describes woman w 's optimal point of truncation, given that the other women submit their true preference lists. If there are multiple optima, we conservatively select that which involves the least truncation. We now have the following theorem.

Theorem 1. *Let woman w have uniform beliefs and preferences linear in rank (or any strictly increasing, convex transformation of such preferences). Then $\lim_{N \rightarrow \infty} \frac{k^*(N)}{N} = 0$.*

[Theorem 1](#) states that as the market size grows large, the fraction of the list that an individual woman optimally truncates goes to 100%.¹⁴ Note that under uniform beliefs, Roth and Rothblum's optimality theorem applies. Hence, the aggressive truncation strategies described in the theorem are the best overall strategies, not just the optimal truncation strategies.

The intuition behind this theorem can be gleaned from statistical facts about the most and least preferred achievable mates for women. In large markets where preferences are uniform, the expected rank of the most preferred achievable mate of a woman (which is the same as the expected rank of her mate under *WP-DA*) is very low relative to the length of her list; it asymptotes to $\log N$ ([Pittel, 1989](#)). This suggests that a woman may safely truncate a large fraction of her list with little risk of becoming unmatched. Furthermore, the expected rank of a woman's match under *MP-DA* is significantly worse, asymptoting to $\frac{N}{\log N}$ ([Pittel, 1989](#)). In fact, for large markets, [Pittel \(1992\)](#) proved that the worst-off wife will be matched with a husband at the bottom of her list with probability approaching 1. This large gap in a woman's expected most and least preferred achievable mates suggests that not only is it safe to truncate a large fraction of one's list in large markets, but that a woman will also generate gains from such a truncation. See [Appendix A](#) for the details of the proof, as well as a discussion of the unbalanced case in which these properties cease to hold ([Ashlagi et al., 2013](#)).

To get a sense of the impact of truncation and to see examples of optimal truncation levels, we simulate markets of size $N = 10, 100, 1000$, and 10000 ([Fig. 1](#)). In each market, we randomly generate a full preference list for each agent, and calculate an individual woman's payoff from truncating at each point in $\{0, \dots, N\}$, where a woman's payoff is given by $(N + 1) -$ her partner's rank. We then iterate 100 000 times and average her payoffs.¹⁵

Observe that under truthful reporting, w 's payoff (given by the right hand side intercepts) is very close to $\frac{N}{\log N}$, the asymptotic limit found in [Pittel \(1989\)](#). Even in the largest market we simulated, $N = 10000$, w 's payoff at the peak is roughly 10% higher than her payoff from truthful reporting. Note further that in each market, peak utility is lower than $N + 1 - \log N$, the asymptotic expected rank of a woman's most preferred achievable mate. A woman will never be able to do better than this, even with perfect information about reported preferences of others.

In each of the graphs in [Fig. 1](#), and especially when $N = 1000$ and $N = 10000$, there is a flat area on the right hand side. These lower levels of truncation are unlikely to have any impact on payoffs; indeed establishing this is a crucial step in the proof of [Theorem 1](#). Additional truncation can generate better mates, but still bears little risk of over-truncation. Finally, for extreme levels of truncation there is a high probability of over-truncation, leading to a steep dropoff in payoffs. As N grows larger, the "safe range" increases: we obtain larger flat zones and peaks moving to the left.

4. Truncation equilibria and welfare

In this section we consider the Bayesian game in which agents must submit preference lists to *MP-DA*. We demonstrate that in *equilibria in truncation strategies*, compared to outcomes from truthful preference list reporting, welfare for men is

¹³ But uniformity is not equivalent to $\mathcal{M}$ -symmetry. Under $\mathcal{M}$ -symmetric beliefs, a woman may have specific knowledge about how the men rank her. For example, she may know that all the men prefer her to w_2 . With uniform beliefs, such knowledge is ruled out.

¹⁴ We thank a referee for pointing out that this theorem will continue to hold if all agents drew uniformly at random preferences of length αN for $\alpha \in (0, 1)$. We omit formal proof of this statement which would require slight adaptations to the theorems that we cite in our own proof.

¹⁵ Formally, we are estimating $\mathbb{E}[v(k, \tilde{P}_{-w})]$, the expectation of a random variable, by averaging many independent draws of $v(k, \tilde{P}_{-w})$. With 100 000 draws, the 95% confidence intervals around each estimate of $\mathbb{E}[v(k, \tilde{P}_{-w})]$ are so small that they are imperceptible when drawn on the graphs. For example, when $N = 1000$ and $k = 500$, the estimate of $\mathbb{E}[v(500, \tilde{P}_{-w})]$ is 874.5, and the 95% confidence interval is (873.8, 875.2).

Fig. 1. Simulation results for truncation payoffs. The graphs display $(N + 1)$ – an individual woman's expected partner rank from truncating her list at each point $k \in \{0, \dots, N\}$ and submitting these preferences to MP-DA. Preference lists of the other agents are uniformly random, selected from the set of all possible full length preference list profiles, and payoffs are averaged over 100 000 draws. Markets are of size 10, 100, 1000 and 10000.

lower, welfare for women is greater, and the expected number of matches is lower. When there are multiple equilibria that can be compared in degree of truncation, women prefer the equilibrium where they truncate most, while men prefer the equilibrium where they truncate least. In *uniform markets*, we demonstrate the existence of a symmetric equilibrium in truncation strategies, but asymmetric equilibria may also exist. In a truncation equilibrium where some women truncate more than others, the women who truncate less receive higher payoffs. That is, while across equilibria women prefer to see higher degrees of truncation, within an equilibrium, they prefer not to be the ones bearing the risk of truncating.

4.1. Equilibria under general preferences

Define the *Preference List Submission Game* as follows: Let U be a finite subset of the set of all possible utility profiles for $I = \mathcal{M} \cup \mathcal{W}$ and $\phi(\cdot)$ be a distribution over U . Let the message space of any agent $i \in I$ be $\mathcal{P}_i$, the set of all possible strict preference lists for player i , with $\mathcal{P} = \prod_i \mathcal{P}_i$. Recall that $\mu^M[P]$ gives the MP-DA matching for reported preference lists P . The *Preference List Submission Game* is the Bayesian game described by

$$(I, \mathcal{P}, \mu^M[\cdot], U, \phi(\cdot)).$$

A pure strategy for agent i is a mapping $s_i : U_i \rightarrow \mathcal{P}_i$, and a mixed strategy for i is a mapping $\sigma_i : U_i \rightarrow \Delta(\mathcal{P}_i)$ which describes a randomization over submitted preference lists for each possible type. Define a *truncation strategy* for woman w as a strategy in which w mixes over truncations of her true preference list. For any two truncation strategies σ_w and σ'_w for a woman w , we say that σ_w involves more truncation than σ'_w if the distribution over truncation points induced by σ_w first order stochastically dominates the distribution induced by σ'_w .

We will restrict attention to equilibria where men report preferences truthfully, an assumption motivated by the dominant-strategy result of Dubins and Freedman (1981) and Roth (1982).¹⁶ Define a Bayesian Nash equilibrium $\sigma = (\sigma_{m_1}, \dots, \sigma_{m_N}, \sigma_{w_1}, \dots, \sigma_{w_N})$ in which men report truthfully and women mix over truncation strategies as an *equilibrium in truncation strategies*. The following theorem describes welfare in such equilibria.

Theorem 2. Let σ and σ' be equilibria in truncation strategies in which each woman truncates more under σ than under σ' . Then compared to the outcomes in σ' , under σ ,

¹⁶ If we ignore this requirement, there is always a trivial equilibrium in which all players submit an empty list.

- i) welfare for women is weakly greater
- ii) welfare for men is weakly lower
- iii) the expected number of matches is weakly lower.

Furthermore, under both σ and σ' , i), ii) and iii) hold in comparison to the outcomes from truthful reporting of preferences to MP-DA.

The results of [Theorem 2](#) can be obtained by considering the effect of incremental truncations in light of [Proposition 1](#). An incremental truncation by a woman w can only negatively affect the welfare of men: a “divorced” man will receive a worse mate, and the chain of offers that follows can only lead to worse mates for the other men as well. Since the chain will end in an offer accepted by some woman, or else in no match, the incremental truncation weakly decreases the number of matches. This logic underpins results ii) and iii). At the same time, incremental truncation by a woman has a (weakly) positive spillover on the welfare of other women: rejection of a man can only lead to more offers for other women. The spillover from the truncation of other women, together with her best response requirement, imply that each woman weakly prefers the equilibrium with more truncation, and that any truncation equilibrium is preferred to truthful reporting.

[Theorem 2](#) is similar in spirit to [Kojima \(2006\)](#) and [Konishi and Ünver \(2006\)](#) who show that in games of capacity manipulation in hospital-intern markets, every hospital prefers a Nash equilibrium to any reported profile of larger capacities. [Theorem 2](#) also brings to mind the welfare result in [Coles et al. \(2013\)](#), in which signaling equilibria with varying cutoffs are compared. In each of these settings, actions by one side of the market – capacity reduction in [Kojima \(2006\)](#) and [Konishi and Ünver \(2006\)](#), signaling by women in [Coles et al. \(2013\)](#), and truncation in this paper – serve to “shift the balance of power.” When there are equilibria with varying degrees of action, the sides of the market are at odds over which equilibrium is preferred, and whether any action is desirable at all.¹⁷

4.2. Equilibria in uniform markets

Let a *uniform market* be the setting in which each agent is equally likely to have any full preference list. Additionally, agent utility depends on partner rank, agents identically value a match with their r th ranked choice $\forall r \in \{1, \dots, N\}$, and have identical value to being unmatched.

Theorem 3. *In uniform markets, there exists a symmetric equilibrium $((\sigma_m)^N, (\sigma_w)^N)$ where men each use the strategy σ_m of truthful reporting and women each use the strategy σ_w , which is a mixture over truncation strategies.*

Proof. We begin by constructing an auxiliary game. In this game, the set of players is the same as in the original game, the set of pure strategies for each woman is $\{0, 1, \dots, N\}$, and men all have one strategy, $\{N\}$. States of the world are profiles of preferences, which are realized with the same probability distribution as in the original game, but now players learn neither the preferences of others, nor their own preferences. Payoffs are defined according to the same utility function as in the original game, where each player receives the payoff from being matched to his stable partner under the profiles truncated at levels corresponding to the pure strategies chosen.

A standard argument due to [Nash \(1951\)](#) shows that the auxiliary game has an equilibrium, symmetric with respect to women. It is easy to see that this remains an equilibrium in the game where players observe their own preferences (but not the preferences of others) before choosing an action (truncation). Finally, returning to the unrestricted game, we recall Roth and Rothblum’s optimality of truncation theorem from [Section 2.4.1](#). Since men are playing dominant strategies, and since the strategies yield $\mathcal{M}$ -symmetric beliefs, we conclude that the profile of strategies that we found is an equilibrium in the game where strategies are unrestricted. If it were not, then some woman could do strictly better by using a non-truncation strategy. But since each woman w ’s beliefs in this setting are $\mathcal{M}$ -symmetric, a truncation strategy weakly dominates this non-mixed strategy, which yields a contradiction.¹⁸ $\square$

In addition to symmetric equilibria, asymmetric equilibria may exist. The following example illustrates this.

Example 2. Consider a 2×2 uniform market. Suppose each woman derives utility 10 from being matched to her top choice, 1 from being matched to her second choice, and 0 otherwise. We first calculate the probability of three events:

¹⁷ Another paper that bears mention is [Ashlagi and Klijn \(2012\)](#), which considers “group manipulations in truncation strategies” by women in the MP-DA. Such manipulations weakly benefit other women and harm other men. The results in [Ashlagi and Klijn \(2012\)](#) differ from ours, as we focus on equilibria and on incomplete information.

¹⁸ Using the same proof technique and following [Ehlers \(2008\)](#), we can also prove that in uniform matching markets that use an anonymous mechanism satisfying positive association, individual rationality, and independence of truncations, there exists a non-trivial equilibrium in which all men play the same truncation strategy and all women play the same truncation strategy. This class of mechanisms includes all priority mechanisms and all linear programming mechanisms introduced in British entry-level medical markets and in public school choice in some American cities.

$A = \{\text{There is a unique stable matching, and in this matching } w_1 \text{ is matched to her top choice}\}.$

$B = \{\text{There is a unique stable matching, and in this matching } w_1 \text{ is matched to her second choice}\}.$

$C = (A \cup B)^c = \{\text{There are two stable matchings}\}.$

A simple calculation shows that $P(A) = \frac{5}{8}$, $P(B) = \frac{2}{8}$ and $P(C) = \frac{1}{8}$.

Now suppose agents report preferences to *MP-DA*. If the other agents are truthful, w_1 should truncate her list to include only her most preferred man (thus earning $\frac{6}{8} \times 10 > \frac{5}{8} \times 10 + \frac{3}{8} \times 1$). But if w_1 truncates her list in this manner, w_2 has no incentive to truncate at all. Even if it turns out there was room for beneficial truncation (event C), w_1 has already done the “hard work” of truncating. She bears the risk of becoming unmatched, but also shifts the outcome from one matching to the other, improving payoffs for both women. $\square$

Several observations can be made from this example. First, payoffs are higher in this equilibrium than under truthful reporting, as predicted by [Theorem 2](#). Second, when w_1 truncates more, w_2 prefers to truncate less. While we don’t have a result that truncation under uniform preferences is a case of strategic substitutes, this example (and simulation evidence in [Fig. 3](#)) stand in contrast to the complete information world where truncation strategies are strategic complements. There, when some woman w_i truncates, this can only improve (or leave unchanged) the most preferred achievable mate for each w_j . Under complete information, this translates to a (weakly) greater optimal degree of truncation.

Another observation from the example is that w_2 ’s utility is $\frac{6}{8} \times 10 + \frac{2}{8} \times 1$, which is greater than the utility of w_1 , which is $\frac{6}{8} \times 10$. That is, the agent who truncates less has greater utility than the agent who truncates more. In uniform markets, this result generalizes: within asymmetric truncation equilibria, we have a crisp preference among women against truncation. [Theorem 4](#) encapsulates this.

Theorem 4. *Consider any asymmetric equilibrium in a uniform market where w_1 truncates more than w_2 (in the sense of first order stochastic dominance). Then*

- i) *if w_1 and w_2 swap strategies, the resulting profile will also be an equilibrium and*
- ii) *w_2 prefers the original equilibrium, in which she truncates less.*

Intuitively, w_1 and w_2 face the same opposition except for one feature: each woman “competes” with the other, but not with herself. Woman 2, who truncates less, benefits from facing competition in which the other woman truncates more. Given that w_1 is willing to take the risk of this truncation, w_2 no longer feels compelled to do so herself.¹⁹

4.3. Simulations: Finding a symmetric equilibrium

In this section we run simulations to explore equilibria in a uniform market. We assume that agents care about the rank of their partners (as in [Fig. 1](#), graphed as $N + 1$ -rank, so that the graphs display maxima rather than minima), and we examine a market with $N = 30$. The simulations suggest that under these assumptions, there is a pure strategy symmetric equilibrium with a common truncation point that involves a non-trivial degree of truncation.

We first examine how returns to truncation for w change when other women also truncate their lists. In [Fig. 2](#) we examine the effect on w ’s payoffs when women $\mathcal{W} \setminus \{w\}$ all truncate at a common point j , where j takes on various values. For lower j , curves for w are higher. This follows from the positive spillover of truncation: when $\mathcal{W} \setminus \{w\}$ truncate their lists, this benefits woman w .

In [Fig. 2](#) it is also apparent that as w ’s opponents truncate more, e.g., from $j = 30$ to $j = 25$ to $j = 15$, etc., w should truncate less: as j increases, peaks move to the right. As in [Example 2](#), this stands in contrast to the complete information result where truncation strategies are strategic complements. Note that at the extreme, when $j = 1$, w can never benefit from truncation. Since truncation still bears risk, her optimal degree of truncation in this case is $N = 30$.

Observe also that when $\mathcal{W} \setminus \{w\}$ submit very short lists, e.g., $j \in \{1, 5, 10\}$, w ’s optimal truncation point is hard to observe because her payoff curve becomes very flat. The reasons for this flatness are two-fold: When other women submit very short lists, the expected number of stable matchings is known to be small (see [Immorlica and Mahdian, 2005](#)). Hence, there are minimal opportunities for beneficial truncation. At the same time, when other women truncate, the expected partner rank for w is very low (favorable). This leaves little danger that moderate levels of truncation will leave w unmatched. The minimal rewards and risks to truncation lead to the flatness of payoff curves.

By running a very large number of iterations, we identify the peaks of the curves in [Fig. 2](#). This exercise corroborates the hypothesis that under uniform preferences, truncation is a case of strategic substitutes. As illustrated in [Fig. 3](#), the optimal truncation point for w is inversely related to j , the common truncation point of women $\mathcal{W} \setminus \{w\}$. Of course, due to the flatness of the expected payoff graph, optimal truncation points for small j are “just barely” optimal.

¹⁹ A more general version of [Theorem 4](#) also holds. For any Preference Submission Game with two “symmetric” women, i) and ii) remain valid.

Fig. 2. Woman w 's expected payoff in a uniform market as a function of her truncation point when women $\mathcal{W} \setminus \{w\}$ truncate at j . $N = 30$. Iterations = 1 000 000.

Fig. 3. Optimal truncation point of woman w in a uniform market as a function of the common truncation point j of women $\mathcal{W} \setminus \{w\}$. $N = 30$. Iterations = 10 000 000.

By overlaying the 45° line, we identify the point of truncation in a pure strategy symmetric equilibrium to be 14, more than a 50 percent truncation of the entire list.²⁰ When all women truncate at this common point, no single woman sees significant gains from truncation compared to truthful reporting. However, since truncation has a positive externality on other women, the equilibrium payoff is non-trivially greater than the payoff should all women report truthfully (see Fig. 2).

This equilibrium potentially leads to ex-post instability with respect to true preferences. Some women might over-truncate and be left unmatched. More subtly, when *other* women truncate, woman w may be paired with a mate that is not achievable under the true preferences.

These instabilities suggest a possible application of these results: the impact of strategic behavior on a post-market “scramble” for positions. Since truncation can lead to unmatched participants following the match, a second, organized match might be helpful for these agents. Indeed, post-market scrambles have been organized in both the market for medical residents as well as in the job market for new economists.²¹

²⁰ To test whether this is indeed a symmetric equilibrium, we “guess and verify.” We repeatedly sample w 's payoffs under (k, j) where k is w 's truncation point and j is the common truncation point of $\mathcal{W} \setminus \{w\}$, and establish that w 's payoffs under $(14, 14)$ are sufficiently distant from those under $(15, 14)$, $(13, 14)$, and other profiles. To do this, we construct confidence intervals around $(14, 14)$ and, using different draws, around other profiles. We observe that these intervals have an empty intersection.

²¹ The NRMP offers the “Supplemental Offer and Acceptance Program” (SOAP), which replaced a somewhat less orderly scramble (see <http://www.nrmp.org/2012springmeeting.pdf>). The American Economic Association organizes the “Scramble” in which candidates seeking jobs and employers with positions open late in the job market can announce their availability on the AEA website (see <http://www.aeaweb.org/joe/scramble>).

At first observation, an organized scramble would reduce the downside to remaining unmatched in the primary match. But this might induce additional risk-taking behavior (more truncation) by participants. Such behavior would increase the pool size in the second match, raising the value of being unmatched, inducing even more truncation. A secondary match might ultimately enjoy high participation levels, but only because it has drawn participants away from the primary match, complicating overall welfare analysis.

5. Truncation and risk aversion

While gains to truncation can be significant, truncation is nevertheless a risky strategy. When w 's opponents truncate, truncation for w is particularly risky: compared to truthful reporting, optimal truncation offers minimal benefit, and over-truncating can lead to large losses. One might expect agents with more conservative attitudes toward risk to shy away from this proposition. In this section, we ask how a woman's truncation behavior varies as we vary her attitude towards risk.

We consider a general setting, with arbitrary preferences for woman w and beliefs about reported preferences of others. Let $\psi(\cdot)$ be any strictly increasing, concave transformation. We will show that for any beliefs about others, woman w with preferences $u_w(\cdot)$ will truncate more than a woman w_ψ who has identical beliefs, but preferences given by $\psi(u_w(\cdot))$.

Recall that when we fix w 's preferences to be $u_w(\cdot)$, we defined shorthand

$$v(k, P_{-w}) \equiv u_w(\mu^M[P_w^k, P_{-w}](w)),$$

her payoff from submitting truncated preference list P_w^k . Now define

$$v_\psi(k, P_{-w}) \equiv \psi(u_w(\mu^M[P_w^k, P_{-w}](w))),$$

the payoff from submitting truncated preference list P_w^k for a woman w_ψ with preferences $\psi(u_w(\cdot))$.

The following theorem states that if w prefers truncating less to more, then w_ψ definitely prefers truncating less to more.

Theorem 5. Let $\tilde{P}_{-w}$ be any random variable distributed over $\mathcal{P}_{-w}$. Then $\forall k \in \{1, \dots, N-1\}, \forall t \in \{1, \dots, N-k\}$ we have

$$\mathbb{E}[v(k, \tilde{P}_{-w})] \leq \mathbb{E}[v(k+t, \tilde{P}_{-w})] \Rightarrow \mathbb{E}[v_\psi(k, \tilde{P}_{-w})] \leq \mathbb{E}[v_\psi(k+t, \tilde{P}_{-w})].$$

Furthermore, if

- i) $\psi(\cdot)$ is strictly concave, and
- ii) under $\tilde{P}_{-w}$, each man is achievable for w with positive probability,

then the second inequality is strict.

The constructive proof nicely illustrates incremental truncation analysis, so we provide it in-text.

Proof. We begin with the proof for $t = 1$. An analogous argument works for all other t in the given range, with necessary proof adjustments described at the end. Our technique focuses on two lotteries over outcomes. Let Q^{k+1} be the lottery over mates for w when she truncates at $k+1$, and let Q^k be the lottery when she truncates at k . Our goal is now to show that if Q^k is mean-decreasing as compared to Q^{k+1} from w 's perspective (in terms of her von Neumann–Morgenstern utility), then it will be mean-decreasing from w_ψ 's perspective as well.

Distributions Q^{k+1} and Q^k are shown in Fig. 4. Recalling Proposition 1, k -truncating is equivalent to $(k+1)$ -truncating followed by k -truncating. That is, lottery Q^k is equivalent to starting with lottery Q^{k+1} , then rolling the die again if w receives her $(k+1)$ ranked choice. Hence,

$$q_i^k \geq q_i^{k+1} \quad \forall i \in \{1, \dots, k\} \cup \{w\}.$$

Let shorthand $u_i(Q)$, $i \in \{w, w_\psi\}$ describe i 's expected utility from lottery Q . Suppose first that $u_w(Q^k) = u_w(Q^{k+1})$, that is, $\mathbb{E}[v(k, \tilde{P}_{-w})] = \mathbb{E}[v(k+1, \tilde{P}_{-w})]$, so that from w 's perspective, Q^k is a *mean-preserving spread* of Q^{k+1} . Then by Jensen's inequality, $u_{w_\psi}(Q^k) \leq u_{w_\psi}(Q^{k+1})$. If $\psi(\cdot)$ is strictly concave and $Q^k \neq Q^{k+1}$ (which follows from ii)), then $u_{w_\psi}(Q^k) < u_{w_\psi}(Q^{k+1})$.

Now suppose that $u_w(Q^k) < u_w(Q^{k+1})$, so that from w 's perspective, Q^k is *mean-decreasing* as compared to Q^{k+1} . We will now construct an intermediate lottery Q' such that

1. Q^{k+1} P_w -stochastically dominates Q' and
2. From w 's perspective, Q^k is a mean preserving spread of Q' .

Fig. 4. k -Truncation is equivalent to $(k + 1)$ -truncation followed by an extra gamble: divorcing man $k + 1$.

Define lottery Q' so that Q' is identical to Q^{k+1} , except that we replace outcome $k + 1$ (w 's $(k + 1)$ ranked choice) with lottery $\alpha(k + 1) + (1 - \alpha)w$. Choose $\alpha \in [0, 1]$ such that w has $u_w(Q') = u_w(Q^k)$. Such an α must exist: when $\alpha = 1$, $Q' = Q^{k+1}$, and when $\alpha = 0$, $u_w(Q') \leq u_w(Q^k)$. Our desired α follows from the Intermediate Value Theorem.

By construction, Q^{k+1} P_w -stochastically dominates Q' . With respect to w 's utility, we also have that Q^k is second order stochastically dominated by Q' . To see this, observe that Q' was constructed to have the same mean as Q^k , and that compared to Q' , Q^k shifts probability mass to the extremes: $q_j^k \geq q_j'$ for $j \in \{1, \dots, k\} \cup \{w\}$.

Since $P_w = P_{w_\psi}$, by P_w -stochastic dominance, w_ψ also strictly prefers Q^{k+1} to Q' . By Jensen's inequality, w_ψ weakly prefers Q' to Q^k . Hence,

$$u_{w_\psi}(Q^k) < u_{w_\psi}(Q^{k+1}),$$

so the theorem is proved for $t = 1$.

When $t > 1$, we may again construct an intermediate lottery Q' , this time that transfers weight from $\{k + 1, k + 2, \dots, k + t\}$ to the unmatched option w . Just as before, we can construct Q' to ensure that Q^{k+1} P_w -stochastically dominates Q' , and that from w 's perspective, Q^k is a mean preserving spread of Q' . The key insight is that truncation transfers probability mass to the extremes: the most preferred mates, as well as the unmatched option. $\square$

We can now use Theorem 5 to sort optimal truncation points based on degree of concavity.

Corollary 1. Let k_i^l be the minimum optimal truncation point (by rank) and let k_i^h be the maximum optimal truncation point for woman $i \in \{w, w_\psi\}$. Then $k_w^l \leq k_{w_\psi}^l$ and $k_w^h \leq k_{w_\psi}^h$. Furthermore, if conditions i) and ii) from Theorem 5 hold, then $k_w^h \leq k_{w_\psi}^l$.

Proof. If k_w^l is w 's minimum optimal truncation point, then w strictly prefers truncation at k_w^l to truncation at any $k < k_w^l$. Following the reasoning of Theorem 5, w_ψ must then prefer truncation at k_w^l to truncation at any $k < k_w^l$. Hence, $k_w^l \leq k_{w_\psi}^l$. A similar argument can be used to show $k_w^h \leq k_{w_\psi}^h$.

If k_w^h is w 's maximum optimal truncation point, then w (weakly) prefers truncation at k_w^h to truncation at any $k < k_w^h$. If conditions i) and ii) hold, then w_ψ must strictly prefer truncation at k_w^h to truncation at any $k < k_w^h$. Hence, $k_w^h \leq k_{w_\psi}^l$. $\square$

Thus, when facing the same environment, players who are more risk averse truncate less, with the set of optimal truncation points overlapping at the very most at one point.

The key insight in the analysis is the interpretation of truncation as a risky lottery, and then mapping the additional risk associated with incremental truncation to an *extra lottery* a woman must face. If a woman doesn't like to face the extra lottery, then certainly a woman with more concave preferences will not want to face it.

Note that despite pertaining to risk aversion, the results in this section do not restrict the structure of $u_w(\cdot)$ in any way. For example, we do not require $u_w(\cdot)$ to be "concave." Rather, it is the *relative* concavity that is crucial. For example, if we restrict ourselves to the class of functions that are s -shaped in rank, we know that within this class, concave transformations induce less truncation.

In a general sense, this result can be taken as advice to participants. Players can observe the patterns of behavior of others, size up their own attitudes toward risk, and truncate *more* or less accordingly. In markets where there is a steep dropoff in utility from a woman's most preferred partner to her second choice, a smaller dropoff from choice two to three, and so forth, we may anticipate more aggressive truncation. On the other hand, if participants are largely content with any of the available choices, but see great disutility from being unmatched, truncation is not advisable.

This result can also offer advice to a market designer. If an objective is to maximize the number of matches, a market designer may wish to choose the less risk averse side to be the "proposers" in the Deferred Acceptance Algorithm. If the two sides of the market are identical in all regards except for their risk preferences, the more risk averse side will be less

likely to truncate, even if manipulations increase their expected partner rank. Lower levels of truncation will increase the number of realized matches, and consequently, reduce the number of participants left unmatched. However, in making this choice, the market designer should take other market features into consideration as well, as we demonstrate in the next section.

6. Correlated preferences

In the Preference List Submission Problem for women, we now let woman w believe other women in the market have preferences similar to hers. We consider how woman w should vary her degree of truncation as the degree of similarity varies. We provide evidence, both theoretical and simulation-based, that the greater the similarity in the preferences of other women to her own, the less woman w should truncate.

6.1. Perfectly correlated preferences

We consider first the case of perfectly correlated preferences on the women's side of the market.

Remark. When women have identical preferences, there is a unique stable matching.

To see this, note that the top-ranked man, as agreed upon by all women, must be matched with his most preferred partner in any stable matching, or else these two would constitute a blocking pair. The second-ranked man must then be matched to his most preferred remaining woman, and so on. *MP-DA* reduces to a serial dictatorship, determined by the common ranking of the men.

Since there is a unique stable matching in this setting, an individual woman's misrepresentation of her preference list can never improve her match. In fact, if a woman is certain that other women share her preferences (and are reporting truthfully), but is uncertain about what men will submit to the algorithm, truncation can very well lower her outcome by leaving her unmatched.

6.2. Partially correlated preferences

In this section, we introduce a notion of partial correlation of preferences indexed by a single parameter α . We will show that the greater the degree of correlation, the less a woman should truncate.

Consider the Preference List Submission Problem for woman w with preferences u_w and beliefs $\tilde{P}_{-w}$ about reported preference lists of opponents. Let $p(\cdot, \cdot)$ be the probability mass function for w 's beliefs. That is,

$$p(P_{\mathcal{M}}, P_{\mathcal{W} \setminus \{w\}})$$

gives the likelihood that the men will report preference lists $P_{\mathcal{M}}$ and women $\mathcal{W} \setminus \{w\}$ will report preference lists $P_{\mathcal{W} \setminus \{w\}}$. Define the marginal probability over men's preference profiles by $p^M(\cdot)$.

Given $p(\cdot, \cdot)$, define beliefs $p^C(\cdot, \cdot)$ by

$$p^C(P_{\mathcal{M}}, P_{\mathcal{W} \setminus \{w\}}) \equiv \begin{cases} p^M(P_{\mathcal{M}}) & \text{if } P_{\hat{w}} = P_w \forall \hat{w} \in \mathcal{W} \setminus \{w\} \\ 0 & \text{otherwise.} \end{cases}$$

$p^C(\cdot, \cdot)$ is the distribution that preserves the marginal distribution over men's preferences $p^M(\cdot)$, but where the other women share the preferences of w .

Define beliefs $p^\alpha(\cdot)$ by

$$p^\alpha(P_{-w}) \equiv (1 - \alpha)p(P_{-w}) + \alpha p^C(P_{-w}).$$

Hence, as α varies from 0 to 1, p^α ranges from p to p^C . The marginal distribution over men's preferences remains fixed, while the correlation of women's preferences steadily increases (the distribution remains constant if $p = p^C$).

The set of optimal truncation points for woman w with preferences u_w and beliefs indexed by α is given by

$$k^*(\alpha, p, u_w) \equiv \arg \max_{k \in \{0, \dots, N\}} \mathbb{E}_{p^\alpha} [v(k, \tilde{P}_{-w})].$$

Notice that since the choice set is finite, $k^*(\cdot, \cdot)$ will be non-empty.

Let $k^h(\alpha, p, u_w) = \max[k^*(\alpha, p, u_w)]$ and $k^l(\alpha, p, u_w) = \min[k^*(\alpha, p, u_w)]$, the optimal choices involving the least and most truncation respectively.

The following proposition states that for any preferences u_w and beliefs p , as we increase the degree of correlation α , woman w should truncate less.

Proposition 2. Let $\alpha, \alpha' \in [0, 1]$ with $\alpha' > \alpha$. Then $k^l(\alpha', p, u_w) \geq k^l(\alpha, p, u_w)$ and $k^h(\alpha', p, u_w) \geq k^h(\alpha, p, u_w)$.

Fig. 5. The graphs display $(N + 1) -$ an individual woman's expected partner rank from truncating her list at each point $k \in \{0, \dots, N\}$ and submitting these preferences to MP-DA. Preference lists of men are uniformly random, and lists for women are randomly generated using the procedure described in the text. Payoffs are averaged over 100 000 draws.

The proof relies on the fact that when there is a unique stable matching, it can never hurt to submit a full list. Using this fact, we can show that if under low correlation, w prefers truncating less to more, then under high correlation w definitely prefers truncating less to more. This is enough to sort optimal truncation points.

Intuition for this result is related to the size of the set of stable matchings. Truncation can yield improvement only when there are multiple stable matchings. The greater the degree of correlation, the smaller this set, and the lower the likelihood that a window for gain from truncation exists.

The anticipated level of correlation in the environment might influence the advice a market designer offers participants. If correlation is high, the designer can safely advise participants to report truthfully, and it is in their best interest to do so. With low correlation (sufficiently heterogeneous preferences), players may anticipate gains from truncation, which if acted on, could lead to unstable matchings.

6.3. Noisy preferences

In Section 6.2, a woman believes it is possible that opponents have preference lists *identical* to hers. In this section, woman w believes women have preference lists *similar* to hers, but not necessarily identical. We model such beliefs for women by generating noisy deviations from a common preference list. By performing simulations, we corroborate the theoretical results in Section 6.2; more correlation means a woman should truncate less.

We generate correlated preferences as follows. Each man m_i is assigned a random number $r_i \sim U[0, 1]$, and this value is agreed upon by all women. For each man m_i , each woman w_j also assigns an idiosyncratic (noise) component, $q_{ij} \sim U[0, 1]$. Woman w_j 's rankings over men are then determined by the sum $\alpha \cdot r_i + (1 - \alpha)q_{ij}$, where $\alpha \in [0, 1]$ is a parameter that we will vary. Observe that from the perspective of any woman w , the preferences of other women are noisy versions of her own rankings. Values of α close to one imply low noise, so α measures the degree of correlation. Men are assumed to have uniformly random rankings over the women.²²

The process just described is used only to determine preference orderings. We further assume that w 's payoff is given by $(N + 1 - \text{partner rank})$, and being unmatched is just worse than being matched to her least preferred man, so we can compare outcomes to those depicted in Fig. 1.

Fig. 5 graphs the return to truncation for various values of α . For each value, we randomly generated 100 000 preference list profiles and for each k , we graph woman w 's average payoff from k -truncation, when other agents are truthful.

When $\alpha = 0$ (the top curve), this corresponds to uniform beliefs for w , the case studied in Section 3.4. When $\alpha = 1$, all women rank men the same way, the stable matching will be unique, and truncation cannot be helpful (as in Section 6.1).

From Fig. 5, we make two key observations. First, woman w dislikes correlation. This fact is easy to explain. If all women agree on who the top men are, they "compete" for them as mates. The lower the correlation, the less the competition, and the better the expected mate for w . Second, w 's optimal truncation point increases as correlation increases. This corroborates the result in Section 6.2: when there is more correlation, w should truncate less.

²² The common starting point for preferences might be an aggregate ranking based on available data, like the US News and World Report's annual ranking of universities. Caldarelli and Capocci (2000) simulate preferences in a one-to-one model similarly. In their model, the common component r_i is a man's "beauty," which in their view, evidently, is not in the eye of the beholder.

7. Discussion and conclusion

In this paper, we study optimal strategic behavior in one-to-one matching markets that are based on the Deferred Acceptance Algorithm, when agents have incomplete information about the preferences of others. We focus on truncation strategies. Among classes of strategies for preference list misrepresentation, truncation is an attractive option because it is guaranteed to weakly increase the likelihood of matching with one's more preferred partners. By contrast, more complicated strategies, such as swapping the order of agents in a preference list, may require detailed information about the preferences reported by others, and their outcomes are more difficult to predict.

Recent work by Immorlica and Mahdian (2005), Kojima and Pathak (2009) and others demonstrate that in large markets where agents submit short preference lists, opportunities for manipulation are limited. Lee (2011) presents a random utility model and shows that, in some sense, gains from manipulation become small in large markets.²³ In light of these findings one may ask whether agents – especially agents with little detailed information – can ever substantially gain from manipulation. Our paper answers in the affirmative. When agents view reported preference lists of others as being drawn uniformly from the set of all possible full length preference lists, they may truncate their lists with little risk of being unmatched, but with the potential to see large gains in terms of the expected partner rank. Importantly, we show that while according to Lee (2011), utility gain from manipulation may be small, the optimal truncation may still be substantial. This finding provides an essential qualification to his results.

For many of the settings in which the Deferred Acceptance Algorithm has been successfully applied, notably in the NRMP and in the Boston and New York school systems, the markets do reflect large numbers and short preference lists. But the high levels of optimal truncation demonstrated in this paper raise a key issue: in large markets where agents submit short preference lists, can we be sure that the short lists were not simply the result of optimization? Costliness of information discovery often places natural limits on the length of submitted preference lists. Flyouts are costly for medical students; perhaps somewhat less so for hospitals. Nevertheless, this paper illustrates the theoretical possibility that even with full information about one's own preferences, substantial truncation (submission of short lists) may simply be utility-maximizing strategic behavior.

Appendix A

A.1. Proofs

Proof of Remark 1. Observe first that any matching that is stable with respect to (P_w^k, P_{-w}) and matches woman w to a man must be stable with respect to P , and that any matching $\tilde{\mu}$ that is stable with respect to P with $\tilde{\mu}(w)$ ranked $\leq k$ must be stable with respect to (P_w^k, P_{-w}) . Hence, setting

$$\mathcal{M}_1 = \{m \in \mathcal{M} \mid m \text{ achievable for } w \text{ under } (P_w^k, P_{-w})\}$$

and

$$\mathcal{M}_2 = \{m \in \mathcal{M} \mid m \text{ achievable for } w \text{ under } P \text{ and } m \text{ ranked } \leq k \text{ in } w\text{'s list}\}$$

we have $\mathcal{M}_1 = \mathcal{M}_2$. By the Gale–Shapley result, $\mu^M[P_w^k, P_{-w}](w)$ is w 's least preferred element of $\mathcal{M}_1$, and hence of $\mathcal{M}_2$. Should both sets be empty, then $\mu^M[P_w^k, P_{-w}](w) = w$. $\square$

Proof of Proposition 1. We introduce Algorithm 1 below and prove that given the same input, Algorithm 1 and the Divorcing Algorithm generate the same output, which in each case is the MP-DA outcome described in the statement of the proposition.

Like the Divorcing Algorithm, Algorithm 1 takes as its input a profile P of preference lists, a woman w , and a truncation point $k \in \{1, \dots, N\}$, and outputs a matching. Algorithm 1 is adapted from an algorithm due to McVitie and Wilson, which differs from MP-DA in that the men make offers one at a time instead of in rounds, but is nevertheless outcome equivalent (McVitie and Wilson, 1970). Algorithm 1 is identical to McVitie and Wilson's, except that we explicitly delay selecting man $\mu^M[P]$ until absolutely necessary. By McVitie and Wilson (1970), the algorithm plainly produces $\mu^M[P_w^k, P_{-w}]$, the MP-DA outcome when w k -truncates her preference list. $\square$

Algorithm 1.

- **Step 0.** Initialization. Identify the least preferred achievable mate for woman w under (P_w, P_{-w}) and call this man m_l . For example, we may identify this man by running MP-DA, setting $m_l = \mu^M[P_w, P_{-w}](w)$.
Iteration over steps 1 and 2. Preferences in these steps are given by (P_w^k, P_{-w}) .

²³ Lee's model is more general than ours, in many respects. Note, however that his model requires agent utilities to be bounded, while our model does not exclude unbounded functions.

- **Step 1.** Pick any single man other than m_l who has not exhausted his preference list. If no such man exists, pick m_l . If we have picked m_l , and m_l is not single, or if m_l has exhausted his preference list, terminate.
- **Step 2.** The man chosen in the previous step makes an offer to the most preferred woman on his preference list who has not already rejected him. If this woman finds the man acceptable and prefers him to her current mate (or if she is single), she holds his offer and divorces her previous mate (if necessary). Return to step 1.

Let $\mu^1(P, k, w)$ be the output of Algorithm 1 and recall that $\mu^{DIV}(P, k, w)$ is the output of the divorcing algorithm.

To establish outcome equivalence of the algorithms, begin by letting l be the rank of w 's least preferred achievable mate m_l under P .

- If $k \geq l$, both algorithms clearly produce μ^M , the men-optimal matching under P .
- If $k < l$, then the algorithms will reach a point where they coincide. That is, there will be a point where the sequences of single men chosen coincide, as do the temporary matchings and preference lists.

In Algorithm 1, we claim that (1) at some point, m_l will make an offer to w , which will be rejected. (2) From this point forward, the algorithm coincides with the Divorcing Algorithm, just after its initialization step.

1. Under MP-DA, when w k -truncates her list, men are (weakly) worse off than if she reports truthfully (see Gale and Sotomayor, 1985). This means that in $\mu^M(P_w^k, P_{-w})$, m_l must be matched with a candidate worse than w , or possibly with no woman at all. Hence, in Algorithm 1, m_l must have made an offer to w (since he makes offers from his list in order of preference), and this offer must have been rejected.

2. When in Algorithm 1, m_l makes his offer to w , no better ranked man has yet done so. Otherwise, let m' be the first man ranked higher than m_l to make an offer to w and backtrack to the point in the algorithm where this offer is made. Note that up to this point, the path of the algorithm is consistent with w having l -truncated her preferences, since she has not faced any man ranked k through l . But this implies that if w l -truncated her list, she would receive a mate at least as good as m' , not m_l . This contradicts Proposition 1.

By the choice-of-proposer rule in the algorithm, we know that when m_l proposes to w , he must be the only single man who has not yet exhausted his list. If w accepted m_l 's offer, the path of the algorithm would be consistent with w having l -truncated her list, and the algorithm would terminate with matching μ^M . Hence, by instead rejecting m_l , we arrive at exactly the position of the Divorcing Algorithm, following step 0.

Thereafter, the algorithms coincide, thus yielding identical outcomes. $\square$

Proof of the lemma in Section 3.3. For each i and k , define

$$\mathcal{P}_2^i(k) \equiv \{P_{-w} \mid P_{-w} \in \mathcal{P}_2(k), \mu^M(P_w^k, P_{-w})(w) = m_i\}.$$

We wish to show that w finds $\mathcal{P}_2^i(k)$ and $\mathcal{P}_2^j(k)$ equally probable, for all k and all $i, j \leq k$. We proceed by finding a bijection from $\mathcal{P}_2^i(k)$ to $\mathcal{P}_2^j(k)$ which is probability preserving with respect to w 's beliefs.

For $i, j \in \{1, \dots, k\}$, we define a mapping $f_{ijk}: \mathcal{P}_2(k) \rightarrow \mathcal{P}_2(k)$. Let $f_{ijk}(P_{-w}) \equiv P'_{-w}$ be given by the following:

1. Switch m_i and m_j everywhere. Switch the positions of m_i and m_j in each woman's list, and swap m_i and m_j 's preference lists (this is like relabeling).
2. Switch back m_i and m_j in w 's list.

Notice that this is equivalent to swapping m_i and m_j in w 's list only, and then relabeling i and j .²⁴

Suppose $P_{-w} \in \mathcal{P}_2^i(k)$. The fact that w finds P_{-w} and P'_{-w} equally probable follows directly from the definition of $\mathcal{M}$ -symmetry. We will show that $P'_{-w} \in \mathcal{P}_2^j(k)$. Note that it is not immediately clear that we even have $P'_{-w} \in \mathcal{P}_2(k)$, that is, that under P'_{-w} , k -truncation still yields an improvement for w .

We think of the matching as arising from MP-DA. Since $P_{-w} \in \mathcal{P}_2(k)$, if w does not truncate, she will be matched with a man worse than m_k . Hence, during the process of the algorithm, she will not receive an offer from any man $m_1, \dots, m_k$. Hence, rearranging these men in w 's list will not affect the outcome, and in particular, swapping m_i and m_j will not affect the outcome (the stable matching). Furthermore, since $P_{-w} \in \mathcal{P}_2^i(k)$ we know that under P_{-w} , k -truncation leaves w matched with m_i . Using Proposition 1, we know that during the chain of proposals following an "ex-post" k -truncation by w , the first man to make an offer to w will be m_i . Hence, this will still be true if w swaps the position of m_i and m_j in her list.

Thus, we have that if w switches m_i and m_j in her list, k -truncation will yield an improvement and she will again be matched with m_i . But now relabeling m_i and m_j (so that w 's list is $(m_1, m_2, m_3, \dots)$), we have that $P'_{-w} \in \mathcal{P}_2^j(k)$.

Hence, $f_{ijk}(\cdot)$ is a bijection from $\mathcal{P}_2^i(k)$ to $\mathcal{P}_2^j(k)$, which is probability preserving with respect to w 's beliefs. This is sufficient to prove the proposition. $\square$

²⁴ For the trivial case, $i = j$, we use the identity mapping.

To prove [Theorem 1](#), we begin with a lemma demonstrating that even upon submitting a vanishingly small truncation of one's list (relative to the length of one's full preference list), we still see gains relative to truthful reporting. We examine the case where a woman's payoff is given by her partner rank, and being unmatched is treated as rank $N + 1$. At the end of the proof, we show that the result also holds for the more general preferences described in the statement of [Theorem 1](#).²⁵

Lemma 1. *There exists N^* such that for every $N > N^*$, the gain to woman w from truncating at $7 \log^2 N$ relative to truthful reporting is strictly greater than zero. Furthermore, the expected rank of w 's mate is lower than (better than) her expected mate rank from truthful reporting by at least $\frac{1+N}{2+\log N} - 7 \log^2 N - 2$.*

Proof. First, in the case of no truncation, we know from [Pittel \(1989, p. 545\)](#) that the expected rank of w 's husband, $R_w(\mu_m)$ satisfies:

$$\mathbb{E}[R_w] \geq \frac{1+N}{1+H_N} \geq \frac{1+N}{2+\log N},$$

where $H_N = \frac{1}{1} + \frac{1}{2} + \dots + \frac{1}{N}$, the N th harmonic number. Let D be the highest (worst) rank some woman gets under WP-DA when all agents report their preferences truthfully. Using [Theorem 6.1](#) from [Pittel \(1992\)](#), we observe that for N large enough,

$$\Pr(D \leq 7 \log^2 N) \geq 1 - \frac{1}{N}.$$

Therefore, truncating at $7 \log^2 N$ ensures an expected rank of at most $7 \log^2 N \times (1 - \frac{1}{N}) + \frac{1}{N} \times (N + 1)$. Hence, the expected gain (in rank) from truncation, Δ , satisfies:

$$\Delta \geq \frac{1+N}{2+\log N} - 7 \log^2 N - 2.$$

The right hand side approaches infinity as N grows to infinity, so for N large enough, $\Delta > 0$. $\square$

In [Lemma 1](#) we have established that truncating at $7 \log^2 N$ ensures a gain (in terms of expected partner rank) relative to truthful revelation that grows arbitrarily large as $N \rightarrow \infty$. Note that this gain is an absolute measure. As measured as a fraction of the expected payoff from truthful revelation, the gains from truncation go to zero.²⁶

It remains to establish that as a fraction of the market size N , the degree of optimal truncation goes to 0. To do this, we will first show that any truncation of a constant fraction of one's list is (asymptotically) outperformed by the level of truncation found in [Lemma 1](#).

Lemma 2. *For any fraction $\alpha \in (0, 1)$,*

- i) *there exists $N(\alpha)$ such that for every $N > N(\alpha)$, a woman's payoff from truncating at αN is lower than that from truncating at $7 \log^2 N$;*
- ii) *there exists $N^*(\alpha)$ such that for every $N > N^*(\alpha)$, a woman's payoff from truncating at xN is lower than that from truncating at $7 \log^2 N$ for every $x \in [\alpha, 1]$.*

Proof. We begin by proving i) for the case of $\alpha = 1 - \frac{1}{e}$.

Let $\bar{\Delta}$ be the expected difference between the rank of the mate under truncation at αN and the rank of the mate under truthful revelation. Let $\epsilon > 0$ be a small number. Let A_N be the event $\{w \text{ gets fewer than } (1 - \epsilon) \log N \text{ offers, or else more than } (1 + \epsilon) \log N, \text{ before MP-DA stops}\}$. Let $P_N = P(A_N)$. We then have:

$$\begin{aligned} \bar{\Delta} &\leq P_N \times N + (1 - P_N) \times \Pr\left\{\text{Rank}(\mu^M(w)) > \left(1 - \frac{1}{e}\right)N \mid \neg A_N\right\} \times N \\ &\leq P_N \times N + \Pr\left\{\text{Rank}(\mu^M(w)) > \left(1 - \frac{1}{e}\right)N \mid \neg A_N\right\} \times N \\ &\leq P_N \times N + \left(\frac{1}{e}\right)^{(1-\epsilon) \log N} \times N. \end{aligned} \tag{A.1.1}$$

Note that truncation may only matter in the event $\{\text{Rank}(\mu_M(w)) > (1 - \frac{1}{e})N\}$, which is included in the event $B = \{\{\text{Rank}(\mu_M(w)) > (1 - \frac{1}{e})N\} \cap \neg A_N\} \cup A_N$. In the first inequality, we have replaced the conditional benefits from truncation,

²⁵ Throughout the proofs of [Lemmas 1, 2](#), and [Theorem 1](#), for any fractional $x \in \mathbb{R}^+$, we treat x -truncation as $\lfloor x \rfloor$ -truncation.

²⁶ Recall that as a fraction of N , the expected partner rank for women (as well as for men) converges to 0 as N grows large.

be they positive or negative, with N , and considered the event B . In the last inequality, we treat offers w receives as independent draws (and invoke the Principle of Deferred Decisions), when in fact the draws are “without replacement”, which would yield a lower probability. These substitutions are all acceptable as we are finding an upper bound on $\bar{\Delta}$.

Using Eq. (2.16) from [Pittel et al. \(2007\)](#), we know that there exists some $c > 0$ such that for N large enough, $P_N \leq \exp(-c \cdot \log^{\frac{1}{3}} N)$. Hence,

$$\begin{aligned}\bar{\Delta} &\leq \exp(-c \cdot \log^{\frac{1}{3}} N) \times N + \left(\frac{1}{e}\right)^{(1-\epsilon) \log N} \times N \\ &= \frac{N}{\exp(c \cdot \log^{\frac{1}{3}} N)} + N^\epsilon.\end{aligned}\tag{A.1.2}$$

We now must show that for large N ,

$$\frac{N}{\exp(c \cdot \log^{\frac{1}{3}} N)} + N^\epsilon \leq \frac{1+N}{2+\log N} - 7\log^2 N - 2,$$

which by [Lemma 1](#) will imply $\Delta \geq \bar{\Delta}$.

We have

$$\frac{N}{\exp(c \cdot \log^{\frac{1}{3}} N)} + N^\epsilon \leq \frac{2N}{\exp(c \cdot \log^{\frac{1}{3}} N)}$$

for N large enough, since

$$N^\epsilon \leq \frac{N}{\exp(c \cdot \log^{\frac{1}{3}} N)} \iff 1 \leq \frac{N^{1-\epsilon}}{\exp(c \cdot \log^{\frac{1}{3}} N)} \iff c \cdot \log^{\frac{1}{3}} N \leq (1-\epsilon) \log N,$$

which clearly holds for large N .

Hence, it is sufficient to prove that:

$$\frac{2N}{\exp(c \cdot \log^{\frac{1}{3}} N)} \leq \frac{1+N}{2+\log N} - 7\log^2 N - 2.$$

Since for large N , $7\log^2 N + 2 < N^\epsilon$, it suffices to show that $\frac{3N}{\exp(c \cdot \log^{\frac{1}{3}} N)} \leq \frac{1+N}{2+\log N}$, which is implied by

$$\frac{3N+3}{\exp(c \cdot \log^{\frac{1}{3}} N)} \leq \frac{1+N}{2+\log N} \iff 3 \leq \frac{\exp(c \cdot \log^{\frac{1}{3}} N)}{2+\log N}.$$

Observe that

$$\lim_{N \rightarrow \infty} \frac{\exp(c \cdot \log^{\frac{1}{3}} N)}{2+\log N} = \lim_{x \rightarrow \infty} \frac{\exp(c \cdot x)}{2+x^3} = \infty$$

since c is greater than 0. This completes the proof for the case of $(1 - \frac{1}{e})N$.

To show that i) holds, we now consider general $\alpha \in (0, 1)$. Let $r \equiv \frac{1}{1-\alpha} > 1$, so that $1 - \frac{1}{r} = \alpha$. An analogous proof holds with truncation at $(1 - \frac{1}{r})N$. Probability P_N will remain unchanged, and in Eq. (A.1.1), instead of $(\frac{1}{e})^{(1-\epsilon) \log N}$ we have $(\frac{1}{r})^{(1-\epsilon) \log N} = (\frac{1}{e})^{\log r \cdot (1-\epsilon) \log N} = (\frac{1}{N})^{(1-\epsilon) \log r} = (\frac{1}{N})^{\delta(\alpha)}$, where $\delta(\alpha) \equiv (1-\epsilon) \log r = (1-\epsilon) \log \frac{1}{1-\alpha} > 0$. We may then replace N^ϵ with $N^{1-\delta(\alpha)}$ in Eq. (A.1.2), and the remaining argument will hold.

To show that ii) holds, observe that the critical appearance of α is in the inequality $N^{1-\delta(\alpha)} \leq \frac{N}{\exp(c \cdot \log^{\frac{1}{3}} N)}$. For every $x > \alpha$, we have $N^{1-\delta(x)} \leq N^{1-\delta(\alpha)}$. Hence, for any N large enough so that $N^{1-\delta(\alpha)} \leq \frac{N}{\exp(c \cdot \log^{\frac{1}{3}} N)}$, we have that $N^{1-\delta(x)} \leq \frac{N}{\exp(c \cdot \log^{\frac{1}{3}} N)}$ holds as well, demonstrating ii). $\square$

Proof of Theorem 1. By way of contradiction, assume that $\lim_{N \rightarrow \infty} \frac{k^*(N)}{N} = 0$ does not hold. This implies that there exists a subsequence $\{N_j\}$ such that $\lim_{j \rightarrow \infty} \frac{k^*(N_j)}{N_j} = b > 0$, so for N_j large enough, $\frac{k^*(N_j)}{N_j} > b/2$. By [Lemma 2ii\)](#), we know that for large enough N_j , truncating at $7\log^2 N_j$ outperforms truncating at xN_j for any $x \geq b/2$. But this contradicts the optimality of the truncations at $k^*(N_j)$, and so concludes the proof for the case when payoffs are given by partner rank. By applying [Corollary 1](#), we see that the result also holds for any strictly increasing, convex transformation of such preferences. $\square$

To prove [Theorems 2 and 4](#), we show that the following lemma holds:

Lemma 3. Let τ be a profile of strategies where each man reports truthfully and women play truncation strategies. Let σ^* be an equilibrium in truncation strategies, such that every woman in $\mathcal{W} \setminus \{w\}$ truncates more at any state of the world (in the sense of FOSD) and men report truthfully. Then woman w is weakly better off under σ^* than under τ .

Proof. Since other women truncate more under σ^* , it is clear that the payoff to w from the profile (τ_w, σ_{-w}^*) is weakly higher than her payoff under τ . Moreover, since σ_w^* is a best response to σ_{-w}^* , the payoff to w from $(\sigma_w^*, \sigma_{-w}^*)$ is weakly greater than that under (τ_w, σ_{-w}^*) . $\square$

Proof of Theorem 2. i) is a direct consequence of Lemma 3. Proofs for ii) and iii) were given in-text. $\square$

Proof of Theorem 4. i) follows from symmetry. ii) is a direct consequence of Lemma 3. $\square$

Proof of Proposition 2. To prove the proposition, we first show that if under low correlation, we prefer truncating less to more, than under high correlation we definitely prefer truncating less to more.

First, observe that

$$\begin{aligned} \mathbb{E}_{p^{\alpha'}}[v(k, \tilde{P}_{-w})] &= \sum_{P_{-w}} p^{\alpha'}(P_{-w}) v(k, P_{-w}) \\ &= \sum_{P_{-w}} [(1 - \alpha') p(P_{-w}) + \alpha' p^C(P_{-w})] v(k, P_{-w}) \\ &= \sum_{P_{-w}} \left[(1 - \alpha') p(P_{-w}) + \alpha \frac{1 - \alpha'}{1 - \alpha} p^C(P_{-w}) + \frac{\alpha' - \alpha}{1 - \alpha} p^C(P_{-w}) \right] v(k, P_{-w}) \\ &= \left(\frac{1 - \alpha'}{1 - \alpha} \right) \mathbb{E}_{p^{\alpha}}[v(k, \tilde{P}_{-w})] + \left(\frac{\alpha' - \alpha}{1 - \alpha} \right) \mathbb{E}_{p^C}[v(k, \tilde{P}_{-w})]. \end{aligned}$$

Now suppose that for $k_1, k_2 \in \{1, \dots, N\}$ with $k_2 > k_1$, we have

$$\mathbb{E}_{p^{\alpha}}[v(k_2, \tilde{P}_{-w})] \geq \mathbb{E}_{p^{\alpha}}[v(k_1, \tilde{P}_{-w})]. \quad (\text{A.1.3})$$

Then since

$$\mathbb{E}_{p^C}[v(k_2, \tilde{P}_{-w})] \geq \mathbb{E}_{p^C}[v(k_1, \tilde{P}_{-w})],$$

we must have

$$\mathbb{E}_{p^{\alpha'}}[v(k_2, \tilde{P}_{-w})] \geq \mathbb{E}_{p^{\alpha'}}[v(k_1, \tilde{P}_{-w})]. \quad (\text{A.1.4})$$

If the inequality in (A.1.3) is strict, then so too is the inequality in (A.1.4).

We can now use this payoff comparative static to sort optimal truncation points as follows.

By definition, $k_l(\alpha, p, u_w)$ satisfies

$$\mathbb{E}_{p^{\alpha}}[v(k_l(\alpha, p, u_w), \tilde{P}_{-w})] > \mathbb{E}_{p^{\alpha}}[v(k, \tilde{P}_{-w})] \quad \forall k < k_l(\alpha, p, u_w).$$

From (A.1.4), we must then have

$$\mathbb{E}_{p^{\alpha'}}[v(k_l(\alpha, p, u_w), \tilde{P}_{-w})] > \mathbb{E}_{p^{\alpha'}}[v(k, \tilde{P}_{-w})] \quad \forall k < k_l(\alpha, p, u_w),$$

so that $k_l(\alpha', p, u_w) \geq k_l(\alpha, p, u_w)$.

Similarly, $k_h(\alpha, p, u_w)$ satisfies

$$\mathbb{E}_{p^{\alpha}}[v(k_h(\alpha, p, u_w), \tilde{P}_{-w})] \geq \mathbb{E}_{p^{\alpha}}[v(k, \tilde{P}_{-w})] \quad \forall k < k_h(\alpha, p, u_w).$$

From (A.1.4), we must then have

$$\mathbb{E}_{p^{\alpha'}}[v(k_h(\alpha, p, u_w), \tilde{P}_{-w})] \geq \mathbb{E}_{p^{\alpha'}}[v(k, \tilde{P}_{-w})] \quad \forall k < k_h(\alpha, p, u_w),$$

so that $k_h(\alpha', p, u_w) \geq k_h(\alpha, p, u_w)$. $\square$

Fig. 6. Simulation results for truncation payoffs. The graphs display 101 – an individual woman's expected partner rank from truncating her list at each point $k \in \{0, \dots, 100\}$ and submitting these preferences to *MP-DA*. Preference lists of the other agents are uniformly random, selected from the set of all possible full length preference list profiles, and payoffs are averaged over 100,000 draws. All markets have 100 men, and the number of women varies between 90 and 110.

A.2. Unbalanced markets

Recently, [Ashlagi et al. \(2013\)](#) have shown that the requirement that the number of men and women is balanced can be crucial for the existence of a large gap in the expected rank of an agent's best and worst achievable mate. Unlike in balanced markets, in a typical realization of preferences in an unbalanced uniform market only a small number of agents (relative to the size of the market) have multiple partners. Their insight raises the question: to what extent does the optimal level of truncation depend on imbalances in the number of agents on each side of the market? The easy part of the answer is that all of our findings apart from [Theorem 1](#) are completely robust to such imbalances.

As for [Theorem 1](#), we first observe that under complete information, [Ashlagi et al. \(2013\)](#) implies that for a typical realization of preferences in an unbalanced uniform market only a small fraction of women have a strict incentive to misrepresent their preferences. But what does this tell us about decision-making under incomplete information, the setting of our analysis? Upon truncating, women face a tradeoff between safety and the prospect of improvement. If truncation is “safe” in the sense that it is unlikely to lead a woman to become unmatched, even a small probability of improvement may induce women to truncate their list. On the other hand, if the probability of becoming unmatched is higher, a higher probability of improvement is required to offset the expected disutility in the event of over-truncation.

The results of [Ashlagi et al. \(2013\)](#) regarding the small core apply regardless of the direction of the imbalance. That is, no matter the size or direction of the imbalance, the fraction of agents who have a more-preferred achievable mate is small. By contrast, simulation evidence suggests that in the setting of [Theorem 1](#), the direction of the imbalance matters. [Fig. 6](#) presents the results from simulating markets with 100 men and 101, 105 and 110 women (women are “underdemanded”) and markets with 100 men and 99, 95 and 90 women (women are “overdemanded”). In each market we generate truncation payoff graphs similar to [Fig. 1](#).

The simulations suggest that when women are underdemanded, the optimal preference list is quite long. In contrast, when women are overdemanded, it is hard to pinpoint their optimal truncation strategy based on the simulations, as the payoffs over a large range appear to be very flat. We formalize both of these observations below.

The following theorem states that long preference list submission is optimal when women are underdemanded.

Theorem (Coles, Gonczarowski, Shorrer). *Given $L > 0$, consider a market with N men and $N + L$ women. Let woman w have uniform beliefs and preferences linear in rank (or any strictly increasing, concave transformation of such preferences) and define*

$$k_*(N, L) \equiv \min \left\{ \left(\arg \max_{k \in \{0, \dots, N\}} \mathbb{E}[u_w(\mu^M[P_w^k, \bar{P}_{-w}](w))] \right) \right\}.$$

Then $\frac{k_(N, L)}{N} \geq \frac{L}{L+1}$ so $\lim_{N \rightarrow \infty} \frac{k_*(N, L)}{N} \geq \frac{L}{L+1}$. In particular $\frac{k_*(N, L)}{N} \geq \frac{1}{2}$ and $\lim_{N \rightarrow \infty} \frac{k_*(N, L)}{N} \geq \frac{1}{2}$.*

Like in [Theorem 1](#), $k_*(N, L)$ describes woman w 's optimal point of truncation, given that the other women submit their true preference lists. If there are multiple optima, we conservatively select that which involves the most truncation. The theorem constitutes a partial converse to [Theorem 1](#) for the case where women are underdemanded. That is, when there are more women than men, women optimally submit at least half their true preference list, both in small markets and in the limit.

The case where women are overdemanded is analytically more involved. We can show, however, that as is suggested by the simulations in [Fig. 6](#), in this case truncation is “safe” for women. More precisely, in large markets, unless a woman submits a very short list (less than $C \cdot \log^2 N$ for some C), the probability that she will become unmatched by truncating becomes extremely small as N grows large. Details of this result are presented formally in [Coles et al. \(2014\)](#).

In summary, significant truncation is optimal when the market is balanced, is “safe” when women are overdemanded (but unlikely to yield significant gains), and is suboptimal when women are underdemanded.

References

- Ashlagi, Itai, Klijn, Flip, 2012. Manipulability in matching markets: Conflict and coincidence of interests. *Soc. Choice Welfare* 39, 23–33.
- Ashlagi, Itai, Kanoria, Yashodhan, Leshno, Jacob D., 2013. Unbalanced random matching markets. In: *ACM Conference on Electronic Commerce*, pp. 27–28.
- Caldarelli, G., Capocci, A., 2000. Beauty and distance in the stable marriage problem. *Phys. A, Stat. Mech. Appl.* 300, 325–331.
- Coles, Peter, Cawley, John, Levine, Phil, Niederle, Muriel, Roth, Alvin E., Siegfried, John J., 2010. The job market for new economists: A market design perspective. *J. Econ. Perspect.* 24 (4), 187–206.
- Coles, Peter, Gonczarowski, Yannai, Shorrer, Ran, 2014. Strategic behavior in unbalanced matching markets. *Mimeo*.
- Coles, Peter, Kushnir, Alexey, Niederle, Muriel, 2013. Preference signaling in matching markets. *Amer. Econ. J. Microecon.* 5 (2), 99–134.
- Dubins, L.E., Freedman, D.A., 1981. Machiavelli and the Gale–Shapley algorithm. *Amer. Math. Mon.* 88 (7), 485–494.
- Dzierżawa, Michael, Oméro, Marie José, 2000. Statistics of stable marriages. *Physica A* 287, 321–333.
- Ehlers, Lars, 2004. In search of advice for participants in matching markets which use the deferred-acceptance algorithm. *Games Econ. Behav.* 48 (2), 249–270.
- Ehlers, Lars, 2008. Truncation strategies in matching markets. *Math. Oper. Res.* 33 (2), 327–335.
- Ehlers, Lars, Massó, Jordi, 2007. Incomplete information and singleton cores in matching markets. *J. Econ. Theory* 136, 587–600.
- Featherstone, Clayton R., Mayefsky, Eric, 2010. Stability and deferred acceptance: Strategic behavior in two-sided matching. Working paper.
- Gale, D., Shapley, L.S., 1962. College admissions and the stability of marriage. *Amer. Math. Mon.* 69, 9–15.
- Gale, David, Sotomayor, Marilda, 1985. Some remarks on the stable matching problem. *Discrete Appl. Math.* 11 (3), 223–232.
- Immorlica, Nicole, Mahdian, Mohammad, 2005. Marriage, honesty, and stability. In: *Proceedings of the Sixteenth Annual ACM–SIAM Symposium on Discrete Algorithms*.
- Knuth, Donald E., 1976. *Mariages stable et leurs relations avec d'autres problèmes combinatoires: Introduction à l'analyse mathématique des algorithmes*. Les Presses de l'Université de Montréal.
- Kojima, Fuhito, 2006. Mixed strategies in games of capacity manipulation in hospital-intern markets. *Soc. Choice Welfare* 27, 25–28.
- Kojima, Fuhito, Pathak, Parag, 2009. Incentives and stability in large two-sided matching markets. *Amer. Econ. Rev.* 99, 608–627.
- Konishi, Hideo, Ünver, Utku, 2006. Games of capacity manipulation in hospital-intern markets. *Soc. Choice Welfare* 27 (1), 3–24.
- Lee, Robin S., SangMok, 2011. Incentive compatibility of large centralized matching markets. Working paper.
- Lee, Robin S., Schwarz, Michael, 2012. Interviewing in two-sided matching markets. Working Paper.
- McVitie, D.G., Wilson, L.B., 1970. Stable marriage assignments for unequal sets. *BIT* 10, 295–309.

- McVitie, D.G., Wilson, L.B., 1971. The stable marriage problem. *Commun. ACM* 14 (7), 486–490.
- Nash, J., 1951. Non-cooperative games. *Ann. Math.* 54 (2), 286–295.
- Pittel, B., 1989. The average number of stable matchings. *SIAM J. Discrete Math.* 2 (4), 530–549.
- Pittel, Boris, 1992. On likely solutions of a stable marriage problem. *Ann. Appl. Probab.* 2 (2), 358–401.
- Pittel, Boris, Shepp, L., Veklerov, E., 2007. On the number of fixed pairs in a random instance of the stable marriage problem. *SIAM J. Discrete Math.*, 947–958.
- Roth, Alvin E., 1982. The economics of matching: Stability and incentives. *Math. Oper. Res.* 7, 617–628.
- Roth, Alvin E., 1985. The college admissions problem is not equivalent to the marriage problem. *J. Econ. Theory* 36 (2), 277–288.
- Roth, Alvin E., Peranson, Elliot, 1999. The redesign of the matching market for American physicians: Some engineering aspects of economic design. *Amer. Econ. Rev.* 89 (4), 748–780.
- Roth, Alvin E., Rothblum, Uriel G., 1999. Truncation strategies in matching markets – in search of advice for participants. *Econometrica* 67, 21–43.
- Roth, Alvin E., Vande Vate, John H., 1991. Incentives in two-sided matching with random stable mechanisms. *Econ. Theory* 1 (1), 31–44.
- Roth, Alvin E., Xing, Xiaolin, 1994. Jumping the gun: Imperfections and institutions related to the timing of market transactions. *Amer. Econ. Rev.* 84 (4), 992–1044.
- Sönmez, Tayfun, 1997. Manipulation via capacities in two-sided matching markets. *J. Econ. Theory* 77, 197–204.