

Gale-Shapley Stable Marriage Problem Revisited: Strategic Issues and Applications

Chung-Piaw Teo • Jay Sethuraman • Wee-Peng Tan

*Department of Decision Sciences, Faculty of Business Administration, National University of Singapore,
FBA 1-15 Law Link, Singapore 117591*

*Department of Industrial Engineering and Operations Research, Columbia University,
New York, New York 10027*

*Department of Decision Sciences, Faculty of Business Administration, National University of Singapore,
FBA 1-15 Law Link, Singapore 117591*

fbateocp@nus.edu.sg • jay@ieor.columbia.edu

We study strategic issues in the Gale-Shapley stable marriage model. In the first part of the paper, we derive the optimal cheating strategy and show that it is not always possible for a woman to recover her women-optimal stable partner from the men-optimal stable matching mechanism when she can only cheat by permuting her preferences. In fact, we show, using simulation, that the chances that a woman can benefit from cheating are slim. In the second part of the paper, we consider a two-sided matching market found in Singapore. We study the matching mechanism used by the Ministry of Education (MOE) in the placement of primary six students in secondary schools, and discuss why the current method has limited success in accommodating the preferences of the students, and the specific needs of the schools (in terms of the “mix” of admitted students). Using insights from the first part of the paper, we show that stable matching mechanisms are more appropriate in this matching market and explain why the strategic behavior of the students need not be a major concern.

(Stable Marriage; Strategic Issues; Gale-Shapley Algorithm; Student Posting Exercise)

1. Introduction

This paper is motivated by a study of the mechanism used to assign primary school students in Singapore to secondary schools. The current assignment process requires that the primary school students submit a rank-ordered list of six schools to the Ministry of Education. Students are then assigned to secondary schools based on their preferences, with priority given to those with higher examination scores in the Primary School Leaving Examination (PSLE). The current matching mechanism is plagued by several problems. We argue that a satisfactory resolution of these problems necessitates the use of a *stable*

matching mechanism (a complete description appears in §4). In fact, two well-known stable matching mechanisms—the *students-optimal* stable mechanism, and the *schools-optimal* stable mechanism—appear to be natural candidates; however, neither of these mechanisms (nor any other stable matching mechanism) is capable of eliciting truthful participation from *all* of the participants. Our main purpose is to show that the students have very little room to misrepresent their preferences, under either mechanism. While the strategic issues facing the students under the students-optimal mechanism are well understood, the insights under the schools-optimal mechanism

appear to be new. The most suitable stable matching mechanism for the Singapore Posting Exercise will ultimately be a variant of these two well-known matching mechanisms, as local issues and features have to be incorporated. However, our results on simulation experiments under the students-optimal and schools-optimal mechanisms suggest that, regardless of the choice of the stable matching mechanism, the students have little incentive to misrepresent their preferences. This addresses one of the major concerns of the participants in the Singapore Posting Exercise.

A distinguishing feature of our model is that the participants are required to submit *complete* preference lists. Strategic issues in the stable marriage problem have been explored in settings that allow incomplete preference lists; but, to the best of our knowledge, no systematic study has been made when complete preference lists are required. As we shall see, such a requirement imposes a severe restriction on the strategic choices available to the participants, resulting in substantially reduced incentives to cheat. To study the immunity of the schools-optimal mechanism to cheating by the students, we need to first determine the optimal cheating strategy in the classical stable marriage model where the participants are only allowed to cheat by reversing the true preference order of acceptable partners; They are not allowed to truncate the preference list, as is the case usually considered in the literature. The main result of this paper shows that this problem can be solved in polynomial time, using a simple extension of the classical Gale-Shapley algorithm.

Preliminaries. Stable matching problems were first studied by Gale and Shapley (1962). In a stable marriage problem there are two finite sets of participants: the set of men (M) and the set of women (W). We assume that every member of each sex has strict preferences over the members of the opposite sex. In the model that allows rejection, the preference list of a participant can be “truncated” in the sense that participants have the option of declaring some others as unacceptable; in this model, we also include the possibility that a participant may be unmatched, i.e., the participant’s assigned partner in the matching is himself/herself. In contrast, in the *Gale-Shapley* model, the preference lists of the participants are required to

be *complete*, and no one is to be declared as unacceptable. A matching is just a one-to-one mapping between the two sexes such that a man m is mapped to a woman w if and only if w is mapped to m , and m and w are acceptable to each other. (In the economics literature, this is defined as an *individually rational* matching.) A matching μ is said to be *unstable* (under either model) if there is a man-woman pair, who both prefer each other to their partners in μ ; this pair is said to block the matching μ , and is called a *blocking pair* for μ . A *stable matching* is a matching that is not unstable. The significance of stability is best highlighted by a system where acceptance of the proposed matching is voluntary. In such a setting, an unstable matching cannot be expected to remain intact, as the blocking pair(s) may discover that they could both improve their match by joint action: The man and woman involved in a blocking pair could just “divorce” their respective partners and “elope.”

In addition to formulating several versions of the stable matching problem, Gale and Shapley (1962) described a simple algorithm that always finds a stable matching for any instance of the stable marriage problem. This elegant result sparked the interest of many researchers, resulting in a thorough investigation of the structural properties of the stable marriage model. A property that is especially relevant to our study is that the set of all stable matchings for an instance of the stable marriage problem forms a lattice, with the extremal elements being the so-called “men-optimal” and “women-optimal” stable matchings. In fact, the algorithm of Gale and Shapley (1962) that established the existence of a stable marriage constructs a men-optimal stable matching. This algorithm is commonly known as the *men-propose* algorithm because it can be expressed as a sequence of “proposals” from the men to the women. Note that the Gale-Shapley algorithm can be easily adapted to yield a women-optimal stable matching by simply interchanging the roles of men and women; this is commonly called the *women-propose* algorithm. All of the results in this paper will be stated under the assumption that the men-optimal stable matching mechanism is used; by interchanging the role of men and women, analogous results can be derived for the women-optimal mechanism.

Strategy. Consider a market in which men and women submit their preference lists to a centralized agency that matches the participants by computing the men-optimal stable matching. An important difficulty that arises is that this matching mechanism is manipulable by the women: Some women can intentionally submit false preferences and obtain better partners. Such (strategic) questions have been studied for the stable marriage problem by mathematical economists and game theorists, with the goal of quantifying the potential gains of a deceitful participant. An early result in this direction is due to Roth (1982), who proved that when the men-optimal mechanism is used, none of the men benefits by submitting false preferences, regardless of how the other participants report their preferences. By submitting a false preference list, a man can, at best, obtain his (true) optimum stable partner, which he obtains anyway by submitting his true preference list. In game-theoretic parlance, stating true preferences is a *dominant strategy* for the men. However, Gale and Sotomayor (1985b) showed that the women can still force the men-optimal mechanism to return the women-optimal solution. The optimal cheating strategy in this case is simple: Each woman w submits a preference list that is the same as her true preference list, except that she declares men who rank below her women-optimal partner as unacceptable. Our main goal in this paper is to study the analogous question in the Gale-Shapley model, in which all the participants are required to submit complete preference lists. We emphasize that we consider only a centralized market that computes the men-optimal matching, and all the men and women can only manipulate the outcome by permuting their preference lists. We do not consider a decentralized market, where manipulation possibilities are richer, nor do we consider other ways of manipulating the outcome.

Motivation. The motivation for studying strategic questions in the Gale-Shapley model is the Singapore school admissions problem (referred to hereafter as the MOE (Ministry of Education) problem) briefly described earlier. Currently, the schools are merely seen as resources by the ministry to be allocated to students. Furthermore, for ease of implementation, each student is only allowed to submit six

choices, and the ministry pads each list by ranking the remaining schools using location as the only criterion (schools close to the student's home are ranked higher). All assignments are done in a centralized manner, and no student is allowed to approach the schools privately for admission purposes.

However, such passive roles for the schools will change in the near future. In preparation for a knowledge-based economy, the ministry has reiterated its intention to shift from the current examination-oriented system to one that focuses on ability-driven education. In line with this goal, the ministry intends to give the schools administrative and professional autonomy. Student assessments will also be reviewed to meet the objective of developing creative independent learners. The new university admissions system to be implemented from 2003 will not rely solely on the GCE A-level exam results, but instead make a holistic evaluation of a student's potential. The current assignment process of students to schools is not consistent with this new focus of the ministry. Furthermore, the design of the current assignment process has also given rise to several operational problems (see §4), and our purpose is to argue that a centralized stable matching mechanism of the type proposed by Gale and Shapley is superior. Our goal, in the end, is to convince the readers that regardless of whether the students-optimal or schools-optimal mechanism is used, there is no significant incentive for the students or the schools to strategize and exploit the system.

There is an obvious relation between the MOE problem and the stable marriage problem introduced earlier: In the MOE problem, the students play the role of "women" and the secondary schools play the role of "men." Observe that there is also a crucial difference: In the MOE problem, many "women" (students) can be assigned to the same "man" (secondary school), whereas in the stable marriage model we require that the matching be one-to-one. Since the students (schools) in this setting are not allowed to declare any school (student) as unacceptable, the MOE problem can be modeled as a many-to-one matching problem with complete preference lists.

The issues of strategic manipulation in the stable marriage model with rejection are well understood (cf. Roth and Sotomayor 1991 and the references

therein). However, little is known in the case of stable marriage models (one-to-one and many-to-one) without rejection, which, as we discussed, is a more suitable representation of the MOE problem; a notable exception is the recent work of Tadenuma and Toda (1998), in which they consider the implementation question, and show that no stable matching correspondence can be implemented in Nash equilibrium as long as $|M| = |W| > 2$. (Stronger nonimplementability results hold for the rejection model, see Kara and Sonmez 1996.)

Results and Structure of the Paper. To derive insights into the strategic behavior of the participants in the MOE problem, we first consider the one-to-one model, and study the following question: In the stable marriage model with complete preferences, with the men-optimal mechanism, is there an incentive for the women to cheat? If so, what is an *optimal* cheating strategy for a woman? Can a woman always force the mechanism to return her women-optimal partner? In §2 we present the main result of this paper: An optimal cheating strategy can be constructed in polynomial time. A related issue concerns the robustness of the men-optimal mechanism. In the rejection model, it is well known that the women can easily manipulate the men-optimal mechanism, and, in fact, almost all the women will submit false preferences. In sharp contrast, our results for the Gale-Shapley model in §3 provide evidence that the men-optimal mechanism is fairly *robust*, and that there is very little incentive for the women to cheat. In particular, restricting the strategic choices of the women drastically reduces their benefits from cheating, thereby reducing the possibility that a woman will cheat. Armed with this theoretical understanding of the Gale-Shapley model, we prescribe some recommendations in §4 to improve the current “matching mechanism” used in the Singapore MOE Secondary School Posting Exercise. In particular, we argue that a stable matching mechanism is more appropriate for the MOE problem, and that some of the other difficulties inherent in the present system can be effectively addressed by a stable matching mechanism.

It is instructive to compare our results to some recent results obtained in a very interesting study by

Roth and Rothblum (1999). In a centralized many-to-one market (with rejection allowed) operating under the schools-optimal mechanism, Roth and Rothblum consider the strategic issues facing the students. In a low information environment, where preferences of the other participants are not known with certainty, they conclude that stating preferences that reverse the true preference order of two acceptable schools is not beneficial to the students, while submitting a truncation of the true preferences may be. Our simulation-based experimental results in this paper suggest that this observation is valid even in the perfect information setting, i.e., knowing the reported preferences of all the other participants may still not allow one to benefit from cheating, if one is only allowed to reverse the order of the schools, but not allowed to submit a truncated list.

Finally, we note that the strategic issues facing the schools, under the schools-optimal mechanism, is a nontrivial problem: Roth (1985) showed that, contrary to the one-to-one case, the schools-optimal matching need not even be weakly-Pareto optimal for the schools, and there is no stable matching mechanism that makes it a dominant strategy for all the schools to state their true preferences. We shall see in §4 why this is not an important issue in the MOE problem.

2. Optimal Cheating in the Gale-Shapley Model

Our standing assumption in this paper is that woman w is the only deceitful participant and she knows the reported preferences of all the other participants, which remain fixed throughout. We shall show, eventually, that cheating opportunities for woman w are uncommon, in spite of the assumption that she has perfect information. We consider only a centralized market, although the algorithm is phrased as a sequence of proposals from the men to the women. We visualize the deceitful woman as running this algorithm in her head and submitting the optimal cheating strategy thus computed as her preference list to the centralized market.

We shall begin by considering the following question: Consider a centralized market in which the men-propose algorithm is literally used to compute the

men-optimal matching. Suppose woman w has *no knowledge* of the preferences of any of the other participants, and that she is the only deceitful participant. Suppose also that she is allowed to reject proposals. Is it possible for her to identify her women-optimal partner by just observing the sequence of proposals she receives? Somewhat surprisingly, the answer is yes! If w simply rejects all the proposals made to her, then the best (according to her true preference list) man among those who propose to her is her *women-optimal* partner. Our algorithm for finding the optimal cheating strategy in the Gale-Shapley model builds on this insight: Woman w rejects as many proposals as possible, while remaining engaged to some man who proposed earlier in the algorithm. Using a backtracking scheme, she uses the matching mechanism repeatedly to find her optimal cheating strategy. Given our standing assumption that woman w has complete knowledge of the reported preferences of the other participants, and that she is the only agent acting strategically, it is clear what she would do: She would run the algorithm to find her optimal cheating strategy privately and submit this (optimal) preference list to the centralized market.

2.1. Finding Your Optimal Partner

We first describe Algorithm *OP*—an algorithm to compute the women-optimal partner for w using the men-propose algorithm. (Recall that we do this under the assumption that woman w is allowed to remain single.)

Algorithm *OP*

1. Run the men-propose algorithm, and reject all proposals made to w . At the end, w and a man, say m , will remain single.
2. Among all the men who proposed to w in Step 1, let the best man (according to w) be m' .

THEOREM 1. m' is the women-optimal partner for w .

REMARK. Gale and Sotomayor (1985b) showed that when each woman declares all the men inferior to her women-optimal partner as unacceptable, then the men-optimal mechanism will be forced to return the women-optimal stable matching. This is because the only stable matching solution for the modified

preference lists is the women-optimal solution (with respect to the original preference lists). To prove our result, however, we have to show that when a woman unilaterally declares all the men as unacceptable, this is enough to induce her optimal partner to propose to her in the course of executing the men-propose algorithm. Furthermore, we need to show that no higher-ranked man on her list will propose to her even after she rejects the proposal from her optimal partner.

PROOF OF THEOREM 1. Let $\mu(w)$ denote the women-optimal partner for w . We modify w 's preference list by inserting the option to remain single in the list, *immediately after* $\mu(w)$. (We declare all men that are inferior to $\mu(w)$ as unacceptable to w .) Consequently, in the men-propose algorithm, all proposals inferior to $\mu(w)$ will be rejected. Nevertheless, since there is a stable matching (with respect to the original preferences) with w matched to $\mu(w)$, our modification does not destroy this solution, i.e., this solution remains stable with respect to the modified list. It is also well known that in any stable matching instance, the set of people who are single is the same for all stable matchings (cf. Roth and Sotomayor 1990, p. 42). Thus, w must be matched in *all* stable matchings with respect to the modified preference list. The men-optimal matching for this modified preference list must match w to $\mu(w)$, since $\mu(w)$ is now the worst partner for w with respect to the modified list. In particular, $\mu(w)$ must have proposed to w during the execution of the men-propose algorithm. Note that until $\mu(w)$ proposes to w , the men-propose algorithm for the modified list runs exactly in the same manner as in Step 1 of *OP*. The difference is that Step 1 of *OP* will reject the proposal from $\mu(w)$, while the men-propose algorithm for the modified list will accept the proposal from $\mu(w)$, as w prefers $\mu(w)$ to being single. Hence, clearly $\mu(w)$ is among those who proposed to w in Step 1 of *OP*, and so $m' \succeq_w \mu(w)$.

Suppose $m' >_w \mu(w)$. Consider the modified list in which we place the option of remaining single *immediately after* m' . We run the men-propose algorithm with this modified list. Again, until m' proposes to w , the algorithm runs exactly the same as in Step 1 of *OP*, after which the algorithm returns a stable partner for w who is at least as good as m' according to w (under both the original and the modified lists, since

we have not altered the order of the men before m' on the list). The matching obtained is thus stable with respect to the original list. This contradicts our earlier assumption that $\mu(w)$ is the best stable partner for w . \square

Suppose w is allowed to modify her preference list after each run of the men-propose algorithm, and the algorithm is to be repeated until w concludes that she has found her best possible partner. Theorem 1 essentially says that knowing the set of proposals woman w receives is enough to allow her to construct her optimal cheating strategy, if she is allowed to declare certain men as unacceptable; she need not know the preferences of *any* of the other participants involved, as long as they behave truthfully. In the next section, we shall use this insight to construct an optimal cheating algorithm for w , under the additional restriction that she is not allowed to declare any man as unacceptable.

2.2. Cheating Your Way to a Better Marriage

Observe that the procedure of §2.1 only works when woman w is allowed to remain single throughout the matching process. Suppose the central planner does not give the woman an option to declare any man as unacceptable. How do we determine her best attainable partner by manipulation? This is essentially a restatement of our original question: Who is the best stable partner woman w can have when the men-optimal mechanism is used and when she can lie only by permuting her preference list. Note that the preferences of the remaining participants (except woman w) are fixed throughout.

A natural extension of Algorithm *OP* is for woman w to: (i) accept a proposal first, and then reject all future proposals; (ii) from the list of men who proposed to her but were rejected, find her most preferred partner; repeat the men-propose algorithm until the stage when this man proposes to her; (iii) reverse the earlier decision and accept the proposal from this most preferred partner, and continue the men-propose algorithm by rejecting all future proposals; and (iv) repeat (ii) and (iii) until w cannot obtain a better partner from all other proposals. Unfortunately, this simple strategy does not always yield the best stable partner a woman can achieve

under the Gale-Shapley model. The reason is that this greedy improvement technique does not allow for the possibility of rejecting the current best partner, in the hope that this rejection will trigger a proposal from a *better* would-be partner. Our algorithm in this paper, which is described next, does precisely that. An illustrative example appears soon after the description of the algorithm.

Let $P(w) = \{m_1, m_2, \dots, m_n\}$ be the true preference list of woman w , and let $P(m, w)$ be a preference list for w so that the men-propose algorithm will return m as her men-optimal partner. In the case that man m cannot be obtained by w as her men-optimal partner for any preference list, we set $P(m, w)$ to be the null list, as a matter of convention. Our algorithm constructs $P(m, w)$ (if man m is attainable by woman w) or determines whether $P(m, w)$ is the null list (if man m is not attainable by woman w) iteratively, and consists of the following steps:

1. Run the men-propose algorithm with the true preference list $P(w)$ for woman w . Keep track of all men who propose to w . Let the men-optimal partner for w be m , and let $P(m, w)$ be the true preference list $P(w)$.

2. Suppose m_j proposed to w in the men-propose algorithm. By moving m_j to the front of the list $P(m, w)$, we obtain a preference list for w such that the new men-optimal partner (after running the men-propose algorithm on this modified list) is m_j . Let $P(m_j, w)$ be this altered list. We say that m_j is a *potential* partner for w .

3. Repeat Step 2 to obtain a preference list for every man (other than m) who proposed to woman w in the algorithm; after this, we say that we have exhausted man m , the men-optimal partner obtained with the preference list $P(m, w)$. All potential partners of w that come from running the men-propose algorithm using the list $P(m, w)$ have been found.

4. If a potential partner for w , say man u , has not been exhausted, run the men-propose algorithm with $P(u, w)$ as the preference list of w . Identify new potential partners. Repeat Steps 2–3 with u in place of m .

5. Repeat Step 4 until all potential partners of w are exhausted. Let N denote the set of all potential (and hence exhausted) partners for w .

6. Among the men in N let m_a be the man woman w prefers most. Then $P(m_a, w)$ is an optimal cheating strategy for w .

The men in the set N at the end of the algorithm have the following crucial properties:

- For each man m in N , there is an associated preference list for w such that the men-propose algorithm returns m as the men-optimal partner for w with this list.
- All other proposals in the course of the men-propose algorithm come from other men in N . (Otherwise, there will be some potential partners who are not exhausted.)

With each run of the men-propose algorithm, we exhaust a potential partner, and so this algorithm executes at most n men-propose algorithms before termination.

EXAMPLE 1. Consider the following stable marriage problem:

1 2 3 4 5 1	1 1 2 3 5 4
2 3 4 5 1 2	2 2 1 4 5 3
3 5 1 4 2 3	3 3 2 5 1 4
4 3 1 2 4 5	4 4 5 1 2 3
5 1 5 2 3 4	5 5 1 2 3 4
True Preferences of the Men	True Preferences of the Women

We construct the optimal cheating strategy for Woman 1 using the algorithm described earlier.

Step 1. Run the men-propose algorithm with the true preference list for Woman 1; her men-optimal partner is Man 5. Man 4 is the only other man who proposes to her during the men-propose algorithm. So $P(\text{Man 5, Woman 1}) = (1, 2, 3, 5, 4)$.

Steps 2–3. Man 4 is moved to the head of Woman 1's preference list; i.e., $P(\text{Man 4, Woman 1}) = (4, 1, 2, 3, 5)$. Man 5 is exhausted, and Man 4 is a potential partner.

Step 4. As Man 4 is not yet exhausted, we run the men-propose algorithm with $P(\text{Man 4, Woman 1})$ as the preference list for Woman 1. Man 3 is identified as a new possible partner, with $P(\text{Man 3, Woman 1}) = (3, 4, 1, 2, 5)$. Man 4 is exhausted after this.

Repeat Step 4. As Man 3 is not yet exhausted, we run the men-propose algorithm with $P(\text{Man 3, Woman 1})$ as the preference list for Woman 1. Man 3 will be

exhausted after this. No new potential partner is found, and so the algorithm terminates.

REMARK. Example 1 shows that woman 1 could cheat and get a partner better than her men-optimal partner. However, her women-optimal partner in this case turns out to be Man 1. Hence Example 1 also shows that Woman 1 cannot always assure herself of her women-optimal partner through cheating, in contrast to the case when rejection is allowed in the cheating strategy.

We conclude this section by stating and proving the main result. Recall that m_a is the *best* man according to w 's true preference list among the men in set N constructed by the algorithm.

THEOREM 2. $\pi = P(m_a, w)$ is an optimal strategy for woman w .

PROOF (BY CONTRADICTION). We use the convention that $r(m) = k$ if man m is the k th man on woman w 's true preference list. Let $\pi^* = \{m_{p1}, m_{p2}, \dots, m_{pn}\}$ be a preference list that gives w her *best* stable partner when the men-optimal mechanism is used. Let this man be denoted by m_{pb} , and let woman w strictly prefer m_{pb} to m_a (under her true preference list), i.e., $r(m_{pb}) < r(m_a)$. Observe that the order of the men who do not propose to woman w is irrelevant and does not affect the outcome of the men-propose algorithm. Furthermore, men of rank higher than $r(m_{pb})$ do not get to propose to w , otherwise we can cheat further and improve on the best partner for w , contradicting the optimality of π^* . Thus we can arbitrarily alter the order of these men, without affecting the outcome. Without loss of generality, we may assume that $1 = r(m_{p1}) < 2 = r(m_{p2}) < \dots < b = r(m_{pb})$. Since $r(m_{pb}) < r(m_a)$, m_a will appear somewhere after m_{pb} in π^* : thus, m_a can be any of the men in the list $m_{p(b+1)}, m_{p(b+2)}, \dots, m_{pn}$.

Now, we modify π^* such that all men who (numerically) rank lower than m_a but higher than m_{pb} (under true preferences) are put in order according to their ranks. This is accomplished by moving all these men before m_a in π^* . With that alteration, we obtain a new list $\tilde{\pi} = \{m_{q1}, m_{q2}, \dots, m_{qs}\}$ such that:

- (i) $1 = r(m_{q1}) < 2 = r(m_{q2}) < \dots < s = r(m_{qs})$.

(ii) $m_{q1} = m_{p1} \dots m_{qb} = m_{pb}$, where the position of those men who rank higher than m_{pb} is unchanged.

(iii) $r(m_a) = s + 1, m_a \in \{m_{q(s+1)}, m_{q(s+2)}, \dots, m_{qn}\}$.

(iv) The men in the set $\{m_{q(s+1)}, m_{q(s+2)}, \dots, m_{qn}\}$ retain their relative position with respect to one another under π^* .

Note that the men-optimal partner of w under $\tilde{\pi}$ cannot come from the set $\{m_{q(s+1)}, m_{q(s+2)}, \dots, m_{qn}\}$. Otherwise, since the set of men who proposed in the course of the algorithm must come from $\{m_{q(s+1)}, m_{q(s+2)}, \dots, m_{qn}\}$, and since the preference list π^* retains the relative order of the men in this set, the same partner would be obtained under π^* . This leads to a contradiction as π^* is supposed to return a better partner for w . Hence, we can see that under $\tilde{\pi}$, we already get a better partner than under π .

Now, since the preference list π returns m_a with $r(m_a) = s + 1$, we may conclude that the set N (obtained from the final stage of the algorithm) does not contain any man of rank lower than $s + 1$. Thus $N \subseteq \{m_{q(s+1)}, m_{q(s+2)}, \dots, m_{qn}\}$. Suppose $m_{q(s+1)}, m_{q(s+2)}, \dots, m_{qw}$ do not belong to the set N , and $m_{q(w+1)}$ is the first man after m_{qs} who belongs to the set N . By construction of N , there exists a permutation $\hat{\pi}$ with $m_{q(w+1)}$ as the stable partner for w under the men-optimal mechanism. Furthermore, all of those who propose to w in the course of the algorithm are in N , and hence they are no better than m_a to w . Furthermore, all proposals come from men in $\{m_{q(w+1)}, m_{q(w+2)}, \dots, m_{qn}\}$, since $N \subseteq \{m_{q(s+1)}, m_{q(s+2)}, \dots, m_{qn}\}$.

By altering the order of those who did not propose to w , we may assume that $\hat{\pi}$ takes the following form: $\{m_{q1}, m_{q2}, \dots, m_{q(s-1)}, m_{qs}, \dots, m_{qw}, m_{q(w+1)}, \dots\}$, where the first $qw + 1$ men in the list are identical to those in $\tilde{\pi}$. But, the men-optimal stable solution obtained using $\hat{\pi}$ must also be stable under $\tilde{\pi}$, since w is match to $m_{q(w+1)}$, and the set of men she strictly prefers to $m_{q(w+1)}$ is identical in both $\hat{\pi}$ and $\tilde{\pi}$. This is a contradiction as $\tilde{\pi}$ is supposed to return a men-optimal solution better than m_a .

This implies that π^* does not exist, and so π is optimum and m_a is the best stable partner w can get by permuting her preference list. \square

3. Strategic Issues in the Gale-Shapley Problem

By requiring the women to submit complete preference lists, we are clearly restricting their strategic options; thus many of the strong structural results known for the rejection model may not hold for the Gale-Shapley model. This is good news, for it reduces the incentive for a woman to cheat. In this section, we present several examples to show that the strategic behavior of the women can be very different under the models with and without rejection.

3.1. The Best Man (Obtained by Cheating) May Not Be Women-Optimal

In the two-sided matching model with rejection, it is not difficult to see that the women can always force the men-optimal mechanism to return the women-optimal solution (for instance, each woman declares as unacceptable those men who are inferior to her true women-optimal partner). In the Gale-Shapley model, which forbids rejection, the influence of the women is far less, even if they collude. A simple example is when each woman is ranked first by exactly one man. In this case, there is no "conflict" among the men, and in the men-optimal solution, each man is matched to the woman he ranks first. In this case, the algorithm will terminate with the men-optimal matching, regardless of how the women rank the men in their lists. So ruling out the strategic option of remaining single for the women significantly affects their ability to change the outcome of the game by cheating.

By repeating the above analysis for all the other women in Example 1, we conclude that, by cheating unilaterally, the best possible partners for Women 1, 2, 3, 4, and 5 are, respectively, Men 3, 1, 2, 4, and 3. An interesting observation is that Woman 5 cannot benefit by cheating alone (she can only get her men-optimal partner no matter how she cheats). However, if Woman 1 cheats using the preference list (3, 4, 1, 2, 5), Woman 5 will also benefit by being matched to Man 5, who is first in her list.

3.2. Does Cheating Pay?

Roth (1982) shows that under the men-optimal mechanism, the men have no incentive to alter their true preference lists. In the rejection model, however, Gale

and Sotomayor (1985a) show that a woman has an incentive to cheat as long as she has at least two distinct stable partners. From Pittel (1989), we know that the average number of stable matchings is asymptotically $n \log(n)/e$, and with high probability, the rank of the women-optimal and men-optimal partners for the women are, respectively, $\log(n)$ and $n/\log(n)$. Thus, in typical instances of the stable marriage game under the rejection model, most of the women will not reveal their true preferences.

Many researchers have argued that the troubling implications from these studies are not relevant in practice, as the model assumes that the women have complete knowledge of all the other participants, as well as their preference lists. For the Gale-Shapley model we consider, it is natural to ask whether it pays (as in the rejection model) for a typical woman to solicit information about the preferences of the remaining participants in the game.

EXAMPLE 2. We have designed a Visual Basic package for the stable marriage problem with complete preference lists. The following example, in one of the simulation experiments, illustrates the dramatic difference in the behavior of the women in the models with and without rejection:

In the example, the men and women are labelled from 1 to 8. Both man i and woman i have identical preference lists. The number next to the list in Figure 1

represents the men-optimal and women-optimal partners of each participant. For instance, Man 1 and Woman 1 have preference lists "64782351" (i.e., Man 1 prefers Woman 6 to Woman 4 to Woman 7, etc., and similarly Woman 1 prefers Man 6 to Man 4 to Man 7, etc.). Under the men-optimal mechanism, Man 1 is matched to Woman 6, and Woman 1 is matched to Man 7. Thus the women-optimal partner for Woman 1 is Man 6 (from symmetry), and so she has an incentive to lie about her preferences. In fact, this example shows that 6 out of the 8 women benefit from cheating in the rejection model. The situation for the Gale-Shapley model is very different: given that the women can only misrepresent their preferences by order reversals, no woman can benefit from cheating.

We ran the algorithm on 1,000 instances, with preferences of each individual generated uniformly at random, for $n = 8$. The number of women who benefit from cheating is tabulated in Table 1.

Interestingly, the number of women who benefit from cheating is surprisingly low. In fact, in 74% of the instances, the men-optimal solution is their only option, no matter how they cheat. The average percentage of women who benefit from cheating is merely 5.06.

To look at the typical strategic behavior on larger instances of the stable marriage problem, we ran the heuristic on 1,000 random instances for $n = 100$. The cumulative plot is shown in Figure 2. In particular, in more than 60% of the instances at most 10 women (out of 100) benefited from cheating, and in more than 96% of the instances at most 20 women benefited from cheating. The average number of women who benefited from cheating is 9.515%. Thus, the chances that a typical woman can benefit from acquiring perfect information (i.e., knowing the preferences of the other participants) is slim in the Gale-Shapley model.

We have performed similar experiments for large instances of the Gale-Shapley model. Because of computational requirements, we can only run the experiment on 100 random instances of the problem with

Figure 1 Stable Marriage Model with Rejection Versus One Without Rejection





Table 1

Number of women who benefited	0	1	2	3	4	5	6	7	8
Number of observations	740	151	82	19	7	1	0	0	0

Figure 2 Benefits of Cheating



500 men and women. The insights obtained from the 100 by 100 cases carry over: the number of women who benefit from cheating is again not more than 10% of the total number of the women involved. In fact, the average was close to 6% of the women population in the problem. This suggests that the number of women who benefit from cheating in the Gale-Shapley model with n women grows at a rate slower than a linear function of n . A detailed probabilistic analysis of this phenomenon is a challenging problem, which we leave for future research.

4. Singapore MOE Posting Exercise

The admission to secondary schools in Singapore is centrally controlled. A distinguishing feature of the Singapore school admissions problem (MOE problem) is that the students are not allowed to remain unassigned at the end of the posting exercise, and that the schools are not allowed to reject any students if there are still places available. In what follows, we first present a summary of the current matching mechanism, along with a brief description of its shortcomings in §4.1. In §4.2, we explain why a stable matching

solution is more appropriate, and how some of the problems under the current system can be addressed using a stable matching mechanism.

4.1. Current Assignment Process

Prior to taking their Primary School Leaving Examination (PSLE), the students are required to submit (around August) their rank ordering of schools to the MOE, which oversees the posting exercise. In the option form, each student is required to list *six* secondary schools she would like to attend, in order of decreasing desirability. The students (rather, their parents) are advised by the MOE to make realistic choices because the preference list, once submitted, cannot be changed. After the PSLE results are announced, the top ten percent performers are given a second option to indicate their preferences for the Independent schools and the Special Assistance Plan (SAP) schools. These are the elite schools in Singapore that admit only the top students.

The posting of students to secondary schools is (partially) computerized. All students are ranked according to their aggregate PSLE scores regardless of

their choices. Currently, the assignment of students to schools is conducted in multiple phases, which are as follows:

Phase 1. Consider the top students who have applied to the independent schools in the second option, and assign them to these schools.

Phase 2. Consider the top students who have applied to SAP schools in the second option, and assign them to SAP schools.

Phase 3. For historical and administrative reasons, certain secondary schools have affiliated primary schools. This is to foster closer collaboration between the primary and secondary schools. Students who opt to go to the affiliated secondary schools will be given higher priority. Consider all students who have selected their affiliated secondary school as their top choice, and assign them to the affiliated schools (subject to availability of places).

Phase 4. Consider the remaining students one by one, and match them to the schools according to the choices they made in August.

Note that in the process of assigning the students to schools, a student who ranks his affiliated secondary school as his first choice is given priority for that school. In this way, the current assignment process has managed to incorporate some of the schools' selection criterion, which actually place more importance on "affiliation" than on performance in the PSLE; not all students, however, can gain admission to their affiliated secondary schools because such admissions are subject to vacancies in these schools. When a student fails to get admission to any of the schools of her choice, she is assigned (manually) to a school (in *nearby postal districts*) that still has vacancies. If there is no such school available, the student is assigned (manually) to a school in some other postal district that still has vacancies. The students are notified of the schools they are assigned to by late December, just in time for the new school semester, which starts early January. The current system is plagued with several problems, some of which are as follows.

Problem 1. Decision making by the parents is complicated by the fact that (i) parents are only given limited options to express their preferences, and (ii) parents have to make their decisions prior to the

release of the PSLE scores. Knowing that many good students will be applying to the good schools, parents cannot afford to reveal their true preferences. If they were to list down just the ideal top six schools as their six choices, they run the risk of being rejected by all these six schools—this happens if their child performs below their expectations in the PSLE and fails to qualify for these schools. Worse, this may actually result in their child getting into a school that is completely out of their consideration, as manual posting would then be carried out. The parents are thus forced to be more "realistic" when choosing the schools for their child. Consequently, the parents are faced with the difficult problem of optimally utilizing the six options provided by the MOE.

The MOE has tried to simplify matters for the parents by providing them with information such as the schools' cut-off points for the previous year to help them make rational decisions. In recent years, the MOE has also provided a rating of all the schools in Singapore so that the parents will be better informed of the schools' strengths and weaknesses. However, this still does not address the parents' major concern in formulating their choices: they have to submit their choices before their children take the examination! In fact, one parent lamented on a national newspaper: "How do we choose when we do not know how they will do at the national level?" (*The Straits Times*, 17 August 1996). *A good matching mechanism should allow the students to submit their true ranking list, without the need to strategize based on their PSLE results.*

We note that it is not possible for the Ministry to solicit the students' options *after* the release of the examination results, due to the tight time span between the release of results and the start of the new school term.

Problem 2. Another problem with the current assignment process is that other than affiliation, the genuine preferences of the schools are not considered. The schools have to accept the students assigned to them under the process. Furthermore, the number of places available in the school is also largely determined by the MOE. To attract the desired students, many schools, especially the newer ones, resort to advertising campaigns such as organizing open

houses, sending brochures to parents, and conducting talks at primary schools. Though such marketing efforts help, schools still depend significantly on the ST School 100 ranking (provided by the MOE) to attract their desired students. However, since the ST ranking is decided mainly by the schools' performance in the GCE "O"-level examination, over-reliance on the ranking could induce the schools to overemphasize the importance of academic results. This is clearly an undesirable repercussion as it actually impedes the Singapore government from achieving the "Thinking Schools, Learning Nation" vision. *A good matching mechanism should allow the schools to rank the students using different criteria, depending on the strength and areas of excellence the schools wish to emphasize.*

Problem 3. A third problem lies in the logic of the current assignment process at Phase 4 of the exercise. When all the six choices are exhausted, the students are then assigned to schools in nearby postal districts. However, according to a 1992 survey (*The Straits Times*, 19 August 1992), only 30% of the respondents think that the *distance to school* is an important factor in their selection of schools. While this is a nonnegligible fraction of the respondents, it is distressing that several other factors, considered *more important* by the respondents, were completely ignored in the current assignment process. The major factors identified in the survey are:

Which Factors Matter Most	Average Score (out of 10)
Quality of O-level passes	7.5
Percentage of O-level passes	7.4
Entry cut-off points	6.6
Value-added	6.6
Percentage of students accepted	6.1
ECA performance	4.9
Other factors considered very important	Percentage of decision makers
Quality of teachers	88%
School discipline	82%
Quality of principal	58%
School reputation	50%
Range of school facilities	35%
Distance from school	30%
Friends in the school	30%

Although a substantial number (30%) of respondents think that distance to school is important, and hence they have no reason to worry about this criterion under the current assignment process, it is interesting to note that 70% of respondents do not consider distance to school as an important factor in their ranking of schools. The current system are not able to cater to these people in its assignment. *A good assignment process should match all students to schools according to the true preferences of the students.*

Problem 4. Finally, since the examination results are released only in late November, the MOE can only obtain the second option form not earlier than early December. As such, the entire assignment process is conducted only in a span of two weeks as the assignments have to be made known to the students by the third week of December. An unfortunate event developed in December 1997 when the Ministry had to publicly apologize for an administrative error:

It has been discovered that there were some discrepancies in the posting of pupils within the Normal (academic) course; 3,278 pupils out of a total of 9,417 are affected. The ministry is re-posting the 3,278 pupils (*The Straits Times*, December 1997).

The ministry apologized for the error and said that the error occurred when wrong data was entered into its computers during the assignment process. The human error was discovered only when "the ministry checked again and the parents gave feedback." This ultimately led to some resentment and confusion among the students and parents, as some of them had already bought the books and met the teachers in their previously assigned secondary school. The fact that the re-assignment was made barely a week before the commencement of the new school term aggravated the seriousness of the human errors. *As far as possible, the matching mechanism should be fully automated to mitigate the chances of human errors and reduce the time taken to post students to schools.*

4.2. Benefits of Stable Matching Mechanisms

The current posting method is appropriate for the old paradigm in which the schools are all alike and, thus, are viewed as resources to be allocated to the students. Also, assigning the students to schools according to merit is a simple and convenient way to solve

the matching problem. However, in recent years, this problem has assumed a new dimension as it is generally recognized that examination results alone are not a good way of assessing the merit of the students. Thus, elaborate mechanisms are now being put in place to assess the students in an all-round fashion. In the future, it will no longer be possible to have a *single* list that ranks all the students according to merit. Also, given that the schools are being encouraged to develop their own identity and strength, it is highly debatable whether the schools will all rank the students in the same way. Currently, besides affiliation, the assignment process cannot handle school-specific criteria. By moving to a new paradigm in which schools are allowed to devise their own criteria to rank the students, and assigning the students to schools via a stable matching mechanism, we can address some of the problems outlined in §4.1. Under this new mechanism, the limit on the number of choices each student is entitled to should be removed, since this is one of the root causes for the problems faced by the students in the current system. Although this will lead to an additional burden on the system (data collection and processing), the benefits that accrue will be well worth the effort.

With these changes, we can have a fully automated matching mechanism that can capture the true preferences of the students, and allow the schools to turn from passive to *active* participants in the assignment process. This immediately addresses Problems 2–4 outlined in §4.1. We argue next that in a stable matching mechanism implementation, it is unlikely that the students will benefit significantly from misrepresenting their preference lists. In particular, we analyze the strategic issues under the students-optimal and schools-optimal mechanisms.

Students-Optimal Mechanism. It follows from classical results in stable matching (cf. Roth 1982) that, in this model, the students have no incentive to misrepresent their preferences. However, the strategic behavior of the schools is not well understood. In fact, to the best of our knowledge, the question of finding the optimal cheating strategy for the schools is still open and appears to be a difficult problem. For the MOE problem, however, this is less of a concern, since proper management mechanisms can be instituted by

the MOE to prevent the schools from misrepresenting their preferences. For instance, by working with the schools to determine their strengths and core competence, the MOE can specify the profile of the prospective students that would fit into the schools' curriculum. This can be used to rank the students according to the various criteria listed down by the schools. This eliminates the need for the schools to submit a ranking list of individual students, thus giving them less scope for strategic manipulation.

Schools-Optimal Mechanism. As in the case of the students-optimal matching mechanism, the danger of schools manipulating the assignment process can be eliminated by appropriate management control. Note that even with the schools-optimal mechanism, it is known that truth revelation is not a dominant strategy for the schools, and so proper management control is still required to ensure that the schools do not misrepresent their preferences. However, it is still possible that some of the students will benefit from lying. In the rest of this section, we examine the extent to which the students benefit by misrepresenting their preferences. We use the algorithm described in §2.2 to obtain some insights into this problem.

We designed an experiment with $n = 60$ students, and k schools each with $60/k$ places, for $k = 1, 3, 5, 10, 20$. The preferences of the schools and the students are generated randomly in the many-to-one model. We then convert the problem into a one-to-one model, with each school represented as k individuals in the new model. All individuals corresponding to the same school will have identical preference lists. Also, for the students in the new model, the preference list is generated so that the individuals corresponding to the same school are clustered together in the students' preference list, capturing the fact that in the many-to-one model, the students basically provide the preference lists for the schools only. In this new one-to-one model (where students correspond to women and schools correspond to men), the preference lists are thus generated in a special manner. We next report the extent to which the women (i.e., students) can cheat in this new one-to-one model, under the men-optimal mechanism. Note that the number of women who can benefit by cheating in this case is not directly equivalent to the number of students who can

benefit by cheating in the many-to-one model, since the one-to-one problem contains artificial strict preferences over slots of the same school. Nevertheless, the numbers reported constitute an upperbound on the number of students who can benefit by cheating in the many-to-one model.

For each set of parameters, we ran 1,000 experiments. The following table shows the average number of women who benefited by cheating.

k	1	3	5	10	20
Percentage of women who benefit by cheating	8.77	11.00	11.09	10.22	4.675

The cumulative distribution of the number of women who benefited by cheating is shown in Figure 3.

For each value of k , the y axis shows the cumulative number of women who can benefit from cheating. For instance, for $k = 1$, in close to 900 out of 1,000 experiments, at least one woman benefits from cheating. Note that in the many-to-one model with only one school, clearly the students cannot benefit from cheating. This number drops to slightly over 800 for $k = 3$, slightly over 700 for $k = 5$, slightly over 500 for $k = 10$, and finally to around 200 for $k = 20$. Interestingly, the number of times that no women benefit by cheating (i.e., optimal strategy is for all the women to tell the truth) seems to increase as k increases (From 100 for $k = 1$ to 800 for $k = 20$). Translated to the many-to-one model, it indicates that, for instance, in the case $k = 20$ (i.e. 20 schools and 60 students), the chances that no students can benefit from cheating is *at least* 80%. It indicates that if there are relatively many schools in the many-to-one model compared to the total number of students, cheating becomes increasingly impossible.


The simulation experiments also showed an interesting feature: The proportion of women who can benefit from cheating hovers around 10% for all choices of k , even though the preference lists are generated in a special manner. The numbers are similar to those in the earlier experiment when the preference lists for the one-to-one model are generated in a random fashion.

5. Concluding Remarks

Our study here focused on two main areas: In the first part, we studied the theoretical aspects of the stable marriage problem. Specifically, we addressed the strategic issues in the Gale-Shapley model (i.e., a one-to-one matching model in which each participant submits complete preference lists)—very little was known about the strategic issues for this model prior to our work. Assuming that the men-optimal mechanism is used, we derived an optimal cheating strategy for the women. We also showed that the strategic behavior of the women can be very different under the models with and without rejection. We saw that in our model, an optimal cheating strategy for a woman did not always guarantee her women-optimal partner. Moreover, in sharp contrast to the rejection model, we showed that the chances of a woman benefitting by cheating are small.

The second part of our study emphasized the importance of adopting stable matching mechanisms in the context of the Singapore Secondary School posting exercise. There is an enormous amount of literature on the college admissions problem, which sheds light on the viability of the matching mechanisms we proposed. For instance, a variant of the student-propose mechanism suggested here has recently been recommended to be the official posting method to assign interns to hospitals in the United States (Roth and Peranson 1997). The high participation rate in this interns/hospitals market, despite the fact that participation is voluntary, shows the relevance and importance of ensuring stability. For the Singapore posting exercise, however, participation rate is not the central theme of the problem, since the students do not have a choice in participation. Interestingly, exploiting this unique feature of the problem (that the student cannot say “no”), we showed that a stable matching mechanism is still appropriate in this context. The current assignment process used by the MOE seems to place undue emphasis on the parents making intelligent, “realistic” choices; this is troublesome for the parents, especially because they have to make their decisions before their children appear for the primary school leaving examination. We argued that adopting a stable matching mechanism results in a better system. We showed further that regardless of

Figure 3 Cheating in the Many-to-One Model



the particular stable matching mechanism used (the students-optimal or the schools-optimal mechanism), there is no significant incentive for the students to cheat, under the perfect information scenario. Since the students have little reason to cheat in the two extremal stable matching mechanisms, it is natural to suspect that this result carries over to all the “intermediate” stable matchings as well; we leave this as a challenging problem for future research.

As for the strategic behavior of the schools, the adoption of any stable matching mechanism would give incentives for some schools to falsify their preferences or even to falsify the number of places available (see Sonmez 1997 on how the schools can manipulate capacities to obtain a better match under both students-optimal and schools-optimal mechanisms). In the Singapore context, however, the latter is relatively unimportant as the schools are tightly regulated by the government, and the number of places available are determined by the Ministry. Furthermore, the ranking of the students, according to the schools, is usually done by some broad selection criteria (e.g., examination results, affiliation, etc.) rather than at an individual level. Hence the issues of manipulation (both rank and capacity) on the schools’ part is less of a concern in this context.

To summarize, our analysis of the strategic issues suggests that stable matching mechanisms can be used to address most of the shortcomings of the current assignment process. In particular, it can be used to address the key concern of most parents: “How do I choose the schools if I do not know my child’s standing at the national level?” The answer to that, regardless of the stable matching mechanism used, is: “Just tell the truth.”

Acknowledgments

The authors thank the editor and the two anonymous referees for pointing many errors and omissions in earlier versions of this paper. In particular, the referees made several suggestions that improved the paper significantly, both in form and substance. This research was supported by the National University of Singapore under grant RP 3970021 and a fellowship from the Singapore-MIT Alliance program.

References

- Dubins, L.E., D.A. Freedman. 1981. Machiavelli and the Gale-Shapley algorithm. *Amer. Math. Monthly* **88** 485–494.
- Gale, D., L.S. Shapley. 1962. College admissions and the stability of marriage. *Amer. Math. Monthly* **69** 9–15.
- , M. Sotomayor. 1985a. Some remarks on the stable matching problem. *Discrete Appl. Math.* **11** 223–232.

- , ———. 1985b. Ms Machiavelli and the stable matching problem. *Amer. Math. Monthly* **92** 261–268.
- Gardenfors, P. 1975. Match making: Assignments based on bilateral preferences. *Behavioural Sci.* **20** 166–173.
- Gusfield, D., R.W. Irving. 1989. *The Stable Marriage Problem: Structure and Algorithms*. MIT Press, Cambridge, MA.
- Irving, R.W. 1994. Stable marriage and indifference. *Discrete App. Math.* **48** 261–272.
- Kara, T., T. Sonmez. 1996. Nash implementation of matching rules. *J. Econom. Theory* **68** 425–439.
- Knuth, D.E. 1976. *Marriages Stables*. Les Presses de l'Université de Montreal, Montreal, Canada.
- McVitie, D.G., L.B. Wilson. 1970. Stable marriage assignments for unequal sets. *BIT* **10**.
- , ———. 1971. The stable marriage problem. *Comm. ACM* **10** 486–490.
- Pittel, B. 1989. The average number of stable matchings. *SIAM J. Discrete Math.* 530–549.
- Roth, A.E. 1982. The economics of matching: Stability and incentives. *Math. Oper. Res.* **7** 617–628.
- . 1984a. The evolution of the labor market for medical interns and residents: A case study in game theory. *J. Political Econom.* **92** 991–1016.
- . 1984b. Stability and polarization of interests in job matching. *Econometrica* **52** 47–57.
- . 1985. The college admissions problem is not equivalent to the marriage problem. *J. Econom. Theory* **36** 277–288.
- . 1990. New physicians: A natural experiment in market organization. *Science* **250** 1524–1528.
- , E. Peranson. 1997. The effects of the change in the NRMP matching algorithm. *J. Amer. Medical Assoc.* **278** 729–732.
- , U. Rothblum. 1999. Truncation strategies in matching markets: In search of advice for participants. *Econometrica* **67** (January), 21–43.
- , M. Sotomayor. 1991. *Two-sided Matching: A Study in Game-theoretic Modeling and Analysis*. Cambridge University Press, Cambridge, MA.
- Sonmez, Tayfun. 1997. Manipulation via capacities in two-sided matching markets. *J. Econom. Theory* **77** 197–204.
- Tadenuma, K., M. Toda. 1998. Implementable stable solutions to pure matching problems. *Math. Social Sci.* **35**(2) 121–132.
- Teo, C.P., J. Sethuraman. 1998. The geometry of fractional stable matching and its applications. *Math. Oper. Res.* **23**(4) 874–891.

Accepted by Bruce L. Golden; received January 1999. This paper was with the authors 10 months for 2 revisions.