

Preference elicitation in matching markets via interviews: A study of offline benchmarks

Baharak Rastegari
School of Computing Science
University of Glasgow
Glasgow, UK
baharak.rastegari@glasgow.ac.uk

Paul Goldberg
Dept. of Computer Science
University of Oxford
Oxford, UK
paul.goldberg@cs.ox.ac.uk

David Manlove
School of Computing Science
University of Glasgow
Glasgow, UK
david.manlove@glasgow.ac.uk

ABSTRACT

The *stable marriage problem* and its extensions have been extensively studied, with much of the work in the literature assuming that agents fully know their own preferences over alternatives. This assumption however is not always practical (especially in large markets) and agents usually need to go through some costly deliberation process in order to learn their preferences. In this paper we assume that such deliberations are carried out via *interviews*, where an interview involves a man and a woman, each of whom learns information about the other as a consequence. If everybody interviews everyone else, then clearly agents can fully learn their preferences. But interviews are costly, and we may wish to minimize their use. It is often the case, especially in practical settings, that due to correlation between agents' preferences, it is unnecessary for all potential interviews to be carried out in order to obtain a stable matching. Thus the problem is to find a good strategy for interviews to be carried out in order to minimize their use, whilst leading to a stable matching. One way to evaluate the performance of an interview strategy is to compare it against a naïve algorithm that conducts all interviews. We argue however that a more meaningful comparison would be against an optimal offline algorithm that has access to agents' preference orderings under complete information. We show that, unless $P=NP$, no offline algorithm can compute the optimal interview strategy in polynomial time. If we are additionally aiming for a particular stable matching (perhaps one with certain desirable properties), we provide restricted settings under which efficient optimal offline algorithms exist.

Categories and Subject Descriptors

Theory of Computation [**Algorithmic game theory and mechanism design**]: Algorithmic game theory

General Terms

Algorithms, Economics, Theory

Keywords

Two-sided matching, preferences, interviews

Appears at: 3rd Workshop on Exploring Beyond the Worst Case in Computational Social Choice. Held as part of the 15th International Conference on Autonomous Agents and Multiagent Systems. May 9th-10th, 2016. Singapore.

1. INTRODUCTION

Two-sided matching markets model many practical settings, such as corporate hiring and university admission [21, 16]. The classical *stable marriage problem* is perhaps the most widely studied matching problem in this class, where participants are partitioned into two disjoint sets – men and women – and each participant on one side of the market wishes to be matched to a candidate from the other side of the market, and has preferences over potential matches. A matching is called stable if no pair of participants would prefer to leave their assigned partners to pair with each other. Gale and Shapley's seminal paper [4] proposed a polynomial-time algorithm for finding a stable matching. The books by Knuth [13], Gusfield and Irving [6], Roth and Sotomayor [21], and Manlove [16] provide excellent introductions and surveys.

A key assumption in much of this literature is that all market participants know their full preference orderings. The classical Gale-Shapley (GS) algorithm [4] and its variants require participants' preferences as input. This assumption is reasonable in some settings. However, as markets grow large (e.g., in the hospitals-residents matching market [20, 8] or college admission market [4, 22]) it quickly becomes impractical for participants to assess their precise preference rankings. Instead, participants usually start out with some partial knowledge about their preferences and need to perform some deliberation in order to learn their precise preference ordering. In this paper we assume that deliberations are carried out via *interviews*, where an interview is a unit operation that involves one agent from each side of the market and is informative to both participants. For example, in the hospitals-residents problem (which models the entry-level labor market in which graduating medical students, or residents, are seeking to be assigned to hospital posts), hospitals are likely to be able to identify their "top-tier" residents, "second-tier" residents and so on, and in order to rank the residents in each tier they need to interview them. An interview between a hospital h and a resident r yields information about the qualities of each party to the other. Thus we initially assume that each agent's preference list is in general expressed in terms of a partial order, and after an agent has interviewed ℓ members of the opposite side of the market, he/she has discovered enough information to rank those elements in strict order.

In order to be able to use the GS algorithm to find a stable matching in this setting, a naïve solution is for each participant to conduct all potential interviews and fully learn their preferences. Interviews however are usually costly both

in terms of time, mental energy, and money, therefore we wish to minimize their usage. Indeed, the naïve approach may impose unnecessary deliberation. For example, in the hospitals-residents problem, one expects some degree of correlation across hospitals in the assessment of the most desirable residents, and likewise residents are expected to have correlated views (at least to some extent) on the desirability of hospitals. Therefore, it is expected that more desirable residents get matched to more desirable hospitals and so on. It is then not hard to see that it is wasted effort if a top-tier resident is to interview low-tier hospitals, or a second-tier resident is to interview top-tier hospitals. For a concrete example, consider a setting with four residents and four hospitals where each hospital can admit (at most) one resident. Assume that residents r_1 and r_2 are top-tier residents and r_3 and r_4 are second-tier residents. Likewise assume that h_1 and h_2 are top-tier hospitals and h_3 and h_4 are second-tier hospitals. The preference lists of all agents are correlated according to these hierarchies, although each agent’s individual strict ranking (initially unknown) within these hierarchies may differ. It is not hard to verify that no matter what the true (initially unknown) preference orderings of the participants are, under a stable matching r_1 and r_2 each gets matched to either h_1 or h_2 , and r_3 and r_4 each gets matched to either h_3 or h_4 . Thus an interview between r_1 and either h_3 or h_4 is unnecessary, for example.

Unfortunately, we cannot always avoid unnecessary interviews. For example consider a setting with two residents and two hospitals, where initially agents have no information on their preference orderings and hence cannot compare the two alternatives. W.l.o.g. assume that h_1 interviews both residents and learns that it prefers r_1 to r_2 . If r_1 additionally interviews h_2 and learns that he prefers h_1 to h_2 , then a stable matching μ is found after 3 interviews, in which h_i is matched to r_i ($1 \leq i \leq 2$). Now imagine that r_1 instead learns that he prefers h_2 to h_1 . It is easy to verify that the identity of a stable matching is not yet revealed and hence more interviews are required. The only remaining interview is between h_2 and r_2 after which one can definitely identify a stable matching. If h_2 learns that it prefers r_2 to r_1 then μ is a stable matching. Assume that r_2 also prefers h_2 to h_1 . In this case the interview between h_1 and r_1 is unnecessary as the three other interviews would have provided enough information – that h_2 and r_2 are each others’ top choice – for μ to be identified as a stable matching. However, a priori we can not always rule the interview between h_1 and r_1 as unnecessary; for example when h_2 ’s top choice is r_2 but r_2 ’s top choice is h_1 .

Any interviewing strategy leads to refinements of the partial orders contained in the original problem instance that represented uncertainty over the true preferences. A key aim could be to carry out sufficient interviews so as to arrive at an instance that admits a super-stable matching μ . Super-stability will be defined formally in the next section, but informally it ensures that μ will be stable regardless of how the remaining uncertainty is resolved. The original instance need not admit a super-stable matching (see [9] for an example) but we are guaranteed that a super-stable matching is always achievable (e.g., by conducting all possible interviews, we will arrive at a strictly-ordered instance, where super-stability and classical stability become equivalent, and the existence of a stable matching is assured [4]).

Thus our aim is to find a good strategy that conducts as

few interviews as possible so as to obtain a refined instance that admits a super-stable matching. In general any such strategy will be an online algorithm, since the next interview to be carried out might depend on the results of previous ones.

This leads to the question of how to evaluate the performance of any given interview strategy. One could compare it against the naïve algorithm described above that conducts all interviews. We argue however, by analogy with online algorithms and their competitive ratio, that it makes more sense to compare it against an optimal “offline” algorithm. Here, the optimal offline algorithm has access to agents’ preference orderings under full information and has to compute the optimal (i.e., minimum) number of interviews required in order to reach an information state under which it can identify a super-stable matching. In this paper we show that unless $P=NP$, no offline algorithm can compute an optimal interview strategy in polynomial time.

Some stable matchings have desirable properties, and we may be interested in refining the preferences further so as to obtain such matchings. For example, in the *man-optimal stable matching*, each man has the best partner that he could obtain in any stable matching, whilst the *woman-optimal stable matching* has a similar optimality property for the women. As described above, after a certain number of interviews we may reach an instance that admits a super-stable matching μ . But by carrying out more interviews, some men, for example, may end up with better partners than they had in μ . This would be the case if μ is not the man-optimal stable matching in the instance with the strict (true underlying) preferences.

If we wish to evaluate the performance of an online algorithm that aims for potential improvements in men’s partners even after a super-stable matching has been identified, then a suitable offline benchmark for the competitive ratio would be the minimum number of interviews required to refine the original instance so as to make a *specified* matching super-stable. In this paper we show that, whilst this problem is NP-hard in general, there are restricted cases that are solvable in polynomial time.

Related work.

Until very recently, the problem of incremental preference elicitation has received little attention. Several works in the past few years however have addressed this problem from different angles [12, 14, 1, 3, 19, 7, 2]. Those closest in spirit to ours are [1, 19, 2].

In [19] the authors introduced a stable matching model in which participants start out with incomplete information about their preferences, in the form of partially ordered sets, and are able to refine their knowledge by performing interviews. They investigated the problem of minimizing the number of interviews required to find a matching that is stable w.r.t. the true underlying strict preference ordering and additionally is optimal for one side of the market. They presented several results among which are the following two: (i) finding a minimum certificate, that is a set of partial preferences that supports an optimal (for one side of the market) stable matching is NP-hard, and (ii) in a setting where participants on one side of the market have the same partially ordered preferences, an optimal interview policy can be found in polynomial time.

In [1], the authors studied a setting where deliberation is

in the form of pairwise comparison queries (that is, a query leads to strict order of preference being determined over two acceptable agents for a given agent). They proposed a method for finding approximately stable matchings, using minimax regret as a measure, while keeping the number of required comparisons relatively low. In [2] the authors combined the comparison query model of [1] with the interview model of [19] and introduced a unified model where both types of elicitation can take place. They provided an efficient (polynomial-time) scheme for generating queries and interviews, and examined the effectiveness of their scheme via empirical evaluation including comparison against the polynomial-time algorithm of [19] for the restricted setting in which participants on one side of the market have the same partially ordered preferences.

Our work is also related to the body of literature studying variants of stability defined in settings where agents' preferences may include ties. As discussed above, super-stable matchings are relevant in the context of incomplete preference information, because they are stable no matter which refinements represent the true (strict) preferences. Polynomial time algorithms have been proposed for finding a super-stable matching, or reporting that none exists, in various two-sided matching markets [9, 15, 10, 18].

In the next section we provide definitions of notation and terminology, leading to formal statements of the problems under consideration in this paper. A roadmap of the remaining sections is then given at the end of Section 2.

2. PRELIMINARY DEFINITIONS AND RESULTS

SMPI, SMTI, *and levels of stability.*

In an instance of the *Stable Marriage problem with Partially ordered preferences and Incomplete lists* (SMPI), there are two sets of agents, namely a set of men $M = \{m_1, m_2, \dots, m_n\}$, and a set of women $W = \{w_1, w_2, \dots, w_n\}$. We assume without loss of generality that $|M| = |W|$ (we can easily reduce the case where the two sets are of different sizes to our setting). Let $[i]$ denote the set $\{1, 2, \dots, i\}$. We use the term *agents* when making statements that apply to both men and women, and the term *candidates* to refer to agents on the opposite side of the market to that of an agent under consideration. Each agent a finds a subset of candidates acceptable – we refer to these as a 's *acceptable candidates*. An agent a 's preferences over his/her acceptable candidates need not be strict. That is, given two candidates, a might not be able to compare them against each other. We denote by p_{m_i} and p_{w_j} the partial orders that represent the preferences of m_i and w_j , respectively. We let $p_{M,W} = (p_{m_1}, \dots, p_{m_n}, p_{w_1}, \dots, p_{w_n})$ and call $p_{M,W}$ a *partial preference ordering profile*.

Let $I = (M, W, p_{M,W})$ be an instance of SMPI, let a be an agent and let c_1 and c_2 be two acceptable candidates for a in I . We say that a *strictly prefers* c_1 to c_2 if $(c_1, c_2) \in p_a$, and we say that a *cannot compare* c_1 and c_2 (or that a finds c_1 and c_2 *incomparable*) if $(c_1, c_2) \notin p_a$ and $(c_2, c_1) \notin p_a$. We sometimes use the graph-theoretic representation of p_a where candidates in p_a correspond to vertices and there is an arc from a candidate c_i to a candidate c_j if and only if $(c_i, c_j) \in p_a$.

An instance $I' = (M, W, p'_{M,W})$ of SMPI is a *refinement*

of I if for each agent a , any strict total order that is a linear extension of p'_a is also a linear extension of p_a . We may also refer to $p'_{M,W}$ being a refinement of $p_{M,W}$ (or indeed I) using the same definition. Also we can define p'_a being a refinement of p_a for some specific agent a similarly.

OBSERVATION 1. *Given two instances I and I' of SMPI, I' is a refinement of I if and only if the following condition holds: for each agent a and every two candidates c_1 and c_2 acceptable to a , if $(c_1, c_2) \in p_a$ then $(c_1, c_2) \in p'_a$.*

A well studied special case of SMPI is the *Stable Marriage problem with Ties and Incomplete lists* (SMTI), in which incomparability is transitive and is interpreted as indifference. In SMTI, each agent has a partition of acceptable candidates into *indifference classes* or *ties* such that he or she is indifferent between the candidates in the same indifference class, but has a strict preference ordering over the indifference classes. In an instance of SMTI, let C_t^a denote the t -th indifference class of agent a , where $t \in [n]$. We assume that $C_t^a = \emptyset$ implies $C_{t'}^a = \emptyset$ for all $t' > t$. The *Stable Marriage problem with Incomplete lists* (SMI) is the special case of SMTI in which each tie is of size one. Similarly the *Stable Marriage problem with Ties* (SMT) is the special case of SMTI in which each man finds each woman acceptable and vice versa.

Given an instance I of SMPI, a *matching* μ is a pairing of men and women such that each man is paired with at most one woman and vice versa, and no agent is matched to an unacceptable partner. If m and w are matched in μ then $\mu(m) = w$ and $\mu(w) = m$. We say that $\mu(a) = \emptyset$ if a is unmatched under μ . Different levels of stability can be defined in the context of SMPI [9, 15]. A *strong blocking pair* is an acceptable (man,woman) pair, each of whom is unmatched or strictly prefers the other to his/her partner. A *weakly stable matching* is a matching with no strong blocking pair. Every instance of SMPI admits a weakly stable matching [17]. An acceptable (man,woman) pair is a *weak blocking pair* if each member of the pair is either unmatched or strictly prefers the other to his/her partner or cannot compare the other with his/her partner, and one member of the pair is either unmatched or strictly prefers the other to his/her partner. A *strongly stable matching* is a matching with no weak blocking pair. Finally a *very weak blocking pair* is an acceptable (man,woman) pair, each of whom is unmatched or strictly prefers to other to his/her partner or cannot compare the other with his/her partner. A *super-stable matching* is a matching with no very weak blocking pair. It can be easily verified that a matching is super-stable if and only if it is weakly stable w.r.t. all strict total orders that are linear extensions of the given partial preference orderings [16, Lemma 3.2.4]. In instances of SMI, weak stability, strong stability and super-stability are all equivalent to classical stability.

Interviews to refine the partial orders.

In a given instance $I = (M, W, p_{M,W})$ of SMPI in this paper, we assume that $p_{M,W}$, the partial preference ordering profile, represents the agents' initial information state. That is, agents may not have enough information initially in order to rank their acceptable candidates in strict order. However in the problem instances that we will later define in this section, we will assume that each agent a has a strict preference ordering \succ_a over his or her acceptable candidates. This rep-

resents the true (and strict) underlying preferences over a 's acceptable candidates, although crucially, a may not (and in general will not) initially be aware of the entire ordering. We let $\succ_{M,W} = (\succ_{m_1}, \dots, \succ_{m_n}, \succ_{w_1}, \dots, \succ_{w_n})$ and call $\succ_{M,W}$ a *strict (true underlying) preference ordering profile*. The task of the agents is to learn enough information about their acceptable candidates in order to refine their preferences, in a manner consistent with $\succ_{M,W}$, to obtain an SMPI instance I' that admits a super-stable matching μ (thus μ will be stable with respect to $\succ_{M,W}$).

Following the model introduced in [19], we assume that instances can be refined through *interviews*. Each interview pairs one man m with one woman w . An interview is informative to both parties involved. Hence saying “ m interviews w ” is equivalent to saying “ w interviews m ”. When agent a interviews ℓ candidates, this results in a new refined SMPI instance which is exactly the same as I except that a now has a strict preference ordering over the ℓ interviewed candidates.

Notice that if an agent interviews only one candidate, no refinement takes place. Note also that not all refinements of I can be reached by a set of interviews. For example, suppose that in I we have one man m_1 and three women w_1 , w_2 , and w_3 . Suppose m_1 finds the three women acceptable and incomparable. Assume that in I' man m_1 prefers w_1 to both w_2 and w_3 , and cannot compare w_2 and w_3 . It is easy to see that I' is a refinement of I , but no set of interviews can reach I' : for m_1 to learn that he prefers w_1 to the other two women he must interview all three women, but then he will have a strict preference ordering over the three of them.

We say that an SMPI instance I' is an *interview-compatible refinement* of an SMPI instance I if I' can be refined from I using interviews. We now show that interview-compatible refinements can be recognized easily.

PROPOSITION 2. *Let I and I' be two instances of SMPI. We can determine in $O(n^3)$ time whether I' is an interview-compatible refinement of I .*

PROOF. To verify whether I' is a refinement of I , it is sufficient to check whether the condition of Observation 1 holds. With a suitable data structure, we can do this in $O(n^3)$ time. For each agent a identify the edges present in p'_a that are not in p_a , and let $S(a)$ be the set of candidates in p'_a that form an endpoint of at least one such edge. For I' to be an interview-compatible refinement of I , it is necessary and sufficient that, for every a , $S(a)$ forms a complete subgraph in the undirected graph corresponding to p'_a . This can be tested in $O(n^3)$ time overall. \square

Let I' be an SMPI instance that is an interview-compatible refinement of a given SMPI instance I . We define the *cost* of I' given I to be the minimum number of interviews required to refine I into I' . The following proposition shows how to compute this cost efficiently.

PROPOSITION 3. *Let I be an SMPI instance and let I' be an interview-compatible refinement of I . We can determine in $O(n^3)$ time the cost of I' given I .*

PROOF. We identify the set of interviews T that refines I into I' as follows. Initially $T = \emptyset$. For each agent a and every two candidates c_1 and c_2 , if a cannot compare c_1 and c_2 under I , but prefers one to the other under I' , a must have interviewed both c_1 and c_2 . Add both of these interviews to T . Notice that we might have already accounted

$$\begin{array}{ll} r_1 : & h_1 \ h_2 & h_1 : & r_2 \ r_1 \\ r_2 : & h_2 \ h_1 & h_2 : & r_2 \ r_1 \end{array}$$

Figure 1: A stable marriage problem instance.

for one or both of these interviews. However since T is a set, no interview is going to be included in T more than once. With a suitable data structure, the aforementioned procedure terminates in $O(n^3)$ time overall, and once it does, $|T|$ denotes the cost of I' . \square

Problem definitions.

The motivating problem is as follows: given an instance $I = (M, W, p_{M,W})$ of SMPI, find an interview-compatible refinement I' of minimum cost such that I' admits a super-stable matching. Since the result of one interview might influence which interview/s to carry out next, any strategy for carrying out interviews should be regarded as an online algorithm.

In fact there may be no online algorithm that can guarantee to provide an optimal solution in all cases. To see this, let us return to the example of Section 1 involving two residents and two hospitals, and suppose that initially everyone finds the two agents on the other side of the market acceptable and incomparable. Let the true underlying preferences be given by Figure 1 (here, preference lists are ordered from left to right in decreasing order of preference). Any online interviewing strategy must start with one interview in the absence of any knowledge; w.l.o.g. suppose that the first interview involves r_1 and h_1 . Then it may be verified that the algorithm is bound to use 3 more interviews before a super-stable matching can be found. (If one interview does not take place then two agents on opposite sides of the market cannot compare the two candidates in their preference list; it then follows that each of the two possible matchings would be blocked according to super-stability.) On the other hand the interview involving r_1 and h_1 was unnecessary and an optimal strategy uses only 3 interviews.

Towards computing bounds for the competitive ratio of an online algorithm, the offline scenario is of interest, and that is what we consider in what follows. In the offline case, the mechanism designer is given $\succ_{M,W}$, the strict (true underlying) preference ordering profile of the agents, and would like to compute an optimal interviewing schedule, i.e., an interview-compatible refinement I' of I , such that $\succ_{M,W}$ refines I' . This is reflected in the definition of the following problem, named MIN-ICR, which is an abbreviation for “Minimum-cost Interview Compatible Refinement problem”.

DEFINITION 4. *An instance of MIN-ICR comprises a tuple $(I, \succ_{M,W})$, where I is an instance of SMPI and $\succ_{M,W}$ is a strict preference ordering profile that refines I . The problem is to find an interview-compatible refinement I' of I such that (i) $\succ_{M,W}$ refines I' , (ii) I' admits a super-stable matching, and (iii) I' is of minimum cost amongst interview-compatible refinements that satisfy (i) and (ii).*

The decision version of MIN-ICR is defined as follows.

DEFINITION 5. *An instance of ICR-DEC comprises a tuple $(I, \succ_{M,W}, K)$, where I is an instance of SMPI, $\succ_{M,W}$*

is a strict preference ordering profile that refines I , and K is a non-negative integer. The problem is to decide whether there exists an interview-compatible refinement I' of I , with cost at most K , such that $\succ_{M,W}$ refines I' and I' admits a super-stable matching.

As discussed in Section 1, it is sometimes the case that we aim for a particular matching, stable under $\succ_{M,W}$, that has some desirable properties, for example the woman-optimal stable matching. The offline problem can then be viewed as a restricted variant of MIN-ICR where, in addition to I and $\succ_{M,W}$, we are also equipped with a matching μ . This is reflected in the definition of the following problem, named MIN-ICR-EXACT, which is an abbreviation for “Minimum-cost Interview Compatible Refinement problem with Exact matching”.

DEFINITION 6. *An instance of MIN-ICR-EXACT comprises a tuple $(I, \succ_{M,W}, \mu)$, where I is an instance of SMPI, $\succ_{M,W}$ is a strict preference ordering profile that refines I , and μ is a matching that is weakly stable w.r.t. $\succ_{M,W}$. The problem is to find an interview-compatible refinement I' of I , such that (i) $\succ_{M,W}$ refines I' , (ii) μ is super-stable in I' , and (iii) I' is of minimum cost amongst interview-compatible refinements of I that satisfy (i) and (ii).*

The decision version of MIN-ICR-EXACT, called ICR-EXACT-DEC, is then defined analogously to the way that ICR-DEC was obtained from the definition of MIN-ICR.

The remainder of this paper is organized as follows. In Section 3 we first show that ICR-DEC is NP-complete even under quite restricted settings. The proof is by reduction from Vertex Cover. We also leverage the same proof to show that ICR-EXACT-DEC is also NP-complete. Then in Section 4 we provide a reverse reduction, from Vertex Cover to ICR-EXACT-DEC, and utilize it to show that MIN-ICR-EXACT is polynomial-time solvable for several restricted settings. Some concluding remarks are presented in Section 5.

3. NP-COMPLETENESS RESULTS

We show that ICR-DEC is NP-complete even if I is an instance of SMTI in which each indifference class is of size at most 3. Further, we prove that ICR-DEC is NP-complete even for SMT instances, and even if all men are indifferent between all women. We first provide a lemma that will come in handy in proving our claims; the proof is straightforward, and is omitted.

LEMMA 7. *Let $G = (V, E)$ be an undirected graph where for each vertex v , $\deg(v) \leq 3$. We can direct the edges in E such that for each v , $\deg^+(v) \leq 2$ and $\deg^-(v) \leq 2$.*

Unlike many problems that are NP-complete, the membership of ICR-DEC in the class NP is not trivial. Hence, we provide a proof via the following lemma.

LEMMA 8. *ICR-DEC is in NP.*

PROOF. To prove this, it is sufficient to show that given SMPI instances I and I' , a strict preference profile $\succ_{M,W}$ and an integer K , we can decide in polynomial time whether (i) I' is an interview-compatible refinement of I , (ii) I' has cost at most K , (iii) $\succ_{M,W}$ refines I' , and (iv) I' admits a super-stable matching.

Both (i) and (ii) are established by Propositions 2 and 3 respectively. For (iii), it is straightforward to check in polynomial time whether $\succ_{M,W}$ refines I' .

Finally to establish (iv), we can use the polynomial time algorithm of [18], SUPER-SMP, to decide whether I' admits a super-stable matching or not. \square

We show that ICR-DEC is NP-hard by reducing from the decision version of the Vertex Cover problem (VC). VC is defined as follows: given a graph $G = (V, E)$ and an integer K , decide whether G admits a vertex cover of size at most K . VC is NP-complete even if each vertex has degree at most 3 [5]; let VC-3 denote this restriction. We denote by MIN-VC the optimization version of VC, that is the problem of finding a minimum vertex cover in a given graph G .

THEOREM 9. *ICR-DEC is NP-complete even for SMTI instances in which each indifference class has size at most 3.*

PROOF. By Lemma 8, ICR-DEC belongs to NP. To show NP-hardness, we reduce from VC-3. Let the undirected simple graph $G = (V, E)$ be given such that $\deg(v) \leq 3$, $\forall v \in V$. Let $V = \{v_1, \dots, v_n\}$. Let $G' = (V, E')$ be a digraph where (i) $\forall (v_i, v_j) \in E'$, $(v_i, v_j) \in E$, (ii) $\forall (v_i, v_j) \in E$, either $(v_i, v_j) \in E'$ or $(v_j, v_i) \in E'$ (but not both), and (iii) $\forall v \in V$, $\deg^+(v) \leq 2$ and $\deg^-(v) \leq 2$. Note that by Lemma 7 such a graph G' exists. We create an instance $I = (M, W, p_{M,W})$ of SMTI as follows:

- For each vertex $v_i \in V$ there is a man $m_i \in M$ and a woman $w_i \in W$. That is $M = \{m_i | \forall v_i \in V\}$ and $W = \{w_i | \forall v_i \in V\}$.
- Each man m_i finds acceptable w_i and all women w_j such that $(v_i, v_j) \in E'$. Moreover, each man m_i is indifferent between all his acceptable women.
- Each woman w_i finds acceptable m_i and all men m_j such that $(v_j, v_i) \in E'$. Moreover, each woman w_i is indifferent between all her acceptable men.

Note that the total length of the men’s preference lists is equal to $|E| + |V|$. Also note that as $\deg^+(v) \leq 2$ and $\deg^-(v) \leq 2$ for all $v \in V$, hence the indifference classes are of size at most 3. Let $\succ_{M,W}$ be a strict preference ordering under which m_i and w_i rank each other at the top of their preference lists. We prove that G has a vertex cover of size at most K if and only if there exists a refinement I' of I , of cost at most $K' = K + |E|$, such that $\succ_{M,W}$ refines I' and I' admits a super-stable matching. Notice that $\succ_{M,W}$ admits only one stable matching, that being $\mu = \{(m_i, w_i) | \forall i\}$. This implies that I' must admit exactly one super-stable matching, that being μ .

Proof of the only-if direction: Assume that G has a vertex cover C of size $k \leq K$. We show that there is a refinement I' of cost $k' = k + |E| \leq K'$ such that $\succ_{M,W}$ refines I' and μ is super-stable in I' . We create I' as follows. For each vertex $v_i \in C$ refine I such that both m_i and w_i now have strict preferences as in $\succ_{M,W}$. This refinement can of course be achieved by having both m_i and w_i interview all candidates in their lists; this includes an interview between m_i and w_i . Notice that since interviews are informative to both parties involved, partial refinements in the preference orderings of those persons whose corresponding vertices are not in C must have taken place as well. For example, consider a case in which $v_i, v_j \in C$, $v_k \notin C$, and $(v_i, v_k), (v_j, v_k) \in E'$.

Then both m_i and m_j interview w_k and hence w_k must now rank m_i and m_j in strict order of preferences as in $\succ_{M,W}$. An interview is either between same indexed agents, e.g., between m_i and w_i , or between agents of different indices, e.g., between m_i and w_j where $i \neq j$. We refer to an interview of the former type as a *same-index* interview and an interview of the latter type as a *different-index* interview. The total number of interviews performed by all agents is going to be k same-index interview plus a number of different-index interviews. The number of different-index interviews under any refinement can be at most $|E|$, and under our proposed refinement is exactly $|E|$ (since C is a vertex cover). Therefore the total number of interviews is exactly $k + |E|$. It remains to show that μ is a super-stable matching in I' . We call a (*man, woman*) pair a *fixed pair* if they are matched in every stable matching of every strict order refinement of I' . We show that (m_i, w_i) is a fixed pair for all $i \leq n$, hence proving that μ is the only stable matching in every strict order refinement of I' and therefore definitely a super-stable matching in I' . Take any pair (m_i, w_i) such that $v_i \in C$. By our construction, m_i and w_i rank each other at top, so clearly (m_i, w_i) is a fixed pair. Now take any pair (m_j, w_j) such that $v_j \notin C$. Since v_j is not in the vertex cover, therefore it must be the case that all neighbors of v_j are in C . Thus, for any v_k that is a neighbor of v_j , it has been already established that (m_k, w_k) is a fixed pair. Therefore (m_j, w_j) is also a fixed pair. Moreover, neither m_j nor w_j can form a very weak blocking pair with a person they are not matched to.

Proof of the if direction: Assume that I has a refinement I' of cost $k' \leq K'$ such that $\succ_{M,W}$ refines I' and μ is super-stable in I' . We show that G admits a vertex cover of size at most $k' - |E| \leq K$. We first show that in order to arrive at I' , every agent should have interviewed every candidate s/he finds acceptable and to whom s/he is not matched. Assume for a contradiction that this is not the case. Take a pair (m_i, w_j) , acceptable to each other, who have not interviewed. Therefore, under I' , m_i is indifferent between w_j and w_i (to whom he is matched in μ), and w_j is indifferent between m_i and m_j (to whom she is matched in μ). Hence (m_i, w_j) constitutes a very weak blocking pair in μ under I' , a contradiction. We have established so far that every agent must have interviewed acceptable candidates to whom s/he is not matched, which means that each agent has interviewed all candidates in his/her list who have a different index from him/her. This amounts to the total of $|E|$ interviews. The only remaining interviews for which we have not yet accounted are those corresponding to matched pairs. Let C be a set of vertices such that vertex v_i is in C if and only if m_i and w_i have interviewed under I' . Notice that $|C| = k' - |E|$. Take any $v_j \notin C$. We show that all neighbors of v_j are in C , establishing that C is a vertex cover. Since $v_j \notin C$, it follows from the construction of C that m_j and w_j have not interviewed under I' . Assume for a contradiction that v_j has a neighbor, say v_k , who too is not in C . Therefore m_k and w_k have not interviewed under I' either. W.l.o.g. assume that $(v_j, v_k) \in E'$. (A similar argument applies if $(v_k, v_j) \in E'$.) Therefore m_j and w_k are acceptable to each other. Furthermore, since neither m_j nor w_k have interviewed their partners in μ , it is the case that m_j is indifferent between w_j (his partner in μ) and w_k , and w_k is indifferent between m_k (her partner in μ) and m_j . Therefore (m_j, w_k) constitutes a very weak blocking pair in

μ under I' , a contradiction. \square

We next show that ICR-DEC is also NP-complete under a different restricted setting by making small alterations to the proof of Theorem 9.

COROLLARY 10. *ICR-DEC is NP-complete even for SMT instances and even if agents on one side of the market are indifferent between all the candidates.*

PROOF. W.l.o.g. assume that all men are indifferent between all women. Modify the reduction presented in the proof of Theorem 9 as follows.

- For each vertex $v_i \in V$ there is a man m_i in M and a woman w_i in W . That is $M = \{m_i | \forall v_i \in V\}$ and $W = \{w_i | \forall v_i \in V\}$.
- Each man m_i finds all women acceptable and is indifferent between them.
- Each woman w_i finds all men acceptable and has two indifference classes. In the top indifference class are m_i and all men m_j such that $(v_i, v_j) \in E$. In the second indifference class are all other men.

Note that the total length of the women's first indifference classes is equal to $2|E| + |V|$. Let $\succ_{M,W}$ be a strict preference ordering under which m_i and w_i rank each other at the top of their preference lists. We prove that G has a vertex cover of size at most K if and only if there exists a refinement I' of I , of cost at most $K' = K + 2|E|$, such that $\succ_{M,W}$ refines I' and I' admits a super-stable matching. Notice that $\succ_{M,W}$ admits only one stable matching, that being $\mu = \{(m_i, w_i) | \forall i\}$. This implies that I' must admit exactly one super-stable matching, that being μ . Modify the proof of Theorem 9 as follows.

In the only-if direction: For each vertex $v_i \in C$ refine I such that m_i has a strict preference ordering, as in \succ_{m_i} , over women in $\{w_i\} \cup \{w_j | (v_i, v_j) \in E\}$ and w_i has a strict preference ordering, as in \succ_{w_i} , over men in $\{m_i\} \cup \{m_j | (v_i, v_j) \in E\}$. Consequently, for all v_j adjacent to v_i , m_i prefers w_i to w_j and likewise w_i prefers m_i to m_j . This refinement can be achieved by having m_i interview w_i and all w_j such that $(v_i, v_j) \in E$, and additionally having w_i interview all m_j such that $(v_i, v_j) \in E$. The number of different-index interviews under any refinement can be at most $2|E|$, and under our proposed refinement is exactly $2|E|$ (since C is a vertex cover). So the total number of interviews is exactly $k + 2|E|$. It remains to show that μ is a super-stable matching in I' . Assume for a contradiction that there exists a very weak blocking pair (m_i, w_j) .

- If $(v_i, v_j) \in E$, then v_i or v_j is in C . If $v_i \in C$ then m_i and w_i have interviewed and therefore $w_i \succ_{m_i} w_j$. If $v_j \in C$ then m_j and w_j have interviewed and therefore $m_j \succ_{w_j} m_i$. Both cases imply that (m_i, w_j) is not a very weak blocking pair, a contradiction.
- If $(v_i, v_j) \notin E$ then $m_j \succ_{w_j} m_i$, therefore (m_i, w_j) is not a very weak blocking pair, a contradiction.

In the if direction: We show that in order to arrive at I' , every man m_i should have interviewed all women w_j such that $(v_i, v_j) \in E$, and likewise every woman w_i should have interviewed all men m_j such that $(v_i, v_j) \in E$. The proof is

similar to that presented in the proof of Theorem 9. Hence we can conclude that at least $2|E|$ different-index interviews must have taken place in the refinement. The rest of the proof is similar to that presented in the proof of Theorem 9, with the difference that $|C| \leq k' - 2|E|$. \square

In the proof of Theorem 9, μ is the unique stable matching under $\succ_{M,W}$. Therefore, it follows from the proofs of Theorem 9 and Corollary 10 that ICR-EXACT-DEC is also NP-complete for the restrictions stated in those results.

COROLLARY 11. *ICR-EXACT-DEC is NP-complete even for SMTI instances, and even when each indifference class is of size at most 3. ICR-EXACT-DEC is also NP-complete even for SMTI instances and even if agents on one side of the market are indifferent between all the candidates.*

We remark that Theorem 4.4 of [19] implies that ICR-EXACT-DEC is NP-complete and, likewise, Corollary 11 implies Theorem 4.4 of [19]. However, Corollary 11 is stronger as it is stated for a more restricted setting.

4. POLYNOMIAL-TIME SOLVABLE VARIANTS

4.1 Preliminaries

In this section we explore the tractability of MIN-ICR-EXACT under various restricted settings. Recall that we have reduced from VC to ICR-DEC and ICR-EXACT-DEC in order to show that these problems are NP-hard. Here we present a reverse reduction, from ICR-EXACT-DEC to VC, that will come in handy in proving our claims.

Let an instance $(I, \succ_{M,W}, \mu)$ of ICR-EXACT-DEC be given. As μ is weakly stable w.r.t. I , it admits no strong blocking pair. If μ is not super-stable w.r.t. I , then μ must admit some very weak blocking pairs. We refer to such blocking pairs as *potential blocking pairs*. We distinguish between potential blocking pairs by the degree of choice one has when attempting to resolve them.

DEFINITION 12 (POTENTIAL BLOCKING PAIR (PBP)). *Given an ICR-EXACT-DEC instance $(I, \succ_{M,W}, \mu)$, a pair (m, w) is a potential blocking pair (PBP) if (m, w) is a very weak blocking pair under I . Each PBP (m, w) belongs to either of the following classes.*

- Potential Blocking Pair of Degree 1 (PBP-D1) if either m or w strictly prefers the other to his or her current partner under $\succ_{M,W}$.
- Potential Blocking Pair of Degree 2 (PBP-D2) if both m and w strictly prefer their partners to each other under $\succ_{M,W}$.

Let I' be an interview-compatible refinement of I . We say that a given potential blocking pair of I , (m, w) , is *resolved* under I' if $(\mu(m), w) \in p'_m$ or $(\mu(w), m) \in p'_w$.

If (m, w) is a PBP-D2, then it must be that m and w cannot compare each other and their current partners under I . Thus in order to resolve (m, w) it is sufficient, and necessary, that m or w learn his/her true preference ordering over his/her partner and the other side.

Let (m, w) be a PBP-D1 and assume that m strictly prefers w to $\mu(m)$ (the argument is similar if $m \succ_w \mu(w)$). Therefore, w must find m and $\mu(w)$ incomparable under I , or

(m, w) either blocks μ or is not a PBP, and $\mu(w) \succ_w m$, or (m, w) blocks μ . Furthermore, in order to resolve this PBP w has to learn that she prefers $\mu(w)$ to m .

In what follows we use *PBP*, *PBP₁*, and *PBP₂* to refer to the set of potential blocking pairs, and those of degree 1 and degree 2 respectively.

PROPOSITION 13. *Let $(I, \succ_{M,W}, \mu)$ be an instance of ICR-EXACT-DEC and I' be an interview-compatible refinement of I . Then μ is super-stable under I' if and only if all PBPs in I are resolved under I' .*

It is easy to see that for a potential blocking pair (m, w) to be resolved, at least one of m or w needs to interview both the other side and his or her current partner and conclude that s/he prefers his or her current partner to the other side. The next proposition then immediately follows.

PROPOSITION 14. *Let $(I, \succ_{M,W}, \mu)$ be an instance of ICR-EXACT-DEC and I' be an interview-compatible refinement of I . Then μ is super-stable under I' only if, for all $(m, w) \in PBP$, m and w have interviewed under I' .*

For each agent $a \in M \cup W$ let $PBP_1(a)$ denote the set of candidates c such that either (a, c) or (c, a) is in PBP_1 and $a \succ_c \mu(c)$.

LEMMA 15. *Let $(I, \succ_{M,W}, \mu)$ be an instance of ICR-EXACT-DEC and I' be an interview-compatible refinement of I . Then μ is super-stable under I' only if a has interviewed $\mu(a)$ under I' for all agents a where $PBP_1(a) \neq \emptyset$.*

PROOF. Assume for a contradiction that there exists an agent a where $PBP_1(a) \neq \emptyset$ and a has not interviewed $\mu(a)$. Therefore for every $c \in PBP_1(a)$ it is still the case that a cannot compare c and $\mu(a)$, and c prefers a to $\mu(c)$. Hence there exists at least one unresolved PBP under I' . \square

4.2 Reduction from ICR-EXACT-DEC to VC

Let $(I, \succ_{M,W}, \mu)$ be an instance of ICR-EXACT-DEC. Let $M' = \{m | PBP_1(m) \neq \emptyset \vee PBP_1(\mu(m)) \neq \emptyset\}$. Let $G(I, \mu) = (V, E)$ be an undirected graph whose vertices V correspond to matched pairs $(m, \mu(m))$. Let $PBP'_2 = \{(m, w) | (m, w) \in PBP_2, m \notin M', \mu(w) \notin M'\}$. Let there be an edge between any two vertices $(m, \mu(m))$ and $(m', \mu(m'))$ if $(m, \mu(m')) \in PBP'_2$ or $(m', \mu(m)) \in PBP'_2$. Remove any vertex with degree zero. Note that for any remaining vertex $(m, \mu(m))$ it is the case that $m \notin M'$.

THEOREM 16. *$G(I, \mu)$ has a vertex cover of size at most K if and only if there exists a refinement I' of I , of cost at most $K' = |PBP| + |M'| + K$, such that $\succ_{m,w}$ refines I' and μ is super-stable in I' .*

PROOF. Assume that $G(I, \mu)$ has a vertex cover C of size at most K . Let I' be a refinement of I under which the following interviews have taken place.

1. Each pair $(m, w) \in PBP$ interview each other – a total of $|PBP|$ interviews.
2. Each $m \in M'$ interviews his partner $\mu(m)$ – a total of $|M'|$ interviews.
3. Each pair $(m, \mu(m)) \in C$ interview each other – a total of K interviews.

The total number of interviews is then equal to $|PBP| + |M'| + K$. As a result of the above interviews, each agent a learns his or her strict preference ordering over the interviewed candidates, as in \succ_a . (Recall that the interviews are informative to both sides.) It is then easy to see that all PBP-D1's are resolved. It is also straightforward to see that for a PBP-D2 (m, w) , if either $m \in M'$ or $\mu(w) \in M'$, then (m, w) is resolved under I' . It remains to show that the remaining PBP-D2's, that is those in PBP'_2 , are resolved as well. Let (m, w) be such a PBP-D2. By the construction of $G(I, \mu)$, V includes $(m, \mu(m))$ and $(\mu(w), w)$ and there is an edge between these two vertices. As C is a vertex cover, at least one of $(m, \mu(m))$ or $(\mu(w), w)$ belongs to C . If $(m, \mu(m)) \in C$ then, following the results of the interviews, m prefers $\mu(m)$ to w under I' . (A similar argument holds for w if $(\mu(w), w) \in C$.) Therefore (m, w) is resolved under I' .

Conversely, assume that I admits an interview-compatible refinement I' of size at most K' such that μ is super-stable in I' . We show that $G(I, \mu)$ admits a vertex cover of size at most $K' - (|PBP| + |M'|)$. Let C be a set of vertices $(m, \mu(m))$ in V where m and $\mu(m)$ have interviewed under I' . Note that as we have removed all vertices of degree zero from $G(I, \mu)$, hence all remaining vertices are adjacent to at least one edge corresponding to a member of PBP'_2 . We show that C is a vertex cover and then prove an upper bound on the size of C .

C is a vertex cover: Let $((m, \mu(m)), (m', \mu(m')))$ be any edge in E . By the construction of $G(I, \mu)$, $(m, \mu(m'))$ or $(m', \mu(m))$ is in PBP'_2 . Assume that $(m, \mu(m')) \in PBP'_2$. (The argument for the case where $(m', \mu(m)) \in PBP'_2$ is similar.) As $(m, \mu(m'))$ is resolved under I' , either m prefers his partner to $\mu(m')$ under I' , or $\mu(m')$ prefers her partner to m under I' . If the former, then m must have interviewed $\mu(m)$ and hence $(m, \mu(m)) \in C$, and if the latter then $\mu(m')$ must have interviewed m' and thus $(m', \mu(m')) \in C$. Thus C is a vertex cover.

C is of size at most $K' - (|PBP| + |M'|)$: We prove this by computing a lower bound on the number of interviews that do not correspond to a vertex in C . It follows Proposition 14 that all PBPs must have interviewed, hence a total of $|PBP|$ interviews. It also follows Lemma 15 that each agent a with $PBP_1(a) \neq \emptyset$ must have interviewed his/her partner. Looking at this from men's perspective, all men m must interview $\mu(m)$ if $PBP_1(m) \neq \emptyset$ or $PBP_1(\mu(m)) \neq \emptyset$ – hence a total of $|M'|$ interviews. Recall that $(m, \mu(m)) \notin V$ if $m \in M'$. Therefore none of the interviews we have accounted for so far, a total of $|PBP| + |M'|$ interviews, correspond to a vertex in C . \square

Theorem 16 essentially tells us that an instance $(I, \succ_{M,W}, \mu)$ of MIN-ICR-EXACT is polynomial-time solvable if MIN-VC is polynomial-time solvable in $G(I, \mu)$. Equipped with this knowledge, we provide three different restricted settings under which ICR-EXACT-DEC, and hence MIN-ICR-EXACT, is solvable in polynomial time.

THEOREM 17. MIN-ICR-EXACT is solvable in polynomial time if one side has fully known strict preference ordering.

PROOF. Assume that women have strict preferences and the target matching is μ . Note that all PBPs must be of degree 1. Therefore $G(I, \mu)$ is an empty graph with vertex cover of size zero. It follows from Proposition 14 and

Lemma 15 that MIN-ICR has a solution of size $|PBP| + |M'|$. \square

THEOREM 18. MIN-ICR-EXACT is solvable in polynomial time under the restriction of SMTI in which indifference classes are of size at most 2.

PROOF. We show that $G(I, \mu)$ is a collection of cycles and paths, and hence its minimum vertex cover can be computed in polynomial time. The size of a minimum vertex cover for any path or cycle of length ℓ is $\lceil \frac{\ell}{2} \rceil$.

Take any vertex $v_1 = (m, \mu(m))$ in V . Recall that if any vertex $v_2 = (m', \mu(m'))$ is a neighbor of v_1 , then it must be that at least one of $(m, \mu(m'))$ or $(m', \mu(m))$ is in PBP'_2 . Note that if $(m, \mu(m')) \in PBP'_2$, then under I man m is indifferent between $\mu(m)$ and $\mu(m')$. Since each indifference class is of size at most 2, at most one such neighbor exists. Likewise, if $(m', \mu(m)) \in PBP'_2$ then $\mu(m)$ is indifferent between m and m' . However, since each indifference class is of size at most 2, at most one such neighbor exist. Thus, each vertex has degree at most 2, hence $G(I, \mu)$ is a collection of cycles and paths. \square

THEOREM 19. MIN-ICR-EXACT is solvable in polynomial time under the restriction of SMTI in which all men are endowed with the same indifference classes, as well as all women. That is $C_i^m = C_i^{m'}$ for all $m, m' \in M$ and all $i \in [n]$, and $C_i^w = C_i^{w'}$ for all $w, w' \in W$ and all $i \in [n]$.

PROOF. We show that $G(I, \mu)$ is a collection of complete graphs, and hence its minimum vertex cover can be computed in polynomial time, since the size of a minimum vertex cover for any complete graph K_ℓ is equal to $\ell - 1$. To prove that $G(I, \mu)$ is a collection of complete graphs, we show that for any three given vertices v_1, v_2 and v_3 , if $(v_1, v_2) \in E$ and $(v_1, v_3) \in E$ then $(v_2, v_3) \in E$.

Take any three vertices $v_1 = (m, \mu(m))$, $v_2 = (m', \mu(m'))$, and $v_3 = (m'', \mu(m''))$. If $(v_1, v_2) \in E$ then, under I , all men are indifferent between $\mu(m)$ and $\mu(m')$, all women are indifferent between m and m' , and $m, m' \notin M'$. If $(v_1, v_3) \in E$ then, under I , all men are indifferent between $\mu(m)$ and $\mu(m'')$, all women are indifferent between m and m'' , and $m'' \notin M'$. Therefore, since I is an instance of SMTI, all men are indifferent between $\mu(m)$, $\mu(m')$ and $\mu(m'')$, and all women are indifferent between m , m' , and m'' . Hence $(m', \mu(m''))$ and $(m'', \mu(m'))$ are PBPs. If $(m', \mu(m''))$ is a PBP-D2 then, as $m', m'' \notin M'$, $(m', \mu(m'')) \in PBP'_2$ and therefore $(v_2, v_3) \in E$. Assume for a contradiction that $(m', \mu(m''))$ is a PBP-D1. Assume that $\mu(m'') \succ_{m'} \mu(m')$ (the argument is similar if $m' \succ_{\mu(m'')} m''$), implying that $PBP_1(\mu(m'')) \neq \emptyset$ and thus $m'' \in M'$, a contradiction. \square

Theorem 18 is likely to be of more theoretical interest. For the setting of Theorem 17, we could envisage a hospitals-residents matching problem where residents are ranked uniformly (i.e., in a "master list" common to all hospitals [11]) according to some known objective value (e.g., which may be based on academic merit, as in the UK) and residents must use interviews in order to determine their true preferences over acceptable hospitals. For the setting of Theorem 19, consider a market with "tiered" preferences, where everybody agrees who/what belongs to each tier (again the membership of these tiers could relate to some objective values), but the precise ordering within these tiers could be subjective, and up to individuals to determine themselves. For

example, students may use national league tables for determining top tier universities, second tier universities and so on, but students' precise ranking over the universities in any given tier may vary.

If I is of one of the restricted forms for which MIN-ICR-EXACT is polynomial time solvable, then one straightforward approach to solving MIN-ICR is to enumerate all matchings that are stable under $\succ_{M,W}$ and then solve MIN-ICR-EXACT for each of them. This approach is practical if $\succ_{M,W}$ admits a polynomial number of stable matchings.

5. CONCLUSION AND FUTURE WORK

In this paper we have studied the complexity of the offline problem relating to computing an optimal interview strategy for a stable marriage market where initially participants have incomplete information, and the aim is to refine the instance using the minimum number of interviews in order to arrive at a super-stable matching. The main direction for future work is to investigate the online case, where the true underlying preferences are not known to the mechanism designer, with respect to measures such as the competitive ratio. Furthermore, an important question for which we do not know an answer yet is whether MIN-ICR is polynomial-time solvable under some restricted setting. Extending the known results on interviewing in stable marriage markets to many-to-one markets such as college admission is another important future direction. It is also interesting to study online algorithms in a setting where elicitation is taking place via comparison queries. In this paper we assume that the objective of the mechanism designer is to minimize the total number of interviews overall. One may however argue that such a strategy may require one or many agents to conduct most of the interviews while the others do none or very little. In the view of fairness and the practicality of such central interview-scheduling schemes, it is also of utmost importance to study settings in which a fair distribution of the interviews is also considered.

Acknowledgments. We thank Rob Irving and Piotr Krysta for the useful discussions and their valuable feedback. We also thank the anonymous reviewers for their insightful comments and suggestions. This work was supported by EPSRC grants EP/K01000X/1 and EP/K010042/1, and COST Action IC1205 on Computational Social Choice.

REFERENCES

- [1] J. Drummond and C. Boutilier. Elicitation and approximately stable matching with partial preferences. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI '13*, pages 97–105. AAAI Press, 2013.
- [2] J. Drummond and C. Boutilier. Preference elicitation and interview minimization in stable matchings. In *Proceedings of the Twenty-Eight AAAI Conference on Artificial Intelligence, AAAI '14*, pages 645–653. AAAI Press, 2014.
- [3] F. Echenique, S. Lee, M. Shum, and M. B. Yenmez. The revealed preference theory of stable and extremal stable matchings. *Econometrica*, 81(1):153–171, 2013.
- [4] D. Gale and L. Shapley. College admissions and the stability of marriage. *American Mathematics Monthly*, 69:9–15, 1962.
- [5] M. R. Garey and D. S. Johnson. The rectilinear steiner tree problem is NP-complete. *SIAM Journal on Applied Mathematics*, 32(4):826–834, 1977.
- [6] D. Gusfield and R. W. Irving. *The stable marriage problem: structure and algorithms*. MIT Press, Cambridge, MA, USA, 1989.
- [7] G. Haeringer and V. Iehle. Two-sided matching with one-sided information. *Working paper*, 2014.
- [8] R. Irving and D. Manlove. Finding large stable matchings. *ACM Journal of Experimental Algorithmics*, 14(2), 2009.
- [9] R. W. Irving. Stable marriage and indifference. *Discrete Applied Mathematics*, 48(3):261–272, 1994.
- [10] R. W. Irving, D. F. Manlove, and S. Scott. The hospitals/residents problem with ties. In *SWAT*, volume 1851 of *Lecture Notes in Computer Science*, pages 259–271, 2000.
- [11] R. W. Irving, D. F. Manlove, and S. Scott. The stable marriage problem with master preference lists. *Discrete Applied Mathematics*, 156(15):2959–2977, 2008.
- [12] S. Kalyanaraman and C. Umans. The complexity of rationalizing matchings. In *Proceedings of ISAAC '08: the 19th International Symposium on Algorithms and Computation*, volume 5369 of *Lecture Notes in Computer Science*, pages 171–182. Springer, 2008.
- [13] D. Knuth. *Marriages Stables*. Montréal: Les Presses de l'Université de Montréal, 1976.
- [14] R. Lee and M. Schwarz. Interviewing in two-sided matching markets. *Working Paper*, 2012.
- [15] D. Manlove. Stable marriage with ties and unacceptable partners. Technical Report TR-1999-29, University of Glasgow, Department of Computing Science, 1999.
- [16] D. F. Manlove. *Algorithmics of matching under preferences*. World Scientific, 2013.
- [17] D. F. Manlove, R. W. Irving, K. Iwama, S. Miyazaki, and Y. Morita. Hard variants of stable marriage. *Theoretical Computer Science*, 276(1-2):261 – 279, 2002.
- [18] B. Rastegari, A. Condon, N. Immorlica, R. Irving, and K. Leyton-Brown. Reasoning about optimal stable matching under partial information. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 431–448. ACM, 2014.
- [19] B. Rastegari, A. Condon, N. Immorlica, and K. Leyton-Brown. Two-sided matching with partial information. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 733–750. ACM, 2013.
- [20] A. Roth. The national residency matching program as a labor market. *Journal of the American Medical Association*, 275(13):1054–1056, 1996.
- [21] A. Roth and M. Sotomayor. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press, 1990.
- [22] A. E. Roth. The college admissions problem is not equivalent to the marriage problem. *Journal of Economic Theory*, 36:277–288, August 1985.