Operations Research Letters 45 (2017) 68–71

Contents lists available at ScienceDirect

Operations Research Letters

journal homepage: www.elsevier.com/locate/orl

A marriage matching mechanism menagerie

James W. Boudreau [a,*], Vicki Knoblauch [b]

[a] *Department of Economics and Finance, University of Texas Rio Grande Valley, United States*

[b] *Department of Economics, University of Connecticut, United States*

ARTICLE INFO

Article history:
Received 30 May 2016
Received in revised form
1 December 2016
Accepted 1 December 2016
Available online 10 December 2016

Keywords:
Matching
Assignment
Social welfare

ABSTRACT

For each of several measures of social welfare we present a marriage matching mechanism that produces a welfare maximizing matching, and our basic approach generalizes to many other welfare measures.

1. Introduction

Matching markets and assignment problems are popular topics with researchers involved in the field of mechanism design. In the case of the former, two distinct sets of agents, each with their own preferences over partners from the other side of the market, must be paired off; in the latter a set of agents with preferences must be matched to indivisible objects that may lack preferences or priorities of their own. Both tend to face difficulties if left unfettered due to the fact that most feature prohibitions on monetary transfers and also a distinct difficulty in terms of re-pairing over time. Along with the fact that assignments in these contexts tend to be both long-term and extremely important to those involved (classic examples include the marriage market, entry-level labor markets, public school assignment, and room or housing assignments) the process by which the allocation of prized partners or possessions is conducted can be quite controversial. It is therefore crucial to consider the social welfare of those involved when designing such procedures.

In this paper we present a family of algorithmic procedures to match two sets of agents in a way that maximizes a wide variety of – or combination of – welfare criteria in polynomial time. Though we specify our results in terms of a one-to-one matching problem and in terms of several specific welfare criteria, the general approach is quite flexible and can easily be adapted for alternative criteria or for assignment problems. Loosely speaking, our approach is to first reduce the graph representing the general matching problem to a collection of subgraphs that satisfy a specified set of specified welfare criteria and then allow agents to choose their partners according to an exogenously specified ordering, resulting in a match that is Pareto-optimal among all matchings that satisfy the desired combination of welfare measures.

As chronicled in the next section, our approach is motivated by the fact that different notions of social welfare are increasingly being recognized as important in the context of matching markets. Deferred acceptance algorithms based on the celebrated work of Gale and Shapley [10] have come to form the basis of most currently implemented centralized matching procedures due to their stability properties (for a history of the use of deferred acceptance algorithms see [21]). In the event that stability is not the primary goal, however, and market designers wish to maximize alternative notions of social welfare, our approach may prove quite useful.

2. Motivation and previous literature

It is well-known that the Gale–Shapley deferred acceptance algorithm produces a stable matching for any marriage matching market (to simplify terminology we use the marriage market interpretation to refer to one-to-one two-sided matching markets, labeling one set of agents men and the other women; many other interpretations are possible). That is, it produces a matching without blocking pairs, meaning there will never be a man–woman pair such that each prefers the other to his or her assigned mate [10]. This type of outcome has been celebrated by economists as it takes advantage of any possible mutual gains and thus results

* Corresponding author.
E-mail addresses: james.boudreau@utrgv.edu (J.W. Boudreau), vicki.knoblauch@uconn.edu (V. Knoblauch).

http://dx.doi.org/10.1016/j.orl.2016.12.001
0167-6377/

in Pareto-efficiency, but it is easy to imagine two-sided matching scenarios in which stability is unimportant, or at least not the most desired property of a matching. For example, a strong central authority or a system of binding contracts may enforce an unstable matching. Or complexity issues, incomplete information, time constraints or other barriers may make it extremely unlikely that members of a potential blocking pair will find each other.

Other than stability, two popular criteria for evaluating matchings are the *utilitarian* and *Rawlsian* welfare measurements. The utilitarian measure evaluates matchings based on the unweighted sum of agents' partner rankings in a match, while the Rawlsian measure considers only the ranking(s) of the agent(s) with the worst (meaning highest) ranked partner in a match (named for the *difference principal* of John Rawls, the Rawlsian criterion of social justice roughly asserts that the worst off in a society should be made as well off as possible [18]). Axtell and Kimbrough [2], for example, argue that "the nearly universal focus on stable matchings in this [matching] literature is misguided at best," and use simulations to highlight the utilitarian losses incurred by the Gale–Shapley algorithm as compared to a distributed, decentralized matching process. Masarani and Gokturk [17], meanwhile, prove the incompatibility of stability and the Rawlsian criterion, and Brams and Kilgour [7] outline a procedure – which can be thought of as a special case of one of the algorithms we present below – to attain Rawlsian-optimal matchings. Boudreau and Knoblauch [5] demonstrate how the suboptimal performance of the Gale–Shapley algorithm, in terms of both the utilitarian and Rawlsian measurements, varies according to properties of a marriage market's preferences, while Hafalir and Miralles [11] find mechanisms to attain utilitarian and Rawlsian optimality in special types of large assignment markets.

The utilitarian measure represents aggregate welfare while the Rawlsian measure focuses on the welfare of the least advantaged. A third form of welfare criterion popular in the literature is one that focuses on the equality of outcomes, which we refer to as *balancedness*. Such considerations are particularly important for matching markets, since it is well known that the Gale–Shapley algorithm favors one side of the market at the expense of the other. Accordingly, much attention has been given to matchings that attempt to balance the interests of both sides of the market. For example, Klaus and Klijn [14,15] establish the existence of "median" stable matchings in various matching formats, and Romero-Medina [19,20] provides algorithms to attain gender-balanced stable matchings in marriage markets. Boudreau and Knoblauch [6], however, show that stability is not generally compatible with balancedness across genders or individuals.

There are many other possible notions of welfare for matching markets (for example those in [6]), and the algorithms we describe below are widely applicable to any which can be described as properties of a graph. For brevity, however, we illustrate our results by focusing on the three basic measurements just described, in addition to combinations thereof.

Similarly, notice that in our use of Rawlsian, utilitarian and balanced measures of social welfare, we are implicitly assuming that participants' preferences over mates are cardinal and we are also making interpersonal comparisons of utility. For example, when using the Rawlsian measure of social welfare in a marriage matching market, a social architect who wishes to maximize the utility of the worst off participant and who does so by choosing a matching that minimizes the highest ranking number any participant attaches to his assigned mate is assuming that if $k < l$ then the utility received by participant i upon being matched with his or her kth choice is greater than the utility received by participant j ($j \neq i$) upon being matched with his or her lth choice.

Budish [8] emphasizes the nuances of different matching environments, and the fact that the design of any individual market should be based on its unique goals rather than those that have been celebrated elsewhere. Our work is in the spirit of that criticism, providing mechanisms that can match agents to achieve a wide variety of goals. Moreover, the structure of our approach is similar to that of Kesten [13] and Tang and Yu [22], who also proceed by iteratively refining a matching problem to proceed toward a desired solution.

3. Preliminaries

A *marriage matching* (or more simply *matching*) for a set $M = \{m_1, m_2, \ldots, m_n\}$ of men and a set $W = \{w_1, w_2, \ldots, w_n\}$ of women is a 1–1, onto function $\mu\colon M \cup W \to M \cup W$ such that $\mu(M) = W$ and for all $m \in M$, $\mu(\mu(m)) = m$. Let MM denote the collection of all marriage matchings for M and W.

For $x \in M \cup W$, let $\succ_x$ be a linear order representing x's preferences over members of $M \cup W$ of the opposite gender. Then $\succ_x$ gives rise to a ranking r_x such that if $y_{i_1} \succ_x y_{i_2} \succ_x \cdots \succ_x y_{i_n}$, then $r_x(y_{i_1}) = 1, r_x(y_{i_2}) = 2, \ldots, r_x(y_{i_n}) = n$. Let $\succ$ be the *preference profile* $(\succ_{m_1}, \succ_{m_2}, \ldots, \succ_{m_n}, \succ_{w_1}, \succ_{w_2}, \ldots, \succ_{w_n})$. Let π be the collection of all preference profiles for M and W. In nearly all that follows, we begin with the assumption that M, W and $\succ$ are fixed.

We will be dealing with certain properties of matchings, including but not limited to the following: a matching μ_0 is *Rawlsian* if it minimizes $\max_{x \in M \cup W} r_x(\mu(x))$, *utilitarian* if it minimizes $\sum_{x \in M \cup W} r_x(\mu(x))$ and *Pareto optimal* or *efficient* if for no $\mu \in MM$ is it the case that $r_x(\mu(x)) \leq r_x(\mu_0(x))$ for all $x \in M \cup W$ and that the inequality is strict for some $x \in M \cup W$. Given $\succ \in \pi$, a property P of matchings is non-null if there exists at least one matching with property P.

A *social welfare function (SWF)* is a real-valued function $f\colon \pi \times MM \to \mathbb{R}$. Whenever $\succ$ is fixed, we will write $f(\mu)$ rather than $f(\succ, \mu)$. It will sometimes be useful to describe a property of matchings in terms of a SWF. For example, a matching μ_0 is Rawlsian if $f_R(\mu_0) = \max_{\mu \in MM} f_R(\mu)$ where $f_R(\mu) = -\max_{x \in M \cup W} r_x(\mu(x))$.

A *bipartite graph* $G = (U, V; E)$ consists of two disjoint finite sets of vertices U and V and a collection E of edges with one endpoint in each set. The Hopcroft–Karp (HK) algorithm [12] inputs a bipartite graph, $(U, V; E)$ with, for our purposes, $|U| = |V|$ and in polynomial time outputs a matching μ for U and V such that $\{u, \mu(u)\} \in E$ for all $u \in U$, or outputs "no" if no such matching exists. The Hungarian (Hun) algorithm [16] inputs a *complete* bipartite graph (one with an edge $\{u, v\}$ for each $u \in U, v \in V$) with an integer assigned to each edge and in polynomial time outputs a matching that minimizes the sum of the edge values.

4. Marriage matching algorithms

For a given preference profile, a property P of matchings is *graph determined* if there exists a bipartite graph $G_0 = (M, W; E)$ such that for every $\mu \in MM$, μ has property P if and only if μ is a subgraph of G_0.

Given a preference profile, a non-null, graph-determined property P with associated graph G_0 and an enumeration $x_1, x_2, \ldots, x_{2n}$ of $M \cup W$, the following algorithm uses repeated applications of the HK algorithm to produce a matching that is Pareto optimal among all matchings with property P.

4.1. Lexicographic Optimizer (LO)

Given a preference profile, a bipartite graph $G_0 = (M, W; E_0)$ associated with a non-null, graph-determined property P and an enumeration $x_1, x_2, \ldots, x_{2n}$ of the men and women, we begin by initializing variables.

$$E \leftarrow E_0 \setminus \{\{x_1, y\}\colon r_{x_1}(y) > 1\}, \qquad G \leftarrow (M, W, E),$$
$$\mu \leftarrow \emptyset, \qquad c \leftarrow 1, \qquad d \leftarrow 1.$$