Journal of Mathematical Economics 54 (2014) 63–73

Contents lists available at ScienceDirect

Journal of Mathematical Economics

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

A many-to-many 'rural hospital theorem'☆

Flip Klijn[a], Ayşe Yazıcı[b,*]

[a] Institute for Economic Analysis (CSIC) and Barcelona GSE, Campus UAB, 08193 Bellaterra (Barcelona), Spain
[b] Durham University Business School, Durham, DH1 3LB, United Kingdom

ARTICLE INFO

Article history:
Received 15 October 2012
Received in revised form
30 August 2014
Accepted 1 September 2014
Available online 10 September 2014

Keywords:
Matching
Many-to-many
Stability
Rural hospital theorem

ABSTRACT

We show that the full version of the so-called 'rural hospital theorem' generalizes to many-to-many matching problems where agents on both sides of the problem have substitutable and weakly separable preferences. We reinforce our result by showing that when agents' preferences satisfy substitutability, the domain of weakly separable preferences is also maximal for the rural hospital theorem to hold.

1. Introduction

We study two-sided matching problems. 'Stability' of outcomes in these problems is considered to be the main property that accounts for the success of matching rules. We identify a large and maximal preference domain for which 'underdemanded' institutions (or agents) have the same partners at each stable outcome. Consequently, no stable rule can implement possibly desirable changes in the set of partners assigned to such institutions.

Our study is motivated by issues raised in certain centralized labor markets. As an illustration, many countries employ each year a centralized mechanism to assign newly graduated medical students to positions in residency programs. Hospitals in rural areas are typically less preferred than those in urban areas by medical graduates, i.e., they are ranked below urban hospitals in a typical student's preference list. Also, graduates from relatively successful programs are more popular among hospitals, i.e., they are ranked above other students in a typical hospital's preference list. Rural hospitals complain that their positions may not be filled by the stable matching rule in use and that they may not be assigned popular graduates. The 'rural hospital theorem' established in several matching models states that the *number* of medical graduates assigned to a hospital and the *set* of graduates assigned to a hospital in a rural area do not vary across stable outcomes. Even though the theorem's name is a useful reminder of its content and origin, the 'rural hospital theorem' equally applies to other labor markets with similar concerns about the numerical distribution of workers or the composition of the workforce of firms.

We study the 'rural hospital theorem' in the context of many-to-many labor markets, i.e., markets where each agent can engage in multiple partnerships. There are several reasons to focus on *many-to-many* markets instead of many-to-one markets (where each worker can be employed by at most one firm). First, a well-known many-to-many market is the medical labor market in the UK. More specifically, each medical graduate in the UK has to seek two positions (a medical position and a surgical position) to be able to register as a medical doctor. Shallcross (2005) mentioned concerns of doctor shortages in rural areas in the UK. Second, as pointed out by Echenique and Oviedo (2006), even if in a labor market most workers are employed by one firm, the presence of a few workers with multiple employers can make a crucial difference. Precisely, Echenique and Oviedo (2006, Example 2.2) showed that the presence of only one worker with two part-time jobs can already change the stable outcome for all other agents. Thus, the functioning of even 'almost many-to-one' labor markets can only be understood through the study of many-to-many matching models. Third, the literature on many-to-many matching markets

☆ We thank William Thomson for his advice, an associate editor and two referees for their useful comments, and Ahmet Alkan, Sophie Bade, Federico Echenique, Lars Ehlers, Bettina Klaus, Fuhito Kojima, Jordi Massó, Duygu Nizamoğulları, İpek Özkal Sanver, Marilda Sotomayor, and seminar participants at the University of Rochester, SED 2011 conference (Montréal), ESF Matching workshop (IAE, Barcelona), ESRC Game Theory Workshop (LSE), University of St. Andrews, Universidad del Rosario and SSCW 2014 conference (Boston) for helpful comments and discussions.

* Corresponding author.
E-mail addresses: flip.klijn@iae.csic.es (F. Klijn), ayse.yazici@durham.ac.uk (A. Yazıcı).

http://dx.doi.org/10.1016/j.jmateco.2014.09.003
0304-4068/

has grown in the last couple of years,[1] but there is still a wide gap with respect to many-to-one markets. Fourth, there are important structural differences between many-to-one and many-to-many matching markets, even if all agents have so-called 'responsive' preferences. For instance, Sotomayor (1999) showed that unlike many-to-one markets, in many-to-many markets the set of stable outcomes need not coincide with the core. Finally, our results are not only novel to the many-to-many framework. Indeed, *the restriction of all our results to the many-to-one framework yields new results and subsumes existing results for that framework* as well.[2]

Next, we describe in more detail the model we study, the existing literature, and our contribution. In a two-sided many-to-many matching problem there are two disjoint sets of agents, which we call 'firms' and 'workers'. Each firm (worker) can only form partnerships with workers (firms). Each agent has a preference order over the set of all subsets of partnerships, i.e., subsets of agents in the other set. For each agent, there is a maximal number ('quota') of partnerships the agent can or is willing to be involved in. An outcome of the problem is a 'match' which consists of a collection of partnerships.

A match is 'stable' if no agent prefers to be matched to a proper subset of its current partners, and no set of firms and workers prefers to establish new partnerships only among themselves and possibly break up some existing partnerships.[3] This definition is more stringent than so-called pairwise stability which is another standard solution concept but that only eliminates blocking by firm–worker pairs. Stability proved to be a crucial property in many entry-level labor markets where workers are matched to firms through a clearinghouse. It has been observed that clearinghouses that use stable rules often perform better than those that use rules that do not necessarily produce stable matches.[4] According to Roth (1991, p. 422) even the weaker stability concept, pairwise stability, can be of primary importance for many-to-many markets as well.

There are many-to-many problems for which no stable match exists. Certain assumptions on preferences have been identified to guarantee that they do. A firm's preferences are 'substitutable' if whenever a worker is chosen from a group of workers by this firm, she is also chosen from any of the group's subsets to which she belongs.[5] Substitutability of workers' preferences is defined similarly. Substitutability is a standard assumption in the literature and it guarantees the existence of a pairwise stable match.[6] Hatfield and Kominers (2012a) showed that for substitutable preferences, stability and pairwise stability are equivalent.[7] Thus, when preferences are substitutable, the set of stable matches is non-empty and coincides with the set of pairwise stable matches. *With the important exception of* Proposition 1, *we assume substitutability throughout.*

Taking the requirement of stability as granted, an important question is whether the choice of a particular stable rule affects the numerical distribution of workers; and if not, whether different matches assign different sets of workers to a firm that does not fill all its positions. For instance, in the context of the assignment of medical graduates in the US, the National Resident Matching Program (NRMP) failed to fill the posts of many hospitals in (typically less preferred) rural areas (Sudarshan and Zisook, 1981). However, provided that the preferences satisfy certain conditions, the problem of the rural hospitals cannot be attributed to the particular stable rule used by the NRMP. Indeed, the results obtained by Gale and Sotomayor (1985) and Roth (1984b, 1986) suggest that any other stable rule would yield (R1) the same numerical distribution of medical graduates and would assign (R2) the same medical graduates to each rural hospital that does not fill all its posts. The two results are known as weak and strong versions of the rural hospital theorem.[8]

Both versions of the rural hospital theorem play a functional role in proving many appealing results. For instance, R1 is used to show the lattice structure of the set of stable matches (Martínez et al., 2001) and the group strategy-proofness (for the workers' side) of the worker-optimal stable rule (Martínez et al., 2004a) in a many-to-one model. Ma (2002) studied refinements of Nash equilibrium based on 'truncations at the match point' for the preference revelation game induced by any stable rule. He used R2 to prove that each equilibrium outcome is stable for the true preferences. Pais (2006) studied ordinal Nash equilibria of the preference revelation game induced by any probabilistic stable rule. She used R2 to show that any equilibrium induces a match that is individually rational for the true preferences. Yazıcı (2012) also employed R2 to extend the latter result to many-to-many matching with a more general preference domain. These results show that the relevance of the rural hospital theorem goes beyond its direct interpretation: it is a powerful tool in establishing structural results and analyzing strategic matching games.

The first papers on the rural hospital theorem (e.g., Gale and Sotomayor, 1985; Roth, 1984b, 1986) studied many-to-one matching problems and assumed firms' preferences to be 'responsive'. A firm's preferences over groups of workers are responsive to its preferences over individual workers if (i) for two groups that only differ in one worker, the firm prefers the one with the preferred worker, and (ii) adding an acceptable (unacceptable) worker to a group that does not fill its quota makes the group better (worse). Responsiveness implies substitutability. Several papers have shown R1 and R2 for preference domains that are strictly larger than the domain of responsive preferences.[9] A firm's preferences are 'separable' if condition (ii) above holds. R1 and R2 hold for substitutable and separable preferences (Martínez et al., 2000, Proposition 2). Since responsiveness implies separability, Martínez et al.'s (2000) result subsumes the previous rural hospital theorem results.[10]

Concerning the many-to-many framework, Alkan (2002, Proposition 6) showed that R1 holds for substitutable and 'cardinally monotonic' preferences. A firm's preferences over groups of workers are cardinally monotonic if whenever the group of workers available to the firm expands, it will not employ fewer workers.[11]

1 Recent papers on many-to-many matching include, among others, Hatfield and Kominers (2012a,b), Jaramillo et al. (2014), Klaus and Walzl (2009), Kojima and Ünver (2008), Kominers (2012), Konishi and Ünver (2006), Ostrovsky (2008), Sotomayor (2004), and Yazıcı (2012).

2 See the discussion that precedes Theorem 3 and Remarks 7 and 8.

3 This is an adaptation of the stability definition in Hatfield and Kominers (2012a).

4 See, for instance, Roth (1991).

5 Substitutability is an adaptation of the gross substitutability property (Kelso and Crawford, 1982) by Roth (1984a) and Roth and Sotomayor (1990) to matching problems without monetary transfers.

6 The existence of a pairwise stable match can be shown via an algorithm for strict preferences (Roth, 1984a) and via a non-constructive proof for non-necessarily strict preferences (Sotomayor, 1999). See also Martínez et al. (2004b) for the computation of the full set of pairwise stable matches.

7 We are thankful to a referee for pointing this out.

8 Since R2 implies R1, R1 (R2) is often referred to as the weak (strong) rural hospital theorem.

9 For the reader's convenience, we refer to the Venn diagram of Fig. 1 (in Section 3) which depicts the inclusion relations among the preference domains we discuss.

10 Kojima (2012) also introduced the domain of separable preferences with so-called affirmative action constraints. This domain is a strict superset of the domain of separable preferences but a strict subset of the domain of cardinally monotonic preferences. Kojima (2012) showed that on his domain an appropriately adjusted version of R2 holds.

11 Cardinal monotonicity is called size monotonicity and law of aggregate demand in Alkan and Gale (2003) and Hatfield and Milgrom (2005), respectively.