

Peer Effects and Stability in Matching Markets

Elizabeth Bodine-Baron, Christina Lee, Anthony Chong, Babak Hassibi, and
Adam Wierman

California Institute of Technology
Pasadena, CA 91125, USA*
{eabodine, chlee, anthony, hassibi, adamw}@caltech.edu

Abstract. Many-to-one matching markets exist in numerous different forms, such as college admissions, matching medical interns to hospitals for residencies, assigning housing to college students, and the classic firms and workers market. In all these markets, externalities such as complementarities and peer effects severely complicate the preference ordering of each agent. Further, research has shown that externalities lead to serious problems for market stability and for developing efficient algorithms to find stable matchings. In this paper we make the observation that peer effects are often the result of underlying social connections, and we explore a formulation of the many-to-one matching market where peer effects are derived from an underlying social network. The key feature of our model is that it captures peer effects and complementarities using utility functions, rather than traditional preference ordering. With this model and considering a weaker notion of stability, namely two-sided exchange stability, we prove that stable matchings always exist and characterize the set of stable matchings in terms of social welfare. To characterize the efficiency of matching markets with externalities, we provide general bounds on how far the welfare of the worst-case stable matching can be from the welfare of the optimal matching, and find that the structure of the social network (e.g. how well clustered the network is) plays a large role.

1 Introduction

Many-to-one matching markets exist in numerous forms, such as college admissions, the national medical residency program, freshman housing assignment, as well as the classic firms-and-workers market. These markets are widely studied in academia and also widely deployed in practice, and have been applied to other areas, such as FCC spectrum allocation and supply chain networks [4,21]

In the conventional formulation, matching markets consist of two sets of agents, such as medical interns and hospitals, each of which have preferences over the agents to which they are matched. In such settings it is important that matchings are ‘stable’ in the sense that agents do not have incentive to change assignments after being matched. The seminal paper on matching markets was by Gale and Shapley [13], and following this work an enormous literature has grown, e.g., [20,27,28,29] and the references therein. Further, variations on Gale

* This work was supported in part by the National Science Foundation under grants CCF-0729203, CNS-0932428 and CCF-1018927, by the Office of Naval Research under the MURI grant N00014-08-1-0747, and by Caltech’s Lee Center for Advanced Networking.

and Shapley’s original algorithm for finding a stable matching are in use by the National Resident Matching Program (NRMP), which matches medical school graduates to residency positions at hospitals [26].

However, there are problems with many of the applications of matching markets in practice. For example, couples participating in the NRMP often reject their matches and search outside the system. In housing assignment markets where college students are asked to list their preferences over housing options, there is often collusion among friends to list the same preference order for houses. These two examples highlight that ‘peer effects’, whether just couples or a more general set of friends, often play a significant role in many-to-one matchings. That is, agents care not only where they are matched, but also which other agents are matched to the same place. Similarly, ‘complementarities’ often play a role on the other side of the market. For example, hospitals and colleges care not only about which individual students are assigned to them, but also that the group has a certain diversity, e.g., of different specializations.

As a result of the issues highlighted above, there is a growing literature studying many-to-one matchings with externalities (i.e., peer effects and complementarities) [10,14,18,19,22,24,3,11,30] and the research has found that designing matching mechanisms is significantly more challenging when externalities are considered, e.g. incentive compatible mechanism design is no longer possible.

The reason for the difficulty is that there is no longer a guarantee that a stable many-to-one matching will exist when agents care about more than their own matching [26,28], and, if a stable matching does exist, it can be computationally difficult to find [25]. Consequently, most research has focused on identifying when stable matchings do and do not exist. Papers have proceeded by constraining the matching problem through restrictions of the possible preference orderings, [10,14,18,19,22,24], and by considered variations on the standard notion of stability [3,11,30].

The key idea of this paper is that *peer effects are often the result of an underlying social network*. That is, when agents care about where other agents are matched, it is often because they are friends. With this in mind, we construct a model in Section 2 that includes a weighted, undirected social network graph and allows agents to have utilities (which implicitly defines their preference ordering) that depend on where neighbors in the graph are assigned. The model is motivated by [3], which also considers peer effects defined by a social network but focuses on one-sided matching markets rather than two-sided matching markets.

We focus on *two-sided exchange-stable* matchings – see Section 2 for a detailed definition. We note that compared to the traditional notion of stability of [13], this is a distinct notion of stability, but one that is relevant to many situations where agents can compare notes with each other, such as the housing assignment or medical matching problem. For example, in [3,4,12], “pairwise-stability” is considered since they consider models where agents exchange offices or licenses in FCC spectrum auctions. Further, consider a situation where two hospital interns prefer to exchange the hospitals allocated to them by the NRMP. If this is a traditional stable matching, the hospitals would not allow the swap, even though the interns are highly unsatisfied with the match. Such a situation has been documented in [15], and has led to a similar type of stability, exchange stability, as defined in [1,8,9,15].

Given our model of peer effects, the focus of the paper is then on characterizing the set of two-sided exchange-stable matchings. Our results concern (i)

the existence of two-sided exchange-stable matchings and (ii) the efficiency of exchange-stable matchings (in terms of social welfare).

With respect to the existence of two-sided exchange-stable matchings (Section 3), it is not difficult to show that in our model stable matchings always exist. Given the contrast to the negative results that are common for many-to-one matchings, e.g., [11,25,26], these results are perhaps surprising.

With respect to the efficiency of exchange-stable matchings (Section 4), results are not as easy to obtain. In this context, we limit our focus to one-sided matching markets and simplify utility functions, but as a result we are able to attain bounds on the ratio of the welfare of the optimal matching to that of the worst stable matching, i.e., the ‘price of anarchy’. We also demonstrate cases where our bounds are tight. When considering only one-sided markets, our model becomes similar to hedonic coalition formation, but with several key differences, as highlighted in Section 4. Our results (Theorems 3 and 4) show that the price of anarchy does not depend on the number of, say, interns, but does grow with the number of, say, hospitals – though the growth is typically sublinear. Further, we observe that the impact of the structure of the social network on the price of anarchy happens only through the clustering of the network, which is well understood in the context of social networks, e.g., [16,32]. Finally, it turns out that the price of anarchy has a dual interpretation in our context; in addition to providing a bound on the inefficiency caused by enforcing exchange-stability, it turns out to also provide a bound on the loss of efficiency due to peer effects.

2 Model and notation

To begin, we define the model we use to study many-to-one matchings with peer effects and complementarities. There are four components to the model, which we describe in turn: (i) basic notation for discussing matchings; (ii) the model for agent utilities, which captures both peer effects and complementarities; (iii) the notion of stability we consider; and (iv) the notion of social welfare we consider.

To provide a consistent language for discussing many-to-one matchings, throughout this paper we use the setting of matching incoming students to residential houses. In this setting many students are matched to each house, and the students have preferences over the houses, but also have peer effects as a result of wanting to be matched to the same house as their friends. Similarly, the houses have preferences over the students, but there are additional complementarities due to goals such as maintaining diversity. It is clear that some form of stability is a key goal of this “housing assignment” problem.

Notation for many-to-one matchings We define two finite and disjoint sets, $H = \{h_1, \dots, h_m\}$ and $S = \{s_1, \dots, s_n\}$ denoting the houses and students, respectively. For each house, there exists a positive integer *quota* q_h which indicates the number of positions a house has to offer. The quota for each house may be different.

Definition 1. A matching is a subset $\mu \subseteq S \times H$ such that $|\mu(s)| = 1$ and $|\mu(h)| = q_h$, where $\mu(s) = \{h \in H : (s, h) \in \mu\}$ and $\mu(h) = \{s \in S : (s, h) \in \mu\}$.¹

¹ If the number of students in $\mu(h)$, say r , is less than q_h , then $\mu(h)$ contains $q_h - r$ “holes” – represented as students with no friends and no preference over houses.

Note that we use $\mu^2(s)$ to denote the set of student s 's housemates (students also in house $\mu(s)$).

Friendship network The friendship network among the students is modeled by a weighted graph, $G = (V, E, w)$ where $V = S$ and the relationships between students are represented by the weights of the edges connecting nodes. The strength of a relationship between two students s and t is represented by the weight of that edge, denoted by $w(s, t) \in \mathbb{R}^+ \cup \{0\}$. We require that the graph is undirected, i.e., the adjacency matrix is symmetric so that $w(s, t) = w(t, s)$ for all s, t .

Additionally, we define a few metrics quantifying the graph structure and its role in the matching. Let the total weight of the graph be denoted by $|E| := \frac{1}{2} \sum_{s \in S} \sum_{t \in S} w(s, t)$. Further, let the weight of edges connecting houses h and g under matching μ be denoted by $E_{hg}(\mu) := \sum_{s \in \mu(h)} \sum_{t \in \mu(g)} w(s, t)$. Note that in the case of edges within the same house $E_{hh}(\mu) := \frac{1}{2} \sum_{s \in \mu(h)} \sum_{t \in \mu(h)} w(s, t)$. Finally, let the weight of edges that are within the houses of a particular matching μ be denoted by $E_{in}(\mu) := \sum_{h \in H} E_{hh}(\mu)$.

Agent utility functions Each agent derives some utility from a particular matching and an agent (student or house) always strictly prefers matchings that give a strictly higher utility and is indifferent between matchings that give equal utility. This setup differs from the traditional notion of ‘preference orderings’ [13,28], but is not uncommon [2,3,4,7,12]. It is through the definitions of the utility functions that we model peer effects (for students) and complementarities (for houses).

Students derive benefit both from (i) the house they are assigned to and (ii) their peers that are assigned to the same house. We model each house h as having an desirability to student s of $D_h^s \in \mathbb{R}^+ \cup \{0\}$. If $D_h^s = D_h^t \forall s \neq t$ (objective desirability), this value can be seen as representing something like the U.S. News college rankings or hospital rankings – something that all students would agree on. This leads to a utility for student s under matching μ of

$$U_s(\mu) := D_{\mu(s)}^s + \sum_{t \in \mu^2(s)} w(s, t) \quad (1)$$

so that the total utility that a student derives from a match is a combination of how “good” a house is as well as how many friends they will have in that house.²³

Similarly, the utility of a house h under matching μ is modeled by

$$U_h(\mu) := D_{\mu(h)}^h, \quad (2)$$

where D_σ^h denotes the desirability of a particular set of students σ for house h (the utility house h derives from being matched to the set of students σ). Note that this definition of utility allows general phenomena such as heterogeneous house preferences over groups of students.

² We note that the utility of any “holes” (such as what happens when a house’s quota is not met), is simply $U_s(\mu) = 0$.

³ Note also that if we remove D_h^s from the utility function and allow unlimited quotas, the matching problem becomes the coalitional affinity game from [7].

Two-sided exchange stability Under the traditional definition of stability, if a student and a house were to prefer each other to their current match (forming a blocking pair), the student is free to move to the preferred house and the house is free to evict (if necessary) another student to make space for the preferred student. In our model, however, we assume that students and houses cannot “go outside the system” (neither can students remain unmatched), like what medical students and hospitals do when they operate outside of the NRMP. As a result, we restrict ourselves to considering swaps of students between houses, similar to [3,4,12].

To define exchange stability, it is convenient to first define a *swap matching* μ_s^t in which students s and t switch places while keeping all other students’ assignments the same.

Definition 2. A *swap matching* $\mu_s^t = \{\mu \setminus \{(s, h), (t, g)\}\} \cup \{(s, g), (t, h)\}$.

Note that the agents directly involved in the swap are the two students switching places and their respective houses – all other matchings remain the same. Further, one of the students involved in the swap can be a “hole” representing an open spot, thus allowing for single students moving to available vacancies. When two actual students are involved, this type of swap is a two-sided version of the “exchange” considered in [1,8,9,15] – *two-sided* exchange stability requires that houses approve the swap. As a result, while an exchange-stable matching may not exist in either the marriage or roommate problem, we show in Section 3 that a two-sided exchange-stable matching will always exist for the housing assignment problem.

Definition 3. A matching μ is *two-sided exchange-stable (2ES)* if and only if there does not exist a pair of students (s, t) such that:

- (i) $\forall i \in \{s, t, \mu(s), \mu(t)\}, U_i(\mu_s^t) \geq U_i(\mu)$ and
- (ii) $\exists i \in \{s, t, \mu(s), \mu(t)\}$ such that $U_i(\mu_s^t) > U_i(\mu)$

This definition implies that a swap matching in which all agents involved are indifferent is two-sided exchange-stable. This avoids looping between equivalent matchings. Note that the above definition implies that if two students want to switch between two houses (or a single student wants to “switch” with a hole), the houses involved must “approve” the swap or if two houses want to switch two students, the students involved must agree to the swap (a hole will always be indifferent). This is natural for the house assignment problem and many other many-to-one matching markets, but would be less appropriate for some other settings, such as the college-admissions model.

Social welfare One key focus of this paper is to develop an understanding of the “efficiency loss” that results from enforcing stability of assignments in matching markets. We measure the efficiency loss in terms of the “social welfare”:

$$W(\mu) := \sum_{s \in S} U_s(\mu) + \sum_{h \in H} U_h(\mu)$$

Using this definition of social welfare, the efficiency loss can be quantified using the *Price of Anarchy* (PoA) and *Price of Stability* (PoS). Specifically, the PoA (PoS) is the ratio of the optimal social welfare over all matchings, not necessarily stable, to the minimum (maximum) social welfare over all stable matchings. Understanding the PoA and PoS is the focus of Section 4.

3 Existence of stable matchings

We begin by focusing on the existence of two-sided exchange-stable matchings. In most prior work, matching markets with externalities do not have guaranteed existence of a stable matching. In contrast, we prove that a 2ES matching always exists in the model considered in this paper. We begin by proposing a potential function $\Phi(\mu)$ for the matching game:

$$\Phi(\mu) = \sum_{h \in H} U_h(\mu) + \sum_{s \in S} D_{\mu(s)}^s + \frac{1}{2} \sum_{s \in S} \left(\sum_{x \in \mu^2(s)} w(s, x) \right) \quad (3)$$

Due to the symmetry of the social network, every approved swap will result in a strict increase of the potential function. As there is a finite set of matches, this results in the existence of a 2ES matching for every housing assignment market.

Theorem 1. *All local maxima of $\Phi(\mu)$ are two-sided exchange-stable.*

If we assume that there are no vacancies in any of the houses and students value houses according to the same rules (i.e., $D_h^s = D_h^t \forall s \neq t$), then each approved swap will result in a strict increase in the *social welfare*. Note that this implies that the maximally efficient matching will be 2ES.

Theorem 2. *If house quotas are exactly met and $D_h^s = D_h^t \forall s \neq t$, all local maxima of $W(\mu)$ are two-sided exchange-stable.*

We omit the exact proofs here; see [5] for details. Note, however, that not all 2ES matchings are local maxima of $\Phi(\mu)$ or $W(\mu)$. Such a case arises when one student, for example, refuses a swap as her utility would decrease, but the other student involved stands to benefit a great deal from such a swap. If the swap were forced, the total potential function (or social welfare) could increase, but only at the expense of the first student.

The contrast between Theorem 1 and the results such as [26] and [28] can be explained by considering a few aspects of the model we study. In particular, we are using a distinct type of stability appropriate to our housing assignment market. Further, the assumption that the social network graph is symmetric are key to guaranteeing existence.

4 Efficiency of stable matchings

To this point, we have focused on the existence of two-sided exchange-stable matchings and how to find them. In this section our focus is on the “efficiency loss” due to stability in a matching market and the role peer effects play in this efficiency loss.

We measure the efficiency loss in a matching market using the price of stability (PoS) and the price of anarchy (PoA) as defined in Section 2. Interestingly, the price of anarchy has multiple interpretations in the context of this paper. First, as is standard, it measures the worst-case loss of social welfare that results due to enforcing two-sided exchange-stability. For example, the authors in [2] bound the loss in social welfare caused by individual rationality (by enforcing

stable matchings) for matching markets without externalities. Second, it provides a competitive ratio for matching algorithms (like those described in [5]). Even even a centralized mechanism with complete information may only find a stable matching, not necessarily the maximally efficient one. The price of anarchy gives us a bound on this worst-case. Third, we show later that the price of anarchy also has an interpretation as capturing the efficiency lost due to peer effects.

The results in this section all require one additional simplifying assumption to our model: *complementarities are ignored and only peer effects are considered*. Specifically, we assume, for all of our PoA results, $U_h(\mu) = 0$, and thus $W(\mu) = \sum_{s \in S} U_s(\mu)$. Under this assumption, the market is one-sided, with only students participating, and our notion of stability is simply exchange-stability. This assumption is limiting, but there are still many settings within which the model is appropriate. Two examples are the housing assignment problem in the case when students can swap positions without needing house approval, and the assignment of faculty to offices as discussed in [3], as clearly the offices have no preferences over which faculty occupy them. In order to simplify the analysis, we also make use of the assumptions in Theorem 2 : (i) $D_h^s = D_h^t \forall s \neq t$ and (ii) house quotas are exactly met.

4.1 Related models

When the housing assignment problem is restricted to a one-sided market involving only students, we note that it becomes very similar to both (i) a hedonic coalition formation game with symmetric additively separable preferences, as described in [6], and (ii) a coalitional affinity game, as described in [7]. For a more detailed discussion of these types of games and their relation to the results in this paper, see [5].

While the one-sided housing assignment problem and hedonic coalition formation games appear to be very similar, there are a number of key differences. Most importantly, the housing assignment problem considers a *fixed* number of houses with a limited number of spots available; students cannot break away and form a new coalition/house, nor can a house have more students than its quota. In addition, our model considers exchange-stability, which is closest to the Nash stability of [6], but is still significantly different in that it involves a pair of students willing and able to swap. Finally, each student gains utility from the house they are matched with, in addition to the other members of that house, which is different from the original formulation of hedonic coalition games.

4.2 Discussion of results

To begin the discussion of our results, note that under our simplifying assumptions the price of stability is 1 for our model because any social welfare optimizing matching is stable. However, the price of anarchy can be much larger than 1. In fact, depending on the social network, the price of anarchy can be unboundedly large, as illustrated in the following example.

Example 1 (Unbounded price of anarchy).

Consider a matching market with 4 students and 2 houses, each with a quota of 2, and two possible matchings illustrated by Figure 1. As shown in Figure

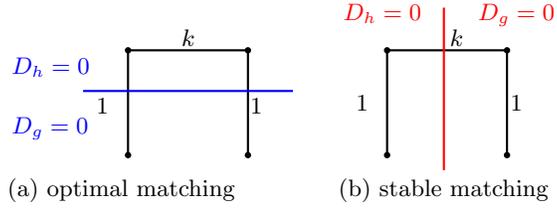


Fig. 1: Arbitrarily bad exchange-stable matching

1 (a) and (b), respectively, in the optimal matching μ^* , $W(\mu^*) = k$; whereas there exists a exchange-stable matching with $W(\mu) = 2$. Thus, as k increases, the price of anarchy grows linearly in k .

Despite the fact that, in general, there is a large efficiency loss that results from enforcing exchange-stability, in many realistic cases the efficiency loss is actually quite small. The following two theorems provide insight into such cases.

A key parameter in these theorems is γ_m^* which captures how well the social network can be “clustered” into a *fixed* number of m groups and is defined as follows.

$$\gamma_m(\mu) := \frac{E_{in}(\mu)}{|E|} \quad (4)$$

$$\gamma_m^* := \max_{\mu} \gamma_m(\mu) \quad (5)$$

Thus, γ_m^* represents the maximum edges that can be captured by a partition satisfying the house quotas. Note that γ_m^* is highly related to other clustering metrics, such as the conductance [17], [31] and expansion [23].

We begin by noting that due to the assumption that $\sum_{h \in H} U_h(\mu) = 0$, we can separate the social welfare function into two components:

$$W(\mu) = \sum_{s \in S} U_s(\mu) = \sum_{h \in H} \sum_{s \in \mu(h)} \left(D_h + \sum_{t \in \mu(h)} w(s, t) \right) = 2E_{in}(\mu) + \sum_{h \in H} q_h D_h.$$

Thus,

$$\frac{\max_{\mu} W(\mu)}{\min_{\mu \text{ is stable}} W(\mu)} = \frac{Q + \max_{\mu} \gamma_m(\mu)}{Q + \min_{\text{stable } \mu} \gamma_m(\mu)} \quad (6)$$

where

$$Q := \frac{\sum_{h \in H} q_h D_h}{2E}. \quad (7)$$

Note that the parameter Q is independent of the particular matching μ .

Our first theorem regarding efficiency is for the “simple” case of unweighted social networks with equal house quotas and/or equivalently valued houses.

Theorem 3. *Let $w(s, t) \in \{0, 1\}$ for all students s, t and let $q_h \geq 2, D_h \in \mathbb{Z}^+ \cup \{0\}$ for all houses h . If $q_h = q$ for all h and/or $D_h = D$ for all h , then*

$$\min_{\text{stable } \mu} W(\mu) \geq \frac{\max_{\mu} W(\mu)}{1 + 2(m-1)\gamma_m^*}$$

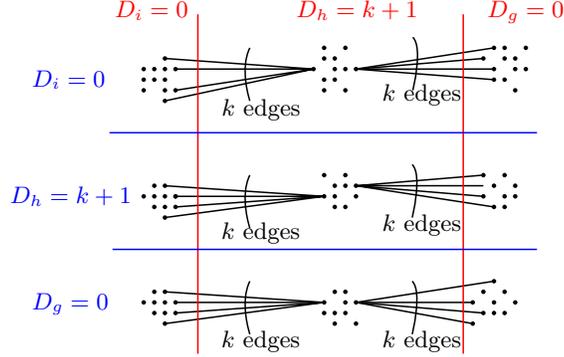


Fig. 2: Network that achieves PoA bound.

The bound in Theorem 3 is tight, as illustrated by the example below.

Example 2 (Tightness of Theorem 3).

Consider a setting with m houses and $q_h = mk$ for all $h \in H$. Students are grouped into clusters of size $k > 2$, as shown for $m = 3$ in Figure 2. The houses have $D_h = k + 1$ and $D_g = D_i = 0$. Each student in the middle cluster in each row has k edges to the other students outside of their cluster (but none within), as shown.

The worst-case stable exchange-matching is represented by the vertical red lines. Note that since $D_h = k + 1$, this matching is stable, even though all edges are cut. Thus $\min_{\mu \text{ stable}} \gamma_m(\mu) = 0$. The optimal matching is represented by the horizontal blue lines in the figure; note that $\gamma_m^* = 1$. To calculate the price of anarchy, we start from equations (6) and (7) and calculate

$$Q = \frac{\sum_{h \in H} q_h D_h}{2|E|} = \frac{mk(k+1)}{2mk(m-1)k} = \frac{k+1}{2(m-1)k},$$

which gives,

$$\frac{\max_{\mu} W(\mu)}{\min_{\text{stable } \mu} W(\mu)} = \frac{Q + \gamma_m^*}{Q + \min_{\mu \text{ stable}} \gamma_m(\mu)} = 1 + 2(m-1) \left(\frac{k}{k+1} \right).$$

Notice that as k becomes large, this approaches the bound of $1 + 2(m-1)\gamma_m^*$.

We note that the requirement $q_h = q$ for all h and/or $D_h = D$ for all h is key to the proof of Theorem 3 and in obtaining such a simple bound; otherwise, Theorem 4 applies. We omit the proofs of these theorems here for brevity; see [5] for the details.

Our second theorem removes the restrictions in the theorem above, at the expense of a slightly weaker bound. Define $q_{max} = \max_{h \in H} q_h$, $w_{max} = \max_{s, t \in S} w(s, t)$ and $D_{\Delta} = \min_{h, g \in H} (D_h - D_g)$, assuming that the houses are ordered in increasing values of D_h .

Theorem 4. *Let $w(s, t) \in \mathbb{R}^+ \cup \{0\}$ for all students s, t and $D_h \in \mathbb{R}^+ \cup \{0\}$, $q_h \in \mathbb{Z}^+$ for all houses h , then*

$$\min_{\text{stable } \mu} W(\mu) \geq \frac{\max_{\mu} W(\mu)}{1 + 2(m-1) \left(\gamma_m^* + \frac{q_{max} w_{max}}{D_{\Delta}} \right)}$$

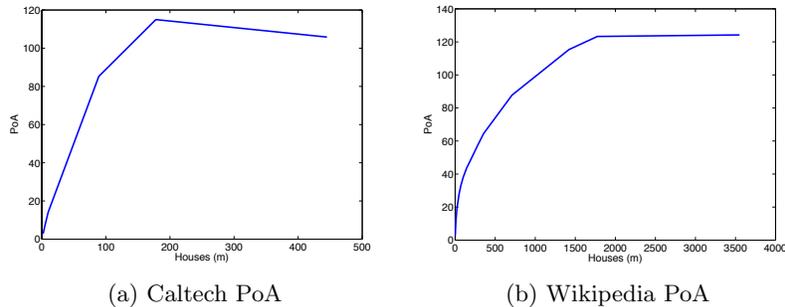


Fig. 3: Illustration of price of anarchy bounds in Theorem 3 for Caltech and Wikipedia networks.

Though Theorem 3 is tight, it is unclear at this point whether Theorem 4 is also tight. However, a slight modification of the above example does show that it has the correct asymptotics, i.e., there exists a family of examples that have price of anarchy $\Theta(m\gamma_m^*q_{max}w_{max}D_{\Delta}^{-1})$.

A first observation one can make about these theorems is that the price of anarchy has no direct dependence on the number of students. This is an important practical observation since the number of houses is typically small, while the number of students can be quite large (similar phenomena hold in many other many-to-one matching markets). In contrast, the theorems highlight that the degree of heterogeneity in quotas, network edge weights, and house valuations all significantly impact inefficiency.

A second remark about the theorems is that the only dependence on the social network is through γ_m^* , which measures how well the graph can be “clustered” into m groups. An important note about γ_m^* is that it is highly dependent on m , and tends to shrink quickly as m grows. A consequence of this behavior is that the price of anarchy is not actually linear in m in Theorems 3 and 4, as it may first appear, it turns out to be sublinear. This is illustrated in the context of real social network data in Figures 3a and 3b. Note that as we are increasing the number of houses, we are in fact creating finer allowable partitions of the network. The social networks used to generate the above plots are described in detail in [5].

Next, let us consider the impact of peer effects on the price of anarchy. Considering the simple setting of Theorem 3, we see that if there were no peer effects, this would be equivalent to setting $w(s, t) = 0$ for all s, t . This would imply that $\gamma_m^* = 0$, and so the price of anarchy is one. Thus, another interpretation of the price of anarchy in Theorem 3 is the efficiency lost as a result of peer effects.

5 Concluding remarks

In this paper we have focused on many-to-one matchings with peer effects and complementarities. Typically, results on this topic tend to be negative, either proving that stable matchings may not exist, e.g., [26,28], or that stable matchings are computationally difficult to find, e.g., [25].

In this paper, our goal has been to provide positive results. To this end, we focus on the case when peer effects are the result of an underlying social network, and this restriction on the form of the peer effects allows us to prove that a two-sided exchange-stable matching always exists. Further, we provide bounds on the maximal inefficiency (price of anarchy) of any exchange-stable matching and show how this inefficiency depends on the clustering properties of the social network graph. Interestingly, in our context the price of anarchy has a dual interpretation as characterizing the degree of inefficiency caused by peer effects.

There are numerous examples of many-to-one matchings where the results in this paper can provide insight; one of particular interest to us is the matching of incoming undergraduates to residential houses which happens yearly at Caltech and other universities. Currently incoming students only report a preference order for houses, and so are incentivized to collude with friends and not reveal their true preferences. For such settings, the results in this paper highlight the importance of having students report not only their preference order on houses, but also a list of friends with whom they would like to be matched. Using a combination of these factors, the algorithms described in [5] and efficiency bounds presented in this paper provide a promising approach, for this specific market as well as any general market where peer effects change the space of stable matchings.

The results in the current paper represent only a starting point for research into the interaction of social networks and many-to-one matchings. There are a number of simplifying assumptions in this work which would be interesting to relax. For example, the efficiency bounds we have proven consider only a one-sided market, where students rate houses similarly and quotas are exactly met. These assumptions are key to providing simpler bounds, and they certainly are valid in some matching markets; however relaxing these assumptions would broaden the applicability of the work greatly.

References

1. J. Alcalde. Exchange-proofness or divorce-proofness? stability in one-sided matching markets. *Rev. Econ. Design*, 1:275–287, 1994.
2. E. Anshelevich, S. Das, and Y Naamad. Anarchy, stability, and utopia: Creating better matchings. In *SAGT*, pages 159–170, 2009.
3. M. Baccara, A. Imrohorglu, A. Wilson, and L. Yariv. A field study on matching with network externalities. 2009. Working paper.
4. P. Bajari and J. Fox. Measuring the efficiency of an fcc spectrum auction. Working Paper 11671, National Bureau of Economic Research, 2009.
5. E. Bodine-Baron, C. Lee, A. Chong, B. Hassibi, and A Wierman. Matching with friends: stability and peer effects in housing assignment. 2011. working paper, available on arxiv.
6. A. Bogomolnaia and M. Jackson. The stability of hedonic coalition structures. *Games Econ. Behav.*, 38(2):201 – 230, 2002.
7. Simina Brânzei and Kate Larson. Coalitional affinity games and the stability gap. In *IJCAI*, pages 1319–1320, 2009.
8. K. Cechlarova and D. Manlove. On the complexity of exchange-stable roommates. *DAM*, 116(3):279 – 287, 2002.
9. K. Cechlarova and D. Manlove. The exchange-stable marriage problem. *DAM*, 152(1-3):109 – 122, 2005.

10. B. Dutta and J. Masso. Stability of matchings when individuals have preferences over colleagues. *J. Econ. Theory*, 75(2):464–475, 1997.
11. F. Echenique and M. Yenmez. A solution to matching with preferences over colleagues. *Games Econ. Behav.*, 59(1):46–71, 2007.
12. J. Fox. Estimating matching games with transfers. 2010. Working paper.
13. D. Gale and L.S. Shapley. College admissions and the stability of marriage. *AMM*, 69(1):9–15, 1962.
14. I.E. Hafalir. Stability of marriage with externalities. *Int. J. Game Theory*, 37:353–370, 2008.
15. R. Irving. Stable matching problems with exchange restrictions. *J. Comb. Opt.*, 16:344–360, 2008.
16. M. Jackson. *Social and Economic Networks*. Princeton University Press, 2008.
17. R. Kannan, S. Vempala, and A. Vetta. On clusterings: Good, bad and spectral. *J. ACM*, 51(3):497–515, 2004.
18. B. Klaus and F. Klijn. Stable matchings and preferences of couples. *J. Econ. Theory*, 121:75–106, 2005.
19. B. Klaus and F. Klijn. Paths to stability for matching markets with couples. *Games Econ. Behav.*, 58:154–171, 2007.
20. F. Kojima and P. Pathak. Incentives and stability in large two-sided matching markets. *Amer. Econ. Rev.*, 99(3):608–627, 2009.
21. M. Ostrovsky. Stability in supply chain networks. *Amer. Econ. Rev.*, 98(3):897–923, 2008.
22. M. Pycia. Many-to-one matching with complementarities and peer effects. 2007. Working paper.
23. F. Radicchi, C. Castellano, F. Cecconi, V. Loreto, and D. Parisi. Defining and identifying communities in networks. *PNAS*, 101(9):2658–2663, 2004.
24. P. Revilla. Many-to-one matching when colleagues matter. 2004. Working paper.
25. E. Ronn. Np-complete stable matching problems. *J. Alg.*, 11:285–304, 1990.
26. A.E. Roth. The evolution of the labor market for medical interns and residents: A case study in game theory. *J. of Polit. Econ.*, 92:991–1016, 1984.
27. A.E. Roth, U.G. Rothblum, and J.H. Vande Vate. Stable matchings, optimal assignments and linear programming. *MOR*, 18(4):803–828, 1993.
28. A.E. Roth and M. Sotomayor. *Two-sided matching: A study in game-theoretic modeling and analysis*. Cambridge University Press, 1990.
29. A.E. Roth and J.H. Vande Vate. Random paths to stability in two-sided matching. *Econometrica*, 58:1475–1480, 1990.
30. H. Sasaki and M. Toda. Two-sided matching problems with externalities. *J. Econ. Theory*, 70(1):93–108, 1996.
31. J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Trans. of Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.
32. S. Wasserman and K. Faust. *Social Network Analysis: Methods and Applications*. Cambridge University Press, 1994.