Allocating Group Housing

Justin Burkett, Francis X. Flanagan, and Amanda L. Griffith*

December 30, 2015

Abstract

We study mechanisms for allocating housing to students with non-trivial preferences over roommates and rooms. In this environment, the mechanism may distort students' preferences over roommates. Compared to certain distortive mechanisms, a non-distortive one always has a stable allocation in our model, selects stable outcomes that are ex ante more efficient under assumptions on the distribution of roommate values, and is less prone to manipulation by strategic student organizations. With outcome data we test the final prediction and conclude that student organizations are successfully manipulating the mechanism, supporting our hypothesis that the theoretical concerns have real implications.

I Introduction

We study the allocation of campus housing to university students. We examine a case in which students are free to choose their own roommates, without interference by the university, but rooms are allocated to roommate groups via a centralized, university controlled mechanism. Since students

*Burkett: Department of Economics, Wake Forest University, Box 7505, Winston-Salem, NC 27109; burketje@wfu.edu. Flanagan: Department of Economics, Wake Forest University, Box 7505, Winston-Salem, NC 27109; flanagfx@wfu.edu. Griffith: Department of Economics, Wake Forest University, Box 7505, Winston-Salem, NC 27109; griffial@wfu.edu.
We thank Thayer Morrill for helpful feedback.

likely care about both their roommates and their rooms, the housing allocation mechanism can affect the students' choice of roommates, and hence the quality of roommate matches.

One mechanism, which we study empirically, provides a simple example of how this may happen. In this mechanism, referred to later as ϕ_L, students individually receive lottery numbers which are used determine their priority for selecting housing on campus. The students are not committed to their roommates when the lottery numbers are distributed, and it is the lowest (best) lottery number in a group of roommates that determines the group's priority for housing selection. As long as students place some weight on where they live, the students receiving better lottery numbers thus become more attractive to others as roommates.

To what extent should such a distortion in roommate preferences lead to worse (or better) overall outcomes? To explore this question, we propose an alternative mechanism (ϕ_R) which removes any potential distortion in roommate preferences by waiting to assign priorities to pairs after they have formed. We then compare this alternative to a class of mechanisms that announce a mapping from matches to priority orders before pairs are finalized, of which ϕ_L is a member, using as criteria desirable characteristics such as stability and efficiency of an allocation, and strategy-proofness of a mechanism.

Our theoretical results show that the amount of information revealed to students about how their roommate groups influence the room distribution can have significant implications for each of these characteristics. We show that when students know with certainty which room they will receive as a function of which roommates they choose (i) it is less likely that a stable roommate grouping exists, (ii) the mechanism is ex ante less efficient under a regularity condition, and (iii) the mechanism is more susceptible to manipulation by organized groups of students, all relative to a mechanism which does not reveal information about the room distribution as a function of roommate groups. We emphasize that the efficiency result in (ii) is restrictive, because it requires strong assumptions on preferences, and we provide counterexamples showing that the ranking may fail when these assumptions are relaxed.

These results derive from a model designed to approximate the environment in which our data are generated. In the model each pair of students receives a common value from living together, all rooms are of capacity two, and all students have the same ordinal preferences over rooms.[1]

Within this model we say a roommate match is stable if no two students would like to leave their roommates to become roommates with each other and use the room that they would receive in the new match (see Definition 1). The stability of a match thus depends on the mechanism. We first show that the existence of a stable match is not guaranteed for the mechanisms we consider. This is simply a consequence of the knowledge of priorities affecting students' preferences over roommates. In the model, under ϕ_L, a student receiving a higher priority makes this student more desirable to all other students. We show that this shift in preferences can lead to there being no stable match with respect to ϕ_L. In contrast, a stable allocation always exists in our model if students form pairs before any information about priorities is revealed, as they do under ϕ_R. In this sense, a mechanism like ϕ_L may introduce instability.

The next objective we consider is constrained efficiency, taking stability as the constraint on the mechanism designer. In our model agents have cardinal payoffs for each possible allocation, and hence the efficiency criterion we use is total social surplus.[2] A simple example shows that in general the social surplus is ambiguous within this class of mechanisms. However, intuition suggests that under the stronger assumption that students have common cardinal preferences over rooms, making social surplus constant in the assignment of *rooms*, a mechanism that does not distort roommate

[1] This last assumption is supported by our data. In a survey of the students we asked for their preferences over different residence halls on campus, and we found that the preferences were consistent across students and with the observed order in which buildings were selected.

[2] The set of Pareto efficient allocations is large and therefore not a useful way to compare mechanisms.

preferences away from their true values (e.g., ϕ_R) may perform better on average than one that does.[3] Assuming independently and identically distributed (i.i.d.) roommate values, we provide sufficient conditions for ϕ_R to be more efficient than the class of mechanisms that includes ϕ_L. Mechanisms which distort roommate preferences cause some pairs to be of lower quality on average. Interestingly, we provide a counterexample showing that the i.i.d. roommate values and common cardinal room values assumptions are not sufficient to guarantee expected surplus maximization. Thus a distortive mechanism like ϕ_L *may* perform better than a non-distortive one even in an environment without ex ante asymmetries.[4]

The final objective, a type of equity objective, is motivated by the belief that students in certain organizations are able to, by virtue of their membership in the organization, secure better housing for themselves.[5] We show that ϕ_L is more susceptible to such manipulation relative to ϕ_R. Under ϕ_L, a roommate pair may choose housing according to the best priority of either individual roommate. Thus this rule can be manipulated by student organizations that are able to reorganize themselves to make better use of priorities. For example, suppose that students s_1, s_2, s_3, and s_4 are members of an organization. If s_1 and s_2 are close friends, and also receive the two best lottery numbers, in some sense a valuable number is wasted if s_1 and s_2 choose to live together. By placing s_1 with, for example, s_3, and s_2 with s_4, the organization can guarantee that all four members occupy one of the two best rooms. Whether or not this improves the allocation for the organization depends on how willing each student is to live with the person they are placed with, but in principle the organization could compensate s_2 for any loss in surplus associated with splitting from s_1.[6]

[3] The ϕ_R mechanism frequently selects inefficient allocations due to the stability requirement, so it is not obvious that this must be true.

[4] It is easy to create examples with ex ante asymmetries in which a properly chosen distortive mechanism performs better (see Footnote 14).

[5] Such organizations may include fraternities and sororities, athletics teams, clubs, etc.

[6] In our model we make explicit assumptions about student preferences which eliminate this problem, but these assumptions can be interpreted as implicit assumptions that organizations have some exogenous means to compensate members. For example, the

Our empirical results strongly suggest that certain student organizations, specifically certain Greek Student Organizations (GSOs),[7] are manipulating their roommate groups in this manner. We observe the priority assigned to each student, along with the priority used by the roommate group that the student is a member of. Figure 1 shows the density estimates for the used priorities broken up by class year[8] and membership in a "strategic" organization. Under the current mechanism priorities are assigned randomly, but members of a higher class year are guaranteed a higher priority than members of a lower class year, thus these ranges are disjoint. It is clear from the figure that second- and third-year students in the strategic organizations are using priorities that are much better than those outside of the organizations.[9] Our detailed empirical analysis is in Section V.1.

The remainder of the paper is organized as follows. In Section II we discuss related literature, followed by the introduction of the model in Section III. Section III also introduces the various mechanisms in question, as well as discusses their stability definitions and properties. Section IV compares the efficiency of each mechanism. The first part of Section V explores the theoretical implications of strategic student organizations. In Section V.1.1, we describe the data we use, while Section V.1.2 presents our initial empirical results. Section V.1.3 gives the evidence that these outcomes are the result of strategic behavior, and Section VI concludes. All proofs are in the Appendix.

compensation may come from future interactions with the mechanism, or from any number of financial, social, or scholarly interactions outside of the mechanism.

[7]These are the organizations that we can reliably identify in the data.

[8]First-year students are not included because they are assigned housing via a different process.

[9]Statistical tests support the pattern seen in the graph. Using Kolmogorov-Smirnoff tests at the 1% level, we reject the null hypothesis that the distribution of used priorities is the same between strategic organizations and all others for either the second- or third-year students. We also can confirm for each of the second-, third- and fourth-year students that there is no significant difference between the priorities received by the members of the strategic organizations and all others, using Kolmogorov-Smirnoff tests at the 5% level, so the observed effect is due entirely to how the priorities are used.

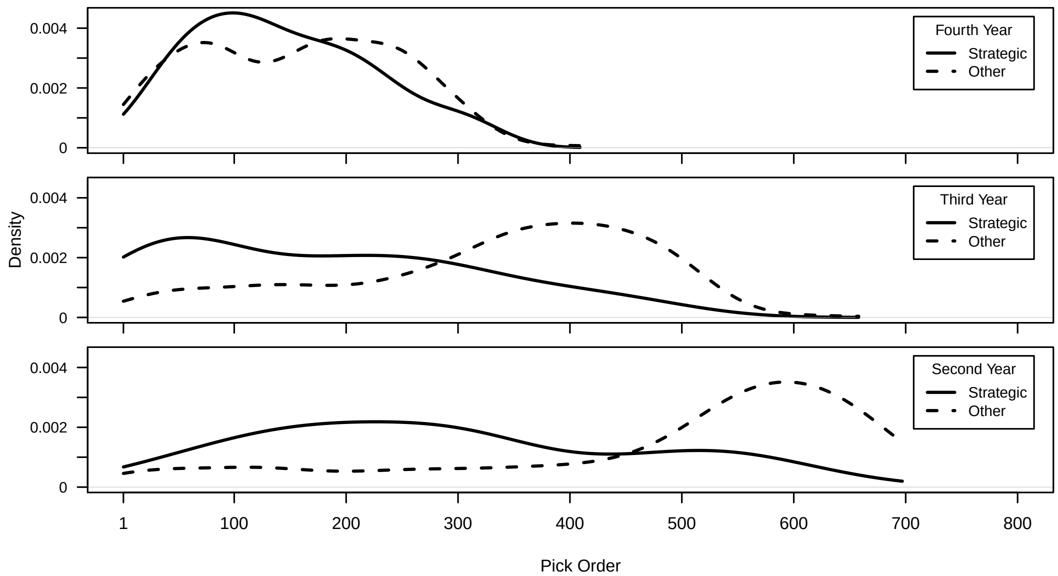

Figure 1: Estimated Densities of Pick Orders

II Related Literature

The theoretical model we introduce in this paper unites two distinct subjects in the matching literature: housing allocation problems and roommate problems. The *house allocation problem*, introduced by Hylland and Zeckhauser (1979), is characterized by a set of n agents who each must be assigned exactly one object (house) from a set of n indivisible objects. This is closely related to the *housing market problem*, introduced by Shapley and Scarf (1974), which is a house allocation problem in which each agent initially owns one object. Abdulkadiroglu and Sönmez (1999) introduce a hybrid model, the *house allocation problem with existing tenants*, in which some, but not all, objects are initially owned. Our model begins as a housing allocation market without existing tenants: initially the objects (rooms) are unassigned. The main difference between previous literature on house allocation and our model is that we wish to assign objects to pairs, rather than individuals. Dogan et al. (2011) looks at the housing market problem with couples, however the authors assume that couples are determined exogenously. In the present paper we are explicitly interested in the coalition formation part of the game, which is why our model is also closely related to the roommates problem.

The *roommates problem*, introduced by Gale and Shapley (1962), is a one-sided matching problem: there are $2n$ agents who must form n pairs. Gale and Shapley introduced this as an example of a market in which there is no guaranteed stable match. To avoid this issue students in our model always receive the same payoff from forming a coalition together; this can be interpreted as a roommate pair producing some surplus which is split evenly, or as student preferences over roommates being represented by a semimetric distance.[10] Student preferences defined by a metric was originally proposed by Bartholdi and Trick (1986) as a restriction on preferences for which there is always a stable match, and Chung (2000) shows that this is part of a more general category of preferences, called *no odd rings* preferences, which always guarantee a stable match in the roommates problem. Likewise, Pycia (2012) shows that if the set of feasible preference profiles is rich enough, then roommate preferences must be pairwise aligned to guarantee a stable coalition structure. Pycia also shows that evenly splitting total surplus generated by a coalition satisfies this condition. For similar reasons we restrict our theoretical results to the situation in which students may only form pairs, rather than larger coalitions. Although our empirical results are collected from a market in which students are allowed to form coalitions larger than two, Arkin et al. (2009) show that even if preferences are generated by distances in a metric space stability is not guaranteed if the size of the roommate coalition is expanded beyond couples.

Lastly, our paper fits into the large and diverse literature on collusion in games (Cf. Graham and Marshall (1987), McAfee and McMillan (1992), Pesendorfer (2000)). Although we do not explicitly model transferable utility, we sometimes allow certain organizations to dictate their members' behavior, with the interpretation that they have some exogenous means of transferring utility among members. However, unlike the related literature,

[10] A semimetric is a metric which does not necessarily satisfy the triangle inequality.

strategic play and coordination are not explicitly against the rules of the mechanisms we study.

III Model

Let R be the finite set of n rooms, and S be the set of $2n$ students. A *(roommate) pair* is a subset of S of size two, denoted (s_i, s_j) with $i \neq j$. A *match* μ is a partition of S into pairs. We refer to student i's roommate in match μ as $\mu(s_i)$, and we have the usual requirement that $s_j = \mu(s_i) \iff s_i = \mu(s_j)$. We refer to the set of feasible matches as $\mathcal{M}$.

Student payoffs are a combination of payoffs over roommates and rooms. *Roommate values* are set according to $w : S^2 \to \mathbb{R}^+$, which assigns a symmetric value to each constituent of a roommate pair: for any $s, s' \in S$, $w(s, s') = w(s', s) \geq 0$. We impose no further restrictions on this function. Observe that the symmetry requirement imposed on w does not imply that if student j is student i's most preferred roommate then student i is student j's most preferred roommate. In other words, it allows for nontrivial differences among students in their preference orderings.

For each student s, *room values* are set according to $v_s : R \to \mathbb{R}^+$. We assume throughout that students have common ordinal preferences over rooms. That is, for any $s, s' \in S$, $r, r' \in R$, $v_s(r) > v_s(r') \iff v_{s'}(r) > v_{s'}(r')$. To avoid triviality, we also assume that there exist two rooms $r, r' \in R$ such that $v_s(r) > v_s(r')$; that is, students are not indifferent between all rooms.[11] Therefore the complete payoff function for student s for being matched with roommate s' and room r is $u_s(r, s') = v_s(r) + w(s, s')$.

Formally, an *allocation* $A = (\mu, M_\mu)$ is a match, μ, and a bijection, $M_\mu : \mu \to R$, where $M_\mu(s, s') = r$ is the room assigned to pair (s, s'). The pri-

[11] Although the common order assumption is strong, our data suggest that students do indeed have similar ordinal preferences over rooms. In a survey conducted prior to the students selecting housing, we asked which suite sizes they were considering choosing, and conditional on these responses we asked them to rank all of the available buildings. These survey responses were consistent with the hypothesis that most students have very similar ordinal preferences over buildings. These reported preferences also closely correspond to the order in which the suites disappear during the selection process.

mary object of interest for our paper is not the allocation, but rather the mechanism through which the allocation is created. The following section introduces the class of mechanisms we study and introduces stability.

III.1 Mechanisms and Stability

A *room allocation mechanism* (or simply a mechanism) is a procedure used to assign students, or student pairs, to rooms. Anecdotal evidence suggests that the most common way colleges distribute rooms to students is to assign priorities via some sort of lottery, and then to allow students to choose rooms according to these priorities.

One variant of a "lottery mechanism" is a mechanism that we refer to as ϕ_L. Under the ϕ_L mechanism, each individual receives a lottery number, and the pick order of a pair is determined by the best (lowest) lottery number in the pair. More precisely, under ϕ_L (s_i, s_j) has higher priority when selecting rooms than (s_k, s_l) if and only if the lower of the two lottery numbers received by s_i and s_j is smaller than the lower of the two received by s_k and s_l. With these rules in place and assuming common knowledge of the lottery numbers received by all of the students, each student can determine their priority given a match. It is important that under ϕ_L students are not committed to a match when lottery numbers are distributed, and hence we allow for the match to adjust in response to different assignments of lottery numbers.

Generalizing, one can think of ϕ_L as a mapping from matches to priority orders and define a class of mechanisms containing all such mappings. If $\mathcal{M}$ is the set of feasible matches, a generic member of this class is a bijection from feasible matches to priority orders, $\phi : \mathcal{M} \rightarrow \{1, \ldots, n\}$. We refer to this class as the Φ_L class of mechanisms. The ϕ_L mechanism is our primary object of study and corresponds to the mechanism from which our data are derived, but most of the theoretical results apply to any $\phi \in \Phi_L$.

A natural alternative to mechanisms in Φ_L that we consider is the one that assigns a priority order to a match *after* the students commit to that match. This mechanism, which we refer to as ϕ_R, determines the priority of a roommate pair uniformly randomly after a match forms. One could implement this mechanism by assigning lottery numbers to pairs after pairs form, instead of students individually. In some sense, the mechanisms in Φ_L and ϕ_R represent two extremes. The mechanisms in Φ_L are deterministic in that the mapping from matches to priority orders is known when the students form matches, while the mapping is unknown at that time under ϕ_R. Note that at this point we make no assumptions on how the lottery numbers are assigned for any $\phi \in \Phi_L$.

The stability notion we use for a match, given in Definition 1, is similar to pairwise stability in other matching models. Informally, a match is not stable if there exist two students who would rather be roommates with each other than their current roommates. Such a pair is called a *blocking pair* or is said to *block* the original match. This corresponds to the usual notion of pairwise stability; however, it is necessary in this model to specify what rooms the students expect to occupy after forming a blocking pair. We require the additional condition that the blocking pair foresees, one step ahead, how its expected room assignment will change after the block. To make this a feasible calculation, we assume that if two students block μ, then the roommates who are "left" by the blocking pair will pair together in the subsequent match. That is, for the roommate pairs $(s,s'),(t,t') \in \mu$, if (s,t) blocks μ, then the new roommate match after the block will be $\mu' = \mu \setminus ((s,s') \cup (t,t')) \cup ((s,t) \cup (s',t'))$. Because of this assumption it is often useful to look at matches which differ in this manner: we say two matches μ and μ' are *adjacent*, written $\mu \leftrightarrow \mu'$, if exactly two pairs are different between them: $\mu \leftrightarrow \mu' \iff |\mu \setminus \mu'| = 2$.

For any mechanism ϕ and any match μ, before the priority order is realized, let the expected payoff for student s be $U_s(\phi,\mu) = V_s(\phi,\mu) + w(s,\mu(s))$, where $V_s(\phi,\mu)$ is the expected room value and $w(s,\mu(s))$ is the roommate value, which is independent of the mechanism. Since the expected payoffs depend on the mechanism, the stability of a match will also depend on the mechanism, as the following definition explains.

Definition 1. Given mechanism ϕ, a match μ is *stable with respect to* ϕ if, for any two pairs (s_1,s_2) and (t_1,t_2) in μ and any adjacent μ' with $(s_1,t_1) \in \mu'$, we have that $U_{s_i}(\phi,\mu') > U_{s_i}(\phi,\mu)$ implies $U_{t_i}(\phi,\mu') \leq U_{t_i}(\phi,\mu)$ for $i \in \{1,2\}$.

We refer to both the match and the resulting allocation as being stable. Theorem 1 shows that under our assumptions a stable match always exists under ϕ_R; however, the same is not true of every $\phi \in \Phi_L$, as the following example illustrates.

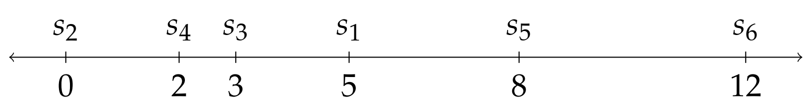

Figure 2: Example with No Stable Match

Example 1. Let roommate values be represented by some constant minus the distance between students on the number line in Figure 2, and let the mechanism be ϕ_L. Let each student have the same room value function, with $v(r_1) = 1.5$, $v(r_2) = 1.25$ and $v(r_3) = 0$, and suppose that their indices represent the lottery numbers received. First look at the match $\mu = \{(s_1,s_5),(s_2,s_6),(s_3,s_4)\}$, listed in order of each pair's priority. This match is the unique stable match under ϕ_R. However, since $v_1(r_1) + w(s_1,s_3) > v_1(r_1) + w(s_1,s_5)$ and $v_3(r_1) + w(s_1,s_3) > v_3(r_3) + w(s_3,s_4)$, the pair (s_1,s_3) blocks this match under ϕ_L. The resulting match, $\mu' = \{(s_1,s_3),(s_2,s_6),(s_4,s_5)\}$, is itself blocked by (s_2,s_4), and one can check that all other possible matches are also blocked.[12]

The essential problem leading to non-existence of a stable match is that a pair may receive different rooms in different matches, leading that pair's utility to depend on the arrangement of the other students. Under ϕ_R however, the expected utility a student assigns to a given pair is always the same. As a consequence, it is straightforward to prove the following.

[12]Blocking pairs are one of $(s_1,s_2),(s_1,s_3),(s_1,s_4),(s_1,s_5),(s_2,s_4)$ and (s_3,s_4).

Theorem 1. *There exists a match which is stable with respect to* ϕ_R.

This result follows from Bartholdi and Trick (1986), Chung (2000), and Pycia (2012): under ϕ_R, $V_s(\phi_R, \mu) = V_s(\phi_R, \mu')$ when $\mu(s) = \mu'(s)$. Therefore if $w(s, \mu(s)) \geq w(s, s')$ for all $s' \neq \mu(s)$, then neither s nor $\mu(s)$ can be part of a pair which blocks μ. Therefore we can construct a match which is stable under ϕ_R by successively identifying the highest value pairs not already formed and placing them in the match.

The possible non-existence of a stable match is not unique to ϕ_L. Our next result, a kind of impossibility result, shows that a stable match is not guaranteed for any $\phi \in \Phi_L$.

Theorem 2. *If* $\phi \in \Phi_L$, *it is possible to construct roommate values so that no stable match exists under* ϕ.

IV Comparison of Mechanisms: Efficiency

A natural question in regard to this class of mechanisms is is there a way to rank the stable outcomes of each mechanism on the basis of efficiency? Since Pareto efficiency does not discriminate between mechanisms in our model, we study the implications of using total surplus as a measure of efficiency. Thus we say that mechanism ϕ is *more efficient* than mechanism ϕ' if, for any allocations (μ, M_μ) and $(\mu', M_{\mu'})$ such that μ is stable with respect to ϕ and μ' is stable with respect to ϕ', we have that $\sum_{s \in S} u_s(M_\mu(s, \mu(s)), \mu(s)) > \sum_{s \in S} u_s(M_{\mu'}(s, \mu'(s)), \mu'(s))$. Since we only compare stable matches, when a mechanism has no stable match it is unclear how to compare surplus. One might favor the mechanism with the stable match in this case, or decide that surplus is incomparable between the two. Without a strong opinion on the matter we will favor the former approach, and since we compare against ϕ_R, which always has a stable match, we are not concerned with the case where no stable match exists in either mechanism.

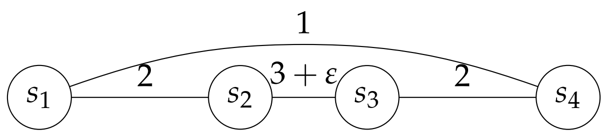

Figure 3: Spatial Distribution of Students

Our results in this section assume that all students have the same cardinal values over rooms. In this setting the match completely determines the total surplus. Given the restrictiveness of our definition of efficiency it is perhaps unsurprising that it turns out that no mechanism is strictly more efficient than all others, which we formalize in Theorem 3. The following example shows that ex post surplus is generally ambiguous.

Example 2. Let there be four students and two rooms, and let all students have common values over rooms, so $v_s(r) = v_{s'}(r) = v(r)$ for all s, s' and r. Let $v(r_1) = v(r_2) + \delta$, and let Figure 3 represent the relevant roommate values between students, with unlisted values set to 0. Notice that it is efficient for s_2 and s_3 to pair as long as $\varepsilon \in (0, \infty)$, and that these students will pair under ϕ_R as long as $\varepsilon \in (-1, \infty)$. That is, within the range $(-1, \infty)$ the unique stable match with respect to ϕ_R is $\mu = \{(s_2, s_3), (s_1, s_4)\}$, even though sometimes this is not efficient. Next consider ϕ_L, which assigns the following priority order to any match: $(s_1, s_i) \succ (s_j, s_k)$. Under this mechanism s_2 will pair with s_1 if $v(r_1) + 2 > v(r_2) + (3 + \varepsilon) \Rightarrow \delta > 1 + \varepsilon$, and the resulting match would be stable. Therefore, under ϕ_L, for large enough δ, when $\varepsilon \in (0, \infty)$ s_2 will choose to pair with s_1 even though it is both not efficient to do so and would not be done under ϕ_R. However, under ϕ_L, for large enough δ, when $\varepsilon \in (-1, 0)$ s_2 will choose to pair with s_1, resulting in the efficient match, but this would not be stable with respect to ϕ_R. Therefore neither mechanism is more efficient than the other.

This example also shows that efficiency and stability of the match are distinct notions, since a stable match (when one exists) is not necessarily efficient, and vice versa. The possible tension between stability and efficiency is a common theme in the matching literature,[13] thus it might be expected that this result holds.

Theorem 3. *No mechanism in* $\{\phi_R\} \cup \Phi_L$ *is more efficient than all others in this set.*

One problem with the mechanism ϕ_R is that it selects a stable match in a "greedy" fashion, always picking the highest value pair, then the next highest value pair, and so on. This procedure will clearly not always choose the match with the largest total surplus. Alternatives like ϕ_L (and any $\phi \in \Phi_L$) potentially distort students' preferences over their roommates by giving some students the ability to select better housing, and can lead to more efficient stable matches than ϕ_R as a result. However, ϕ_L does not change the greedy nature of the mechanism, as the highest value pairs will still tend to form irrespective of other pair's values. Therefore, intuitively, when comparing ϕ_R and ϕ_L we are comparing two greedy mechanisms, one of which is distortive, which suggests that ϕ_R may be more effective at selecting efficient matches at least ex ante under the common cardinal values for rooms assumption.

Surprisingly, a counterexample shows that even when roommate values are i.i.d. it is possible that the expected surplus of a stable match under ϕ_L is higher than the stable pair structure under ϕ_R. However, we show that if we further restrict the possible distributions over roommate values to be log-concave, then the intuition goes through: in this environment ϕ_R indeed yields higher expected surplus than any $\phi \in \Phi_L$.

To evaluate this claim formally, we add an initial stage to the model in which roommate values are drawn from some prior distribution. This prior only affects the mechanism designer's assessment of the expected outcome and has no impact on what is known by students when pairs are formed. Assume that the value associated with each pair is i.i.d. according to the absolutely continuous distribution function F (with density f) on $[0,\infty)$. We assume that students have common cardinal preferences over rooms,

[13] See for example Erdil and Ergin (2008) and Ergin and Sonmez (2006).

but we do not put any distributional restrictions on these values, allowing for any realization of room values.

Our first example illustrates the basic intuition behind why i.i.d. values would seem to make ϕ_R more efficient ex ante; however, it also suggests that the conclusion is not immediate.[14]

Example 3. Let there be $n = 6$ students, and let roommate values be drawn i.i.d. from $U(0,1)$. Let room values be common across all students such that $v(r_1) = 1$ and $v(r_2) = v(r_3) = 0$. Therefore $v(r_1) + w(s_i, s_j) > v(r_l) + w(s_i, s_k)$ for any $i \neq j, k$ and $l = 2, 3$. Under ϕ_R, the stable match includes the pair with the highest roommate value of the fifteen possible pairs, and the pair with the highest roommate value of the six remaining pairs. Thus the expected roommate surplus is $15/16 + (15/16)(6/7) + (1/2)(15/16)(6/7)$. Under ϕ_L, the pair including s_1 and her favorite roommate is in the stable match as well as the pair with the highest roommate value among the remaining six pairs. Thus the expected roommate surplus is $3/4 + 6/7 + (1/2)(6/7)$.

Importantly, under ϕ_L, the value of the "second" pair is not constrained by the value of the pair containing s_1. Therefore, although the expected total surplus under ϕ_R is higher, the expected value of the lowest value two roommate pairs is actually higher under ϕ_L. This leaves open the possibility that the expected surplus under ϕ_L may exceed that of ϕ_R, an example of which is the following.

Example 4. Let $n = 4$ and suppose F puts probability $1/2$ on the interval $I_a = (0, \varepsilon)$ and probability $1/2$ on the interval $I_b = (1 - \varepsilon, 1)$ for some small ε. Let room values be $v(r_1) = 1 - 2\varepsilon$ and $v(r_2) = 0$. Let μ_R^* (resp. μ_L^*)

[14] Without the assumption of identical distributions, it is straightforward to create examples where ϕ_L yields more expected surplus than ϕ_R. To see why note that ϕ_L would allow a student with an exceptionally high value for a particular room know with certainty if she would receive this room as a function of the match. Therefore this student can use this information while forming pairs, which will increase the probability that she receives this room relative to ϕ_R. Similarly, it is easy to construct examples where expected roommate surplus is higher if the mechanism designer disincentivizes certain high value pairs which can only form when all other pairs have much lower value.

be the match which is stable with respect to μ_R (μ_L).[15] Notice that either mechanism will always choose a match containing a roommate pair which has a value in I_b if one exists, and that if μ_L^* and μ_R^* contain the same number of roommate values drawn from I_b the difference in surplus is zero or approximately zero. Therefore there is only a significant difference in surplus between μ_L^* and μ_R^* if one of them has two roommate values from I_b while the other has only one. Finally, notice that whenever the roommate pair in μ_R^* which contains the student with the highest priority under ϕ_L has a value in I_a, then any match with two roommate pairs that each have a value in I_b will block μ_R^* under ϕ_L. It follows that in this setting ϕ_L, relative to ϕ_R, shifts some probability away from selecting a match with one roommate value in I_a and one in I_b, and towards a match with two roommate values in I_b, which increases expected surplus.

Note that by continuity these results hold for settings in which the room values are "almost" common across students, or where there are small differences in the roommate value distributions.

To show that our initial intuition holds in a more limited sense, we give sufficient conditions on the distribution of roommate values such that ϕ_R is more efficient in expectation than any $\phi \in \Phi_L$. If we add the restrictions that F is log-concave, then we can prove the result stated in Theorem 4. A log-concave distribution is one for which $\log F(\cdot)$ is concave, a property which is satisfied by a variety of commonly used distributions (Bergstrom and Bagnoli, 2005).

Theorem 4. *If roommate values are i.i.d. according to the log-concave distribution F, the expected surplus of the unique stable match under ϕ_R is greater than that of any stable match under an alternative $\phi \in \Phi_L$.*

We interpret these results to mean that while the alternatives to ϕ_R in Φ_L introduce distortions in the students' decisions over whom they want to live with, which would seem to favor the non-distortive ϕ_R, we cannot be confident that ϕ_R is more efficient unless the problem is sufficiently

[15] With $n = 4$ students a match which is stable with respect to ϕ_L always exists.

"regular". By regular, we are referring to the fact that log-concavity ensures that $x - \mathrm{E}[Y|Y \leq x]$ is an increasing function of x, which is used in the proof of Theorem 4. Intuitively, this condition places limits on how quickly the conditional expectation of the value of a pair can change.[16]

V Comparison of Mechanisms: Manipulation

Finally, a mechanism designer may be concerned with the ability to manipulate a mechanism. In this section we compare ϕ_L and ϕ_R and show that ϕ_L is more vulnerable to manipulation by organized groups of students.[17] The reason this is the case is that at the time students form pairs under mechanism ϕ_L there is a known one to one mapping from matches to priority order. That is, given any match every pair knows exactly what priority it will receive under ϕ_L. However, under ϕ_R, every pair has equal probability of receiving any priority, thus there is no ability to change the expected room value by changing roommate pairs. All results in this section are conditional on a match forming, but do not require that the reported match under ϕ_L is stable with respect to ϕ_L. Also, we will again assume for the theoretical results that students have common cardinal values over rooms.

Let $\Omega \subset S$ be a set of students such that $|\Omega| = 2m$ for some $2 \leq m < n/2$. Ω is a *society* if it satisfies the following conditions: (i) for all $s, s' \in \Omega$, $w(s,s') = K$, where $K > \max_{r,r' \in R}\{v(r) - v(r')\}$, (ii) for all $s \in \Omega$ and $s' \in S \setminus \Omega$, $w(s,s') = 0$, and (iii) as a group, given mechanism ϕ, Ω attempts to maximize $\sum_{s \in \Omega} U_s(\phi, \mu)$. Since all members of a society always want to pair with each other under any mechanism, and since roommate value for the society is the same in any match in which all members of the society pair with another member of the society, maximizing the expected payoff of Ω

[16]Bergstrom and Bagnoli (2005) discusses the use of the assumption of log-concavity in economic models and derives some useful results to this effect.

[17]Due to the data we have collected, the groups we have in mind in this section are fraternities and sororities, referred to later in the paper as Greek Letter Organizations (GLOs). However there may be other relevant student organizations, such as athletics teams, students within the same major, graduating class, etc.

under ϕ is equivalent to maximizing the sum of the expected room values for each pair. We will call this sum the *room-welfare*, or simply *welfare*, of society Ω under ϕ.

Our first result is that, regardless of the behavior of $S \setminus \Omega$, it is more likely that Ω receives the best $k \leq m$ rooms under ϕ_L relative to ϕ_R.

Theorem 5. *The probability that Ω receives the first $k \leq m$ rooms under ϕ_L is greater than under ϕ_R, regardless of the strategies of $S \setminus \Omega$.*

Corollary 1. *If $S \setminus \Omega$ is a society it is more likely that Ω receives the $k = m$ worst rooms under ϕ_L than under ϕ_R.*

These results show that even in a perfectly symmetric setting in which the set of students is split into two equally sized societies, the strategic play of each society does not "cancel out" the other's actions. Thus ϕ_L results in more extreme room distributions relative to the ϕ_R. However, receiving any specific set of rooms may be a low probability event. The following result shows that for two possible sets of rooms for Ω, it is more likely that Ω receives the set with the better worst ranked room, as long as the worst ranked room in the set is not the worst ranked room in the full set of rooms. This result assumes that $S \setminus \Omega$ is a single society, which, as Theorem 7 shows, can be thought of as a "worst case scenario" for Ω.

Theorem 6. *Under ϕ_L, if $\Omega \setminus S$ is a society, then for any two sets of rooms R' and R'' such that $|R'| = |R''| = m$, and $\min_{r \in R'} v(r) > \min_{r \in R''} v(r) > \min_{r \in R} v(r)$, the probability that the students in Ω occupy the set of rooms R' is greater than the probability that they occupy the set of rooms R''.*

Finally, it should be expected that competing against a single, well organized society is worse for Ω than competing against smaller and less organized groups. The following results formalize this intuition, but first we must define what we mean when we refer to smaller organizations. We assume each student can be a member of at most one society. Thus the set of all societies, together with the set of students not in any society, can be described as a partition of S. If P is a partition of S into societies and students not in any society such that $\Omega \in P$, we may say that Ω *competes against* P. For any two partitions of $S \setminus \Omega$, P and Q, P is *finer* than Q if every element of P is a subset of some element in Q.

Theorem 7. *Let P and B be distinct partitions of S into societies such that $\Omega \in P \cap Q$. If P is finer than Q, then under ϕ_L the expected room-welfare of Ω is strictly higher when competing against P rather than Q.*

Corollary 2. *Let P and Q be partitions of S such that there exist societies $\Omega, \Omega' \in P$ and $\Omega'' \in Q$ such that $\Omega \cup \Omega' = \Omega'' \neq S$, and $P \setminus \{\Omega, \Omega'\} \cup \{\Omega''\} = Q$. Then under ϕ_L the sum of the expected room-welfares of Ω and Ω' when each are competing against P is strictly lower than the expected room-welfare of Ω'' when it competes against Q.*

Therefore a larger society has higher expected welfare than the total welfare of an equivalent sized set of students partitioned into multiple smaller societies. Similarly, if a society were less organized, less concerned about room-welfare, or more naive about the workings of the mechanism, then its expected room welfare would be lower. This is an immediate consequence of the fact that the total room-welfare is the same under any allocation, so any set of students that does not attempt to improve room-welfare will necessarily reduce its own room-welfare, and increase other groups' room-welfare, relative to a situation in which it tries to maximize room-welfare. The following section describes the data and empirical results used to test the theoretical results on manipulation.

V.1 Evidence of Manipulation by Organizations

In ϕ_L, where lottery numbers are assigned to individuals and the best lottery number in a pair determines its priority for housing selection, the potential benefits to coordination across pairs, at least in terms of the ability to improve priorities, are obvious. What is less obvious is the extent to which the societies, as we refer to them in the last section, are able to effectively coordinate. Presumably, in order for a society to restructure its roommate groups[18] to improve its overall priorities it would need to be able to offer some sort of compensation to any members that are harmed;[19] however, in the context of a mechanism with very limited future interactions and no exchange of money, implementing such transfers seems potentially difficult.

We present evidence that despite these difficulties several organizations are able to effectively coordinate and provide organizational rents to their members. We first describe the data and give more detail about the mechanism in use, before presenting our main empirical results.

V.1.1 Description of Data and Details of Allocation Procedure

Our data come from a housing mechanism in use at a medium-sized private university with undergraduate enrollment of roughly 1,200 students per year. On-campus university housing is available for undergraduate students in any year of study. The incoming first-year students (we use the year of student to denote their upcoming year when they choose housing) are assigned housing using a process that is distinct from the other students, thus they are not included in any of the analyses in this paper. Students are required to live on campus for three years, but many choose to live on campus for all four. The second-, third- and fourth-year students (there are no fifth-year or older students in our data) who live on campus select housing from the centralized process that is the focus of this study. Housing is available for every group size from 1 to 11, with the most common sizes being two, four, and six. We refer to all of these as suites in the remainder of the paper, even though some may be more properly called rooms.

About three weeks before students choose where they will live each student is assigned a lottery number. These lottery numbers determine

[18]As noted, in our data students form groups ranging from size 1 to 11. In this section we will refer to these coalitions simply as "groups."

[19]For example, it could be that the society makes promises to these students about how they will be treated in future iterations of the assignment process, or that the society is able to provide external compensation through more favorable treatment in other activities.

the date and time at which they may log into an online system to select housing. They are informed of their login time, but not the lottery number, which has no extra significance. Between 20 and 100 lottery numbers are associated with each selection time, and the login times are spread over five days. The lottery numbers are partitioned into ranges corresponding to the number of semesters a student has lived on campus, and are uniformly distributed within these ranges. For example, two students with four semesters of credit have the same probability of receiving any of the lottery numbers in the range allocated to four-semester students. All of the lottery numbers for n-semester students are better than any lottery number received by an m-semester student if $n > m$. The corresponding login times for the n-semester students are no later than the login times for the m-semester students.[20] Because two different login times might accommodate different numbers of students, the login times are not uniformly distributed within the semester groups, although it is still true that two students with the same number of semesters have the same chances of receiving each of the corresponding login times.

The data we have correspond to the process run in the spring of 2014 for housing in the 2014–2015 academic year. The dataset includes every student that selected housing in that year. In addition to the lottery number of each student and their corresponding assigned login time, we observe the following demographic information for each student: sex, number of semesters, ethnicity and race, indicators for athlete and foreign students, grade point average (GPA), and membership in specific "Greek Letter Organizations".

We construct two variables. First, we rank the selection times from earliest (1) to latest (57) and assign each lottery number in the data to the rank of the corresponding selection time. We refer to this number as the *lottery rank* in our analysis. Any two students will have different lottery

[20]The five-semester students share the latest login time of the six-semester students, the three-semester students share latest login time of the four-semester students, and the one-semester students share the latest login time of the two-semester students. These are the only cases where students with differing numbers of semesters on campus share times.

numbers, but may share the same lottery rank, meaning that they can log in at the same time. We also observe the precise time at which housing was selected for a group of students. We count the number of students who pick ahead of each student to construct *pick order*. This variable takes the same value for all members of the same suite. The students have no control over their own lottery rank but may influence their pick order if they do not select housing at their earliest opportunity.

Upon logging into the selection system, a student can select any of the available suites. The selection occurs in real-time and is effectively final,[21] so anyone who has the ability to log in but has not yet made a selection can remove an available suite from the pool. Groups select suites by giving permission to one person in the group to register all of the group members for a suite. This means that the group can choose a suite as soon as any one of its members can, which then implies that the only lottery rank that impacts the room selection is the lowest (best) one. Groups do not reveal to the university that they want to live together until a suite is selected for them, so while in principle groups may change up until the time they select housing, we only observe the final groupings into suites.

There are three other ways that upper-class students may select a room on-campus. The first is theme housing: students may form a group corresponding in size to one of a few designated "theme houses" and propose a theme for their living area. If chosen, these students do not enter the lottery as described above. Second, a number of the "Greek Letter Organizations" have set blocks of rooms within on-campus housing that they may utilize for their members. These organizations are required to indicate which of their members will live in these blocks prior to the release of the lottery times and numbers, and therefore these students are also not involved in the lottery process.[22] Lastly, there are a small number of

[21] We are told that there are rare exceptions to this rule, but we have no way of measuring how rare they are in the data we have.

[22] Although they are required to do this prior to the lottery, there are incentives at this stage to strategically allocate members to blocks. Simply putting younger students in the blocks will improve the expected lottery times received, for example. We do not have

students that receive accommodations in the selection process for certain health and physical disability issues. These students do pick their housing during the lottery process but are given priority, and some restrictions, on what housing is available. As none of these three groups picks housing in the typical way just described, we omit these students from our analysis.

To obtain further information on the preferences of the students that would be involved in the lottery, the authors conducted a survey prior to the beginning of the selection process in the spring of 2014. The survey was released to all students eligible for the lottery at a point when they were aware of their possible selection times, but had not yet selected a suite. Students were asked to report what size suites they were interested in, and to rank the buildings with available suites of that size. The survey respondents were also asked to report how many hours they had spent preparing for the housing lottery, as well as their membership in on-campus activities such as clubs and other organizations, and how much time they spend weekly with these organizations. 442 students responded, and these data were matched to the housing data provided by the university.

V.1.2 Housing Selection Outcomes

In the data we observe membership in "Greek Letter Organizations" (GLOs) and participation on an athletic team for every student in the housing selection process. Because GLOs play a large role in the daily lives of their members, they are likely to be the organizations most capable of enticing their members to coordinate.

Figure 1 in the introduction shows that members of certain GLOs were able to use better lottery numbers than non-members from the same year.[23] The groups used in this figure are a selection of sororities and fraternities that perform much better in terms of the average used lottery rank than both other GLOs and students who are not members of any GLO.

good information about these students, so we do not do any analysis of this decision, but this may be one reason to focus on the analysis comparing outcomes across organizations.

[23] About half of the students in our data are members of some GLO.

Table 1 shows the results of six OLS regressions, three each for the women and the men.[24] In all six we control for the composition of class years within each organization using the percentage of 2nd and 3rd year students in the organization. For the purposes of this control, we treat all students not in a GLO as one organization. These controls should capture any compositional differences that would allow some GLOs to do better than others, as we would expect the GLOs with more upperclassmen to fare better in the selection process. In (W1) and (M1) we include dummies for the number of semesters completed to capture the relationship between semesters and lottery times awarded, but do not report the coefficients in the table to save space.[25] In (W2) and (M2), we control for the lottery rank that was assigned to each individual, and in (W3) and (M3), we control for the pick order implied by the assigned lottery rank, allowing for ties. These are included as controls as we expect that students assigned better lottery ranks, and therefore earlier pick orders, should have better housing outcomes, regardless of GLO membership or strategic behavior. The reported standard errors for (W2), (W3), (M2) and (M3) are clustered at the suite level to account for the fact that the dependent variable is constant across the observations within a suite.

The lottery rank regressions, (W1) and (M1), provide evidence that the lottery ranks were distributed to individual students independent of their affiliation with an organization. Therefore, any improved outcomes are a function of differences in behavior. In (W2) and (M2), we look at the rank that was used to select the suite that each individual ended up in. The dependent variable in these regressions reflects the best time that an individual could select a suite given his or her group. The dependent variable for (W3) and (M3) is constructed from the actual time at which a suite was selected for an individual. Specifically, it measures the number

[24] Men and women do not compete for the same set of rooms, as suites are determined to be "male" or "female".

[25] There is an average difference of about 16–17 ranks between students with six semesters completed and those with four. Between four and two, there is a difference of approximately 11 ranks.

of students that picked ahead of an individual. An individual may be in a group with a low group rank, reflecting the fact that a member of the group is able to pick a suite relatively early, but a relatively high pick order if they do not actually choose housing when they are able to.[26]

The lottery rank differences indicate improvements in the best available login time of the students, but the number of students are not evenly distributed across login times, so the magnitude of these numbers can be difficult to interpret in terms of how much better off a group of students was as a result of their better login time. The Pick Order regressions, (W3) and (M3), show the average improvement for each of these organizations. There are several sororities and one fraternity whose group ranks are better on average than their fellow classmates outside of their organization, but four stand out in terms of their ability to secure suites well ahead of their competitors. For example, on average members of S2 picked about 9 time slots ahead of non-sorority women, which translates to each member being ahead of 127 students in the selection process. This is in a sample of 655 women.

The sororities and fraternities are arranged by listing the ones we are tentatively identifying as strategic (marked with an (s)) first and the non-strategic ones second. In identifying strategic GLOs, we looked for ones that had both significantly better lottery ranks and pick orders. S7 only improves in one category, for example. Organizations with fewer than five members are included in the regression as non-GLO students. S1, S2 and S3 represent about 31% of the sorority members involved in the process, while the fraction of fraternity members represented by F1 is about 12%. It is interesting to note that we do not see a strong correlation between the number of members in an organization and the relative improvement in outcomes.

These results show that members of several organizations are able to

[26]Pick Order may reflect delays or mistakes made by the students in selecting rooms. A student may have logged in at their assigned time but failed to pick a room right away, falling in the Pick Order. It is also possible that the students in a particular group were interested in selecting a suite that they thought would still be available later in the process.

	(W1) Lottery Rank	(W2) Group Rank	(W3) Pick Order
S1(s)	-0.500	-5.238**	-76.05**
	[4.296]	[2.292]	[25.50]
S2(s)	-0.618	-8.795***	-127.1**
	[2.661]	[2.754]	[40.72]
S3(s)	-1.172	-11.51**	-127.6*
	[10.82]	[5.167]	[56.76]
S4	-1.730	3.737	41.13
	[9.039]	[3.277]	[31.23]
S5	0.675	-2.921*	-31.92
	[2.597]	[1.718]	[34.85]
S6	-1.764	2.747	-1.371
	[1.900]	[2.123]	[38.39]
S7	0.712	-4.065***	51.19**
	[6.797]	[1.301]	[14.34]
S8	1.229	0.584	45.52
	[1.762]	[3.153]	[30.81]
S9	0.116	-0.487	10.46
	[1.871]	[1.933]	[44.48]
S10	2.976	5.365**	104.9**
	[4.709]	[2.681]	[28.78]
Pct. 2^{nd} Yr.	4.370	28.72***	421.3**
	[17.20]	[9.117]	[73.27]
Pct. 3^{rd} Yr.	13.50	-23.16**	-148.6*
	[37.89]	[11.01]	[78.89]
Lottery Rank		0.567***	
		[0.0469]	
Assigned Pick Order			0.634**
			[0.0585]
Constant	4	1.733***	66.03**
	[5.477]	[0.188]	[0.643]
N	655	655	655
R^2	0.887	0.500	0.368

	(M1) Lottery Rank	(M2) Group Rank	(M3) Pick Order
F1(s)	0.812	-8.508***	-100.4**
	[2.409]	[3.055]	[49.74]
F2	7.461	0.961	315.5
	[15.43]	[14.49]	[292.1]
F3	10.24	14.50	719.8
	[26.11]	[24.12]	[488.4]
F4	10.27	12.45	666.7
	[24.99]	[23.84]	[481.8]
F5	9.907	16.21	718.0
	[24.77]	[23.03]	[463.6]
F6	-4.306	-1.303	-191.6
	[12.62]	[11.33]	[230.3]
F7	-1.564	-3.902	54.47
	[4.392]	[4.204]	[85.12]
F8	5.774	10.16	485.0**
	[13.17]	[12.05]	[244.3]
F9	20.25	23.74	1,236
	[49.78]	[45.53]	[927.8]
F10	0.213	-0.299	-77.12
	[5.026]	[7.935]	[112.4]
F11	9.737	10.89	494.1
	[17.22]	[15.52]	[321.6]
Pct. 2^{nd} Yr.	4.493	2.711	-30.44
	[6.924]	[4.079]	[137.5]
Pct. 3^{rd} Yr.	-61.80	-92.78	-4,393
	[167.8]	[154.0]	[3,119]
Lottery Rank		0.597***	
		[0.0413]	
Assigned Pick Order			0.598***
			[0.0467]
Constant	22.76	25.39	1,261
	[41.47]	[38.49]	[770.5]
N	703	703	703
R^2	0.889	0.493	0.407

Standard errors in brackets
(W1) and (M1) control for semesters completed
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

(a) Women (b) Men

Table 1: Outcomes

	Strategic	GLO	Non-Strategic Non-GLO	All NS
Female	0.843	0.538***	0.414***	0.449***
2^{nd} Year	0.313	0.265	0.454***	0.401*
3^{rd} Year	0.304	0.504***	0.349	0.393*
4^{th} Year	0.383	0.231***	0.194***	0.204***
Pct. 2^{nd} Yr. in Org.	0.300	0.239***	0.322***	0.299
Pct. 3^{rd} Yr. in Org.	0.311	0.353***	0.252***	0.280***
Pct. 4^{th} Yr. in Org.	0.390	0.406	0.420***	0.416***
Asian	0.078	0.051	0.117	0.099
Black	0.017	0.037	0.091***	0.076**
Hispanic	0.017	0.037	0.042	0.041
Foreign	0.000	0.014	0.080***	0.062***
Athlete	0.00870	0.0142	0.145	0.108
$GPA - \overline{GPA}$	0.169	0.004***	-0.024**	-0.016***
N	115	351	897	1,248

$\overline{GPA}$ is the sample average GPA; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2: Descriptive Statistics

pick housing significantly earlier than the other students. Given the design of the mechanism this increases the number of choices they have and hence improves their housing choice set relative to other organizations. It should therefore yield more desirable choices of housing. Note that it is possible that some strategic coordination is occurring among other subsets of students, so that the identified organizations are just more sophisticated.[27]

Table 2 compares the means of the demographic variables we have in the data across four groups to examine how the strategic GLOs differ from non-strategic GLOs and non-GLOs. The stars reflect the results of t-tests when these means are compared to the ones in the strategic column. Rows 2–4 show the means of dummy variables for each class year, and hence report the class-year distribution for each column's group. The next three rows show the mean fraction of all members in that student's organiza-

[27]It is also important to note that the statistically significant differences we do find are not likely due to random chance. Within each model we would expect around one significant coefficient at the 5% level, but we find far more.

tion that are in each class year, regardless of whether the member participated in the mechanism (for non-GLO students their "organization" includes all non-GLO students).[28] Differences between rows 2–4 and rows 5–7 within the same column indicate a relationship between class year and the decision to participate in the mechanism. Comparing strategic GLOs to non-strategic GLOs, the table shows a larger fraction of the organizations' fourth year students participate in the mechanism, which is clearly important for outcomes because the fourth year students are guaranteed better priorities. It is interesting to note that the non-strategic GLOs have just as many fourth year students in their organization, so a lack of fourth year students in the organization available to participate in the mechanism does not explain the difference. Relative to the non-GLO students the difference is more stark with a lower fraction of fourth year students in the mechanism but a higher fraction overall.

These differences in distribution seem to be the result of strategic coordination by the organizations, but this evidence does not control for potential preference differences between the organizations. In the analysis that follows, we show that the results are consistent with strategic behavior on behalf of the organizations, after controlling for observable differences.

The above table also shows other demographic differences between the strategic groups and others. Perhaps the most interesting of these is the difference between GPA and the sample average GPA (grade point average measured on a four-point scale), which one might interpret as picking up more intellectual sophistication. The explanation is tempting, but our later results do not suggest a role for GPA or the other demographic characteristics as explanatory variables.

[28]For example, for each student in S1 we calculate the fraction of all members who are second-, third- and fourth-year students in S1, which is obviously the same for all members of S1. The means in Table 2 reflect the result of averaging over these fractions for each student in the mechanism.

V.1.3 Are the organizations behaving strategically?

If only a few organizations can effectively coordinate their members to significantly improve their housing allocation, then ex ante the mechanism allocates rooms unequally across organizations. Thus, through alternative designs which remove or at least reduce these incentives, the university may be able to make outcomes more equitable. The design goal here has been referred to as "leveling the playing field."

To a large extent, this argument relies on the assumption that the organizations are actually behaving strategically, so that they are sensitive to changes in the incentives to strategically coordinate their members. However, if, for example, the room allocations were made independently of the decisions over whom to live with, there would be no ex ante difference in room allocations regardless of whether the differences in outcomes were the result of strategic behavior or some other cause (e.g., a difference in preferences).

In this section, we present tests of the hypothesis that some organizations are strategically coordinating their members to improve their housing allocation. Our main concern in these tests is controlling for the possibility that the outcomes we observe are driven by preference differences between organizations, instead of strategic differences.

To understand the motivation for the tests recall that the best lottery rank within each group of students who live together completely determines their ability to select housing, and hence should be the primary strategic concern. If an organization is able to effectively coordinate students it must do a better job of distributing the best lottery ranks throughout the groups than the individual members can. A good lottery rank that is not the best in the group is effectively wasted if it could have been used to improve the selection of another group in the organization.

The tests all roughly measure the probability that the students within a group have lottery ranks that are close to that of the best lottery rank in the group, and therefore are similar. If students are forming groups due to roommate preferences before picking suites, and any differences in outcomes that we see are due mainly to differences in suite preferences, we should expect that there should be little strategic behavior over the composition of a group and its members. However, if acting strategically, the effective organizations should be able to redistribute students with lottery ranks close to the best in a group to increase their probability of getting more favorable suites. Specifically, the tests address the following questions: controlling for observables, is it more likely that a member of an organization labeled strategic (1) Lives with a student in an older class year? (2–3) Does not have an assigned lottery rank close to the member who selected housing when that member is in the same organization? (4–5) Was assigned a pick order much lower (i.e., worse) than the one used to select housing?

Tests 1–3 focus on the organization of each group and do not directly measure outcomes, with test 1 being the roughest measure. One can imagine the process of selecting housing as occurring in two stages. In the first stage, groups are arranged by the organizations, and in the second these groups select their housing. In the current procedure it is possible that groups were chosen to effectively utilize lottery ranks within the organization, but for some reason the groups did a poor job of selecting housing at the first available times (perhaps simply by mistake). If an organization has a large influence on the first stage but a smaller one on the second, then the outcomes could be poor despite the organization displaying strategic behavior. The measures in 1–3 are designed to pick up on the organization of the groups, ignoring the eventual outcomes.

Tests 4 and 5 are similar to 3 and 4 but incorporate information about the actual outcomes, because the measure is based on the actual order in which the housing was chosen.

Table 3 reports the the marginal effects at the mean observation using a logit specification for each test with the variable of interest, labeled "Strategic", being an indicator variable for whether the student was a member of one of the 4 organizations tentatively identified as strategic in the last section. In each case, we control for observable characteristics of the students shown in Table 2.

In regression (1), we exclude the 4th year students. This regression shows that students in strategic organizations are 32 percentage points more likely to live with a student from an older class year compared to students who are not members of any GLO. There is also not a significant difference in this measure between the non-strategic GLO members and the comparison group (this is the Non-GLO coefficient). The measure in regression (1) is rough, because it only picks up on class year differences, but class years are very important for the purposes of picking housing early, since giving a group its first student from an older class must improve its pick order. It is possible that only members of the organizations we have defined as "strategic" prefer to have mixed class years in their suites, which happens to result in better outcomes, but this seems unlikely.

Regressions (2) and (3) exclude the students who selected housing for their group and ask of the remaining students in the group: What is the probability that their rank is within a certain number of slots of the selector? A lower probability here indicates a better organized group, since it means that it is less likely that a good lottery rank was unused. We ran this regression using cutoffs from 1 to 8 (the remainder are in the appendix), where each lottery rank may be shared by between 20 and 100 students. In regression (2) and for cutoff values less than 5, the Strategic coefficient is insignificant (at the 10% level), but to some extent this is expected as it becomes less likely that any group of students has lottery ranks close together for these cutoff values. The coefficient on Strategic is significant and negative in (3) and for the regression with a cutoff of 7, but becomes insignificant again at a cutoff of 8 (at the 10% level). For reference, the login times spanned five days with four of the five having 11 different login times and the fifth having 12 times.

As with regression (1), regressions (2) and (3) are consistent with the hypothesis that the strategic organizations are more effective at arranging their member's groups. What's interesting about (2) and (3), relative to (1),

	(1) Live with Upperclass.	(2) Rank Within 5	(3) Rank Within 6	(4) Actual Pick > 50	(5) Actual Pick > 100
Strategic	0.322***	-0.0836	-0.158**	0.376***	0.330***
	[0.0642]	[0.0723]	[0.0677]	[0.0652]	[0.0683]
Non-GLO	-0.0601	0.102**	0.112**	0.0604*	0.0257
	[0.0398]	[0.0418]	[0.0471]	[0.0355]	[0.0299]
Female	0.0217	0.0151	0.0185	-0.0695***	-0.0419**
	[0.0332]	[0.0295]	[0.0331]	[0.0235]	[0.0196]
GPA	-0.0263	-0.0274	-0.0468	-0.00659	0.00888
	[0.0343]	[0.0314]	[0.0347]	[0.0244]	[0.0204]
Asian	-0.0147	0.0322	-0.0417	0.0235	0.0225
	[0.0640]	[0.0628]	[0.0614]	[0.0523]	[0.0465]
Black	-0.0143	0.0451	0.0124	-0.0305	-0.0232
	[0.0726]	[0.0580]	[0.0615]	[0.0417]	[0.0347]
Hispanic	0.183**	-0.0133	-0.0315	-0.0147	0.00187
	[0.0854]	[0.0697]	[0.0782]	[0.0585]	[0.0514]
Foreign	0.0880	-0.00172	0.0190	-0.0314	-0.0137
	[0.0706]	[0.0632]	[0.0724]	[0.0462]	[0.0393]
Athlete	0.0898*	0.00755	0.0228	0.0280	-0.00808
	[0.0527]	[0.0499]	[0.0550]	[0.0387]	[0.0300]
2 Sem.	0.120***	-0.288***	-0.389***	0.463***	0.502***
	[0.0335]	[0.0382]	[0.0426]	[0.0466]	[0.0536]
3 Sem.	0.524***	-0.180***	-0.249***	0.670***	0.785***
	[0.0852]	[0.0308]	[0.0254]	[0.0771]	[0.0505]
4 Sem.		-0.240***	-0.337***	0.374***	0.392***
		[0.0337]	[0.0386]	[0.0501]	[0.0587]
Pct. 3^{rd} Yr. in Org.		0.663**	0.533	0.710***	0.483**
		[0.320]	[0.363]	[0.244]	[0.198]
Pct. 4^{th} Yr. in Org.		0.0717	0.142	0.316	0.191
		[0.236]	[0.269]	[0.195]	[0.165]
Observations	1,018	851	851	1,363	1,363

Standard errors are in brackets. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 3: Marginal Effects on Strategic Indicators

is that the Non-GLO coefficient is significant and positive, hinting at some strategic behavior in the other GLO organizations that does not translate into better outcomes. Perhaps these organizations are behaving strategically to some extent, but the members are reversing the benefits by not actually choosing housing at the best possible times (we do see some evidence in the data of a small number of groups that chose housing much later than they could have). It is also possible that non-strategic GLOs are less willing to place students together with as large gaps in class year as strategic GLOs are.

Finally, regressions (4) and (5) are the most similar to the ones in Table 1, because they incorporate outcome data. The dependent variables in these regressions measure the difference between the rank order in which the group actually selected housing and the implied pick order based on their assigned lottery rank (allowing for ties). Again we see a large and significant coefficient on Strategic, but compared to (2) and (3) the coefficient on Non-GLO becomes insignificant at the 5% level, suggesting that there was a problem with how those "non-strategic" organizations chose housing after they organized into groups.

Only a few control variables in these regressions are statistically significant, and none are consistently significant across specifications. Table 2 indicates a possible role for GPA, but these regressions do not indicate any relationship between GPA and strategic behavior. The insignificance of the demographic information we have suggests that the demographic composition of the organizations is not driving and/or explaining this behavior. Finally, we reiterate that we are only considering one partition of students into organizations. The weakly significant coefficient on Athlete in Table 4 indicates that athletes grouped by sport might be another category within which students behave strategically. We have no reason to rule out possibilities like these, but cannot explore them further for lack of data.

Together we view these results as being consistent with the hypothesis that the strategic organizations are effectively organizing members into roommate groups in order to improve housing outcomes.

VI Conclusion

In a model of student housing allocation mechanisms in which students have preferences over roommates and rooms, we can only guarantee that a stable allocation exists if priority for housing is determined after a match forms, as happens in ϕ_R. We also show that in general there is no efficient mechanism, which means there is no mechanism which always results in a weakly more efficient allocation than all other mechanisms. However, if roommate values are independently and identically distributed according to a log-concave distribution, then ϕ_R is more efficient ex ante than any mechanism which determines housing priority through a varying function of the roommate match. Lastly, due to the trade-off between roommate and room values, we show that certain mechanisms allow, or even encourage, student organizations to act strategically.

In outcome data, we identify student organizations that select rooms significantly ahead of all other students in the mechanism. Our tests of strategic behavior suggest that this is the result of these organizations keeping more older students on campus in order to group them with younger students who (by the rules of the mechanism) must receive worse lottery numbers. Our test shows that underclassmen in strategic organizations are approximately 40% more likely to live with an older member of the organization after controlling for observable characteristics and the distribution of students across class year. The result of this strategic behavior is that the average member of a strategic organization chooses suites ahead of a full class year's number of students. In other words, the behavior has significant implications for the ex post allocation of rooms.

References

Abdulkadiroglu, A. and Sönmez, T. (1999). House allocation with existing tenants. *Journal of Economic Theory*, 88(2):233 – 260.

Arkin, E. M., Bae, S. W., Efrat, A., Okamoto, K., Mitchell, J. S., and Polishchuk, V. (2009). Geometric stable roommates. *Information Processing Letters*, 109(4):219 – 224.

Bartholdi, J. and Trick, M. A. (1986). Stable matching with preferences derived from a psychological model. *Operations Research Letters*, 5(4):165 – 169.

Bergstrom, T. and Bagnoli, M. (2005). Log-concave probability and its applications. *Economic theory*, 26:445–469.

Chung, K.-S. (2000). On the existence of stable roommate matchings. *Games and Economic Behavior*, 33(2):206 – 230.

Dogan, O., Laffond, G., and Laine, J. (2011). The core of shapley-scarf markets with couples. *Journal of Mathematical Economics*, 47(1):60 – 67.

Erdil, A. and Ergin, H. (2008). What's the matter with tie-breaking? improving efficiency in school choice. *American Economic Review*, 98(3):669 – 689.

Ergin, H. and Sonmez, T. (2006). Games of school choice under the boston mechanism. *Journal of Public Economics*, 90(1-2):215 – 237.

Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):pp. 9–15.

Graham, D. A. and Marshall, R. C. (1987). Collusive bidder behavior at single-object second-price and english auctions. *Journal of Political Economy*, 95(6):1217 – 1239.

Hylland, A. and Zeckhauser, R. (1979). The efficient allocation of individuals to positions. 87(2):293 – 314.

McAfee, R. P. and McMillan, J. (1992). Bidding rings. *American Economic Review*, 82(3):579 – 599.

Pesendorfer, M. (2000). A study of collusion in first-price auctions. *Review of Economic Studies*, 67(3):381 – 411.

Pycia, M. (2012). Stability and preference alignment in matching and coalition formation. *Econometrica*, 80(1):pp. 323–362.

Shapley, L. and Scarf, H. (1974). On cores and indivisibility. *Journal of Mathematical Economics*, 1(1):23 – 37.

Appendix

Proof of Theorem 1. Construct match μ_n as follows: let $\mu_0 = \emptyset$ and $\mu_i = \mu_{i-1} \cup \{(s,s')_i\}$, where $(s,s')_i = \arg\max_{\{s_j,s_k\} \subset S \setminus \mu_{i-1}} w(s_j,s_k)$. We claim that μ_n is stable with respect to ϕ_R. Under ϕ_R the expected room value of each student is independent of the pair structure, thus any blocking pair (s,t) must be such that $w(s,t) > \max\{w(s,s'), w(t,t')\}$ where $s' = \mu_n(s)$ and $t' = \mu_n(t)$. Clearly this is a contradiction. □

Proof of Theorem 2. We will prove this for the case when all students have common values for all rooms, which is a special case of our model. Keeping ϕ fixed, let $V_s(\phi,\mu) = V_{\mu(s)}(\phi,\mu)$ be the common expected room value of pair $(s,\mu(s))$ in match μ under mechanism ϕ. In the rest of the proof, to simplify notation, we may denote V_{s_i} as V_i.

Lemma 1. *If $\phi \neq \phi_R$, then there exist four matches μ_1, μ_2, μ_3, and μ_4 such that $\mu_1 \leftrightarrow \mu_2 \leftrightarrow \mu_3 \leftrightarrow \mu_4 \leftrightarrow \mu_1$, $|\mu_1 \setminus \mu_3| = 3$, and $V_1(\phi,\mu_2) - V_1(\phi,\mu_1) \neq V_1(\phi,\mu_3) - V_1(\phi,\mu_4)$ with $(s_1,s_2) \in \mu_1 \cap \mu_4$, and $(s_1,s_3) \in \mu_2 \cap \mu_3$*

Proof. Suppose for contradiction that the statement is not true. Take any match μ_1 and pick any three pairs $\{(s_1,s_2),(s_3,s_4),(s_5,s_6)\} \subseteq \mu_1$. Holding all other pairs fixed restrict attention to the subset $\mathcal{M}' \subseteq \mathcal{M}$ of fifteen matches which differ only in how these six students are paired. Let μ_2 be the match adjacent to μ_1 such that $\{(s_1,s_3),(s_2,s_4),(s_5,s_6)\} \subseteq \mu_2$. Let μ_3 and μ_3' be the matches adjacent to μ_2 such that $\{(s_1,s_3),(s_2,s_5),(s_4,s_6)\} \subseteq \mu_3$, $\{(s_1,s_3),(s_2,s_6),(s_4,s_5)\} \subseteq \mu_3'$. Lastly, let μ_4 and μ_4' be such that $\mu_4 \leftrightarrow \mu_3$, $\mu_4' \leftrightarrow \mu_3'$, and $\{(s_1,s_2),(s_3,s_5),(s_4,s_6)\} \subseteq \mu_4$, $\{(s_1,s_2),(s_3,s_6),(s_4,s_5)\} \subseteq \mu_4'$.

Thus $\mu_1 \leftrightarrow \mu_2 \leftrightarrow \mu_3 \leftrightarrow \mu_4 \leftrightarrow \mu_1$ and $\mu_1 \leftrightarrow \mu_2 \leftrightarrow \mu_3' \leftrightarrow \mu_4' \leftrightarrow \mu_1$, which by assumption implies $V_1(\phi,\mu_2) - V_1(\phi,\mu_1) = V_1(\phi,\mu_3) - V_1(\phi,\mu_4) = V_1(\phi,\mu_3') - V_1(\phi,\mu_4')$. Therefore $V_i(\phi,\mu) - V_i(\phi,\mu') = V_i(\phi,\mu'') - V_i(\phi,\mu''')$ for any $\mu \leftrightarrow \mu'$ and $\mu'' \leftrightarrow \mu'''$ such that $\{\mu,\mu',\mu'',\mu'''\} \subset \mathcal{M}'$.

Next we will show that this implies $V_i(\phi,\mu) = V_i(\phi,\mu')$ for any $i = 1\ldots,6$ and $\mu,\mu' \in \mathcal{M}'$. Take any two adjacent matches in $\mathcal{M}'$, and without loss of generality assume they are μ_1 and μ_4. Let μ_5 and μ_5' be the matches in $\mathcal{M}'$ such that $\{(s_1,s_4),(s_2,s_3),(s_5,s_6)\} \subseteq \mu_5$ and $\{(s_1,s_4),(s_2,s_5),(s_3,s_6)\} \subseteq \mu_5'$. Then we have $\mu_1 \leftrightarrow \mu_2 \leftrightarrow \mu_5 \leftrightarrow \mu_1$ and $\mu_4 \leftrightarrow \mu_3 \leftrightarrow \mu_5' \leftrightarrow \mu_4'$. This implies that $V_1(\phi,\mu_4) = V_1(\phi,\mu_4')$. Since this was true for arbitrary μ_1 this implies that $V_{\mu(i)}(\phi,\mu) = V_i(\phi,\mu) = V_i(\phi,\mu') = V_{\mu(i)}(\phi,\mu')$ for all $\mu,\mu' \in \mathcal{M}$ such that $(s_i,\mu(s_i)) \in \mu \cap \mu'$. Therefore there are only fifteen independent expected room values for $\mu \in \mathcal{M}'$, one for each pair. Let the variables $\nu_{i,j}$ with $i,j \in \{1,\ldots,6\}$ represent these values. The system of equations described by these fifteen unknown expected room values and the fifteen equations $\sum_{(s_i,s_j)\in\mu\in\mathcal{M}'} \nu_{i,j} = x$ has a unique solution. It is easy to see that the unique solution is that $\nu_{i,j} = \nu_{k,l}$ for all pairs (s_i,s_j) and (s_k,s_l), which contradicts that $\phi \neq \phi_R$. □

Therefore let μ_1, μ_2, μ_3, and μ_4 satisfy the conditions of the lemma. Let $s_2 = \mu_1(s_1)$ and $s_3 = \mu_2(s_1)$. Without loss of generality we may assume that $V_1(\phi,\mu_2) > V_1(\phi,\mu_1)$ and $V_1(\phi,\mu_2) - V_1(\phi,\mu_1) > V_1(\phi,\mu_3) - V_1(\phi,\mu_4)$. Whenever

$$\begin{aligned} & V_1(\phi,\mu_2) - V_1(\phi,\mu_1) \\ &\quad > w(s_1,s_2) - w(s_1,s_3) \\ &\quad > V_1(\phi,\mu_3) - V_1(\phi,\mu_4) \end{aligned} \tag{1}$$

we have $U_1(\phi,\mu_2) > U_1(\phi,\mu_1)$ but $U_1(\phi,\mu_3) < U_1(\phi,\mu_4)$. Find some M such that

$$M > \max_{i,\mu\leftrightarrow\mu'} |V_i(\phi,\mu) - V_i(\phi,\mu')|,$$

and let $w(s_1,s_2) = 3M$ and $w(s_1,s_3) = 3M - \delta$ to satisfy (1). Let $w(s_3,s_4) = w(s_2,s_5) = 2M$. For all pairs $(s_i,s_j) \in \mu_1 \cap \mu_3$ let $w(s_i,s_j) = M$, and for all other pairs let $w(s_i,s_j) = 0$. We claim that these preferences are such that there does not exist a stable match under ϕ. First note that any match μ in which there is some $s \in S$ such that $\mu(s) \notin \{\mu_1(s), \mu_3(s)\}$ is not stable. Next, since μ_1 and μ_3 differ by exactly three pairs, they are the only two matches in which every $s \in S$ is such that $\mu(s) \in \{\mu_1(s), \mu_3(s)\}$. Finally, note that μ_1 is blocked by μ_2, and μ_3 is blocked by μ_4.

□

Proof of Theorem 3. The proof will be by a counterexample based on Example 2. The example shows that we can produce a market in which s_1 and s_2 form a pair under ϕ_L, even when it is not efficient, and we can produce a market in which s_2 and s_3 form a pair under ϕ_R, even when it is not efficient. Therefore, in order for any mechanism ϕ to be more efficient than all other mechanisms, it would have to be that it always results in the efficient match in this simple market. Therefore when $\varepsilon > 0$, we require that $\max\{V_2(\phi,\mu) + 3 + \varepsilon, V_1(\phi,\mu) + 1\} \quad > \quad \max\{V_1(\phi,\mu') + 2, V_3(\phi,\mu') + 2\}$, and when $\varepsilon < 0$ this inequality is reversed. This implies that for $\varepsilon = 0$ that $\max\{V_2(\phi,\mu) + 3, V_1(\phi,\mu) + 1\} = \max\{V_1(\phi,\mu') + 2, V_3(\phi,\mu') + 2\}$. This is clearly not possible in general since it must be true for all possible values of rooms. □

Proof of Theorem 4. Let the roommate value for the pair (s_i,s_j) be determined by the random variable $X_{ij} \sim F$. Let x be a realization of roommate values corresponding to a stable pair structure under ϕ and let y the realized roommate values of an adjacent pair structure that blocks x under ϕ_R. In events in which no adjacent pair structure blocks x under ϕ then x is also the stable pair structure under ϕ_R, and since we are interested only in the expected difference between the two mechanisms, we can ignore these events.

Because they are adjacent, all but two elements each of x and y are identical. Let the four elements that differ with positive probability be

y_1, y_2, x_1, and x_2, representing, without loss of generality, the roommate values of the pairs (s_1,s_2), (s_3,s_4), (s_1,s_3), and (s_2,s_4) respectively. Let the value of the rooms received by (s_1,s_2) and (s_3,s_4) be denoted by v_{x_1} and v_{x_2}. For (s_1,s_3) and (s_2,s_4), we need the values of the rooms received under ϕ if these pairs form, and we denote those by v_{y_1} and v_{y_2} respectively. At this point we can only rule out that v_{x_1} and v_{x_2} (or v_{y_1} and v_{y_2}) refer to the same room.

The difference in surplus between x and y is $y_1 + y_2 - x_1 - x_2$. Inequalities that characterize the event that y blocks x under ϕ_R but x blocks under ϕ are

$$
\begin{aligned}
&x_1 \leq y_1 \leq x_1 + (v_{x_1} - v_{y_1}) \\
&x_2 \leq \min\{y_1, x_1 + (v_{x_1} - v_{x_2})\} \\
&y_2 \leq \min\{y_1, x_1 + (v_{x_1} - v_{y_2})\}.
\end{aligned}
$$

These follow from observing that it must be that $\max\{y_1,y_2\} \geq \max\{x_1,x_2\}$ and $\max\{y_1 + v_{y_1}, y_2 + v_{y_2}\} \leq \max\{x_1 + v_{x_1}, x_2 + v_{x_2}\}$, and assuming without loss of generality that $x_1 + v_{x_1} \geq x_2 + v_{x_2}$ and $y_1 \geq y_2$. Note that it must be that $v_{x_1} > v_{y_1}$ for this event to have non-zero probability. Consider conditioning on some realization of room values, x_1, and y_1 that does not violate the above inequalities. The conditional expected difference in surplus is

$$
y_1 + \mathrm{E}[Y_2 | Y_2 \leq \min\{y_1, x_1 + (v_{x_1} - v_{y_2})\}] - x_1 - \mathrm{E}[X_2 | X_2 \leq \min\{y_1, x_1 + (v_{x_1} - v_{x_2})\}],
$$

which may be positive or negative depending on the values of v_{y_2} and v_{x_2}. However conditional on x_1 being assigned the room corresponding to v_{x_1} there is by the assumption of anonymity of the mechanism an equal probability that x_2 will be assigned any of the remaining rooms, and similarly for y_1 and y_2 respectively. Conditioning on v_{x_1} and v_{y_1}, and taking expectation over the remaining possibilities yields

$$y_1 - x_1 + \frac{1}{n}\left(\mathrm{E}[Y_2|Y_2 \leq \min\{y_1, x_1\}] - \mathrm{E}[X_2|X_2 \leq \min\{y_1, x_1 + (v_{x_1} - v_{y_1})\}]\right)$$
$$= y_1 - x_1 + \frac{1}{n}\left(\mathrm{E}[Y_2|Y_2 \leq x_1] - \mathrm{E}[X_2|X_2 \leq y_1]\right) > 0,$$

where the inequality follows from the assumption of log-concavity and that $y_1 \geq x_1$, which guarantees that $y - \mathrm{E}[Y|Y < y]$ is a non-decreasing function of y. Next we argue that any such y has lower expected surplus than the unique stable pair structure under ϕ_R. From Theorem 1, the stable pair structure under ϕ_R can be constructed by successively finding the pair with the highest roommate value among all remaining feasible pairs. Let the ϕ_R stable pair structure's roommate values be given by z, which may refer to the same pair structure as x or z. Comparing some z to y, sort the vectors in descending order so that z_1 is the largest element of z for example. If z and y do not refer to the same pair structure, with positive probability there is some i such that either $i = 1$ and $z_1 > y_1$, or $i > 1$ and $\forall j < i, z_j = y_j$ and $z_i > y_i$. That is, they differ at some point in the process of selecting the best pair from the remaining feasible pairs. The pair corresponding to z_i blocks y under ϕ_R and hence there is a pair structure y' incorporating this blocking pair and the pair that results from this block, labeled y'_j. The difference in surplus between y' and y is $z_i - y_i + y'_j - y_j$. Conditioning on z_i and y_i, the difference in expected surplus is positive, since

$$z_i - y_i + \mathrm{E}[Y'_j|Y'_j \leq z_i] - \mathrm{E}[Y_j|Y_j \leq y_i] \geq 0,$$

because $\mathrm{E}[X|X \leq \cdot]$ is an increasing function. If $y' \neq z$, then one can form a y'' adjacent to y' in the same manner. This process can be continued until z is reached.

□

Proof of Theorem 5. The probability that Ω receives the first k rooms under ϕ_L is the probability that students in Ω receive the best k priorities out of a total of $2n$ possible priorities. The probability that Ω receives the best k rooms under ϕ_R is the probability that pairs in Ω receive the best k priorities out of n total priorities. Therefore the relevant probabilities are:

$$\text{Probability } \Omega \text{ receives best } k \text{ rooms under } \phi_R = \prod_{i=0}^{k-1} \frac{m-i}{n-i}$$

$$\text{Probability } \Omega \text{ receives best } k \text{ rooms under } \phi_L = \prod_{i=0}^{k-1} \frac{2m-i}{2n-i}$$

It is easy to show that the second probability is larger than the first as long as $n > m$. □

Proof of Theorem 6. Label the rooms $r_1, \ldots, r_n$ such that $v(r_i) \geq v(r_j)$ for all $i < j$. As in the definition of ϕ_L, label students $s_1, \ldots, s_{2n}$ such that $(s_i, s_j) \succ (s_k, s_l) \iff \min[i, j] < \min[k, l]$.

Lemma 2. *Under ϕ_L, if both Ω and $S \setminus \Omega$ are societies and $r_n \notin R(\Omega)$, then $r_i \in R(\Omega)$ implies that $s_i \in \Omega$.*

Proof. Suppose that for some $r_i \in R(\Omega)$ we have that $s_i \notin \Omega$. First assume that $r_i = r_1$. Since s_1 always chooses r_1 under ϕ_L this is not possible. Next assume $r_i = r_2$. This implies that $\{s_1, s_2\} \subset \Omega \setminus S$ is the pair occupying r_1, and $s_3 \in \Omega$ occupies r_2. Let $\{t, t'\} \subset \Omega \setminus S$ be the pair of students occupying r_n, and let $R(S \setminus \Omega) = \{r_{a_1}, \ldots, r_{a_{n-m}}\}$, where $a_i \in \{1, \ldots, n\}$, and $v(a_i) \geq v(a_j)$ for all $i < j$. Holding all else equal, if $S \setminus \Omega$ instead reports pairs (s_1, t) and (s_2, t'), the set of rooms it receives under ϕ_L would be $R' = \{r_1, r_2, a_{3+1}, \ldots, r_{a_{n-m-1}+1}\}$. The change in room-welfare between $R(S \setminus \Omega)$ and R' is $v(r_2) - v(r_{a_3}) + v(r_{a_{3+1}}) - v(r_{a_4}) + \ldots + v(r_{a_{n-m-1}+1}) - v(r_n) > 0$. Therefore $r_i \neq r_2$. For the inductive step, assume that r_i cannot equal r_j for any $2 \leq j < i$. The proof for this step is analogous to the $i = 2$ case. □

Let r_j be the worst room in $R(\Omega)$, and let R' be any set of rooms such that $|R'| = m$ and $r_j = \arg\min_{r \in R'} v(r)$. From Lemma 2, the probability that Ω occupies the set $R(\Omega) = R'$ under ϕ_L is equal to the probability that the

mth highest priority held by the $2m$ students in Ω is the jth highest priority held by the $2n$ students in S. This probability is given by the following expression:

$$\Pr(R(\Omega) = R'|\phi = \phi_L) = \frac{\prod_{i=0}^{j-m-1}(2(n-m)-i)\prod_{i=0}^{m-1}(2m+i)}{\prod_{i=0}^{j-1}(2n-i)}$$

Therefore, if we let R'' be any set of rooms such that $|R''| = m$ and $r_{j+1} = \arg\min_{r\in R''} v(r)$, $j+1 < n$, then the probability that $R(\Omega) = R''$ is given by:

$$\begin{aligned} Pr(R(\Omega) = R''|\phi = \phi_L) &= \frac{\prod_{i=0}^{j-m}(2(n-m)-i)\prod_{i=0}^{m-1}(2m+i)}{\prod_{i=0}^{j}(2n-i)} \\ &= \Pr(R(\Omega) = R'|\phi = \phi_L)\underbrace{\left(\frac{2n-m-j}{2n-j}\right)}_{<1} \end{aligned}$$

□

Proof of Theorem 7. First we show that the welfare for Ω when competing against P must be weakly higher than when competing against Q. For any society Ω', let $S^*(\Omega')$ be the highest priority $|\Omega'|/2$ students in Ω'. It is always a (weakly) dominant strategy for each society to pair students in $S^*(\Omega')$ with students not in this set. For partition Q, let $S^*(Q) = \{s_{q_1}, \ldots, s_{q_n}\}$ be the ordered set of students such that $q_i < q_j$ for all $i < j$, and $s_{q_i} \in S^*(\Omega')$ for some $\Omega' \in Q$. Similarly define $S^*(P) = \{s_{p_1}, \ldots, s_{p_n}\}$. It is immediate that in partition Q, if $s_{q_i} \in \Omega'$, then $r_i \in R(\Omega')$, and likewise for $s_{p_i} \in \Omega'$ for some $\Omega' \in P$. Let $S^*(\Omega) = \{\omega_1, \ldots, \omega_m\} = \{s_{q_i}, \ldots\} = \{s_{p_j} \ldots\}$. For any set of societies $\{\Omega_1, \ldots, \Omega_k\}$ such that each $\Omega_i \in P$ and $\cup_{i=1}^k \Omega_i = \Omega' \in Q$, every $s_i \in S^*(\Omega')$ is such that $i \leq \max\{j | s_j \in \cup_{l=1}^k S^*(\Omega_l)$. Therefore we have that, for any $\omega_i = s_{q_j} = s_{p_k}$, $|\{s_{p_l} \in S^*(P) | p_l < p_k\}| \leq |\{s_{q_l} \in S^*(P) | q_l < q_j\}|$.

All that is left is to show that there exists some individual student priority order so that Ω has higher welfare when competing against Q rather than P. Let Ω' be any society in Q such that $\Omega' \notin P$. Let Ω'' be the smallest subset of Ω' in P. Let the individual student priority order be such that Ω'' contains the highest $|\Omega''|$ students, and Ω contains the next m s highest priority students. Then under P, $R(\Omega) = \{r_{|\Omega''|+1}, \ldots, r_{|\Omega''|+m}\}$, while under Q $R(\Omega) = \{r_{|\Omega''|/2+1}, \ldots, r_{|\Omega''|/2+m}\}$. $\square$

	(6) Rank Within 3	(7) Rank Within 4	(8) Rank Within 7	(9) Rank Within 8	(10) Actual Pick > 25	(11) Actual Pick > 125	(12) Actual Pick > 150
Strategic	0.00463	0.0285	-0.0132	-0.0202	-0.0976***	-0.0318**	-0.0236*
	[0.0215]	[0.0259]	[0.0352]	[0.0375]	[0.0256]	[0.0154]	[0.0134]
Non-GLO	-0.0182	-0.0636	-0.212***	-0.155	0.404***	0.339***	0.270***
	[0.0640]	[0.0653]	[0.0680]	[0.0967]	[0.0600]	[0.0732]	[0.0724]
Female	0.0616**	0.0935***	0.144***	0.130**	0.0664*	0.0265	0.0198
	[0.0300]	[0.0363]	[0.0501]	[0.0543]	[0.0390]	[0.0230]	[0.0201]
GPA	-0.0420*	-0.0461*	-0.0224	-0.0131	0.00609	0.00852	0.00710
	[0.0223]	[0.0268]	[0.0372]	[0.0394]	[0.0272]	[0.0155]	[0.0133]
Asian	-0.00329	0.00860	-0.0395	-0.0400	0.0421	0.0347	0.0278
	[0.0393]	[0.0501]	[0.0686]	[0.0763]	[0.0573]	[0.0396]	[0.0344]
Black	0.0251	0.0397	-0.0139	-0.0130	-0.0447	-0.00142	-0.00311
	[0.0418]	[0.0507]	[0.0644]	[0.0701]	[0.0450]	[0.0286]	[0.0242]
Hispanic	-0.0485	-0.0484	-0.0611	-0.131	-0.0312	0.00739	0.00281
	[0.0369]	[0.0503]	[0.0834]	[0.0843]	[0.0627]	[0.0409]	[0.0350]
Foreign	0.00878	0.0181	-0.0289	-0.0521	-0.0357	-0.0160	-0.00739
	[0.0496]	[0.0578]	[0.0728]	[0.0762]	[0.0520]	[0.0281]	[0.0250]
Athlete	-0.0364	0.00510	0.0276	0.00833	0.0312	-0.00386	-0.00650
	[0.0307]	[0.0429]	[0.0573]	[0.0593]	[0.0424]	[0.0228]	[0.0191]
2 Sem.	-0.170***	-0.236***	-0.466***	-0.552***	0.414***	0.629***	0.659***
	[0.0293]	[0.0340]	[0.0462]	[0.0494]	[0.0435]	[0.0762]	[0.0967]
3 Sem.	-0.0766**	-0.131***	-0.301***	-0.362***	0.583***	0.843***	0.880***
	[0.0342]	[0.0271]	[0.0240]	[0.0233]	[0.0947]	[0.0569]	[0.0522]
4 Sem.	-0.115***	-0.186***	-0.402***	-0.494***	0.323***	0.530***	0.557***
	[0.0241]	[0.0291]	[0.0433]	[0.0481]	[0.0462]	[0.0884]	[0.115]
Pct. 3^{rd} Yr. in Org.	0.326	0.339	0.365	0.199	0.759***	0.399**	0.202
	[0.231]	[0.300]	[0.394]	[0.413]	[0.269]	[0.157]	[0.138]
Pct. 4^{th} Yr. in Org.	0.113	0.0287	-0.0292	0.405	0.299	0.132	-0.00492
	[0.165]	[0.216]	[0.297]	[0.324]	[0.209]	[0.131]	[0.121]
Observations	850	850	850	850	1,360	1,360	1,360

Standard errors are in brackets. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 4: Marginal Effects on Strategic Indicators