

Sorority Rush as a Two-Sided Matching Mechanism

By SUSAN MONGELL AND ALVIN E. ROTH*

The history and organization of the membership recruitment process of American sororities is studied. Like entry-level labor markets studied previously, this process experienced failures that led to the adoption of a centralized matching procedure in which a matching is determined on the basis of preference lists submitted by the agents. Analysis of the rules of the match and of preference lists from 21 matches reveals an unstable matching procedure that gives agents incentives to behave strategically. The analysis also shows how the agents act on these incentives and how the resulting strategic behavior has contributed to the longevity of the matching system and to the stability of the resulting matches. (JEL C78, D00, J41)

This paper concerns the formal process by which women at American universities join the social organizations called sororities. The history of this process, of the problems it has encountered, and how it has evolved to meet them, have striking similarities to (as well as important differences from) the history and organization of the American labor market for medical interns (see Roth, 1984a) and the several similar entry-level labor markets for physicians in the United Kingdom (see Roth, 1991). Thus, by studying this process, we hope to learn more about other matching processes and hope to assess the generality of various hypotheses about them.

In the medical labor markets, competition for newly graduating medical students and for desirable positions caused the dates at which appointments were finalized to unravel in time, so that by the 1940's in the United States and by the 1960's in the United Kingdom, postgraduation employ-

ment was often arranged well over a year (and sometimes over two years) in advance of graduation. Similarly, by the latter part of the last century, entry into fraternities and sororities, initially reserved for college seniors, had worked its way backward to the freshman class, and in some cases membership was arranged well before matriculation. (This aspect of the competition for members appears to be the origin of the term "rushing," as these membership drives are now called.)

Largely in response to the problems arising out of this kind of unraveling, the parties involved in the different medical labor markets eventually agreed to try a variety of centralized matching procedures, in which participants would not sort themselves out individually but would instead submit rank-orderings of their choices to a central clearinghouse, which would use this information to match students to jobs. Similarly, about 60 years ago, the umbrella organization of American sororities recommended that a centralized procedure be used to match students to sororities on college campuses.

The basic mechanism used to process the rank-orderings submitted by students and sororities is called the "preferential bidding system" (PBS), and it remains in use today. This paper analyzes the PBS algorithm, the setting in which it is employed, the incentives it gives to students and sororities, and the matchings that result.

*Department of Economics, University of Pittsburgh, Pittsburgh, PA 15260. Because of the requirement that the campuses should remain anonymous, we are unable to thank by name the many administrators without whose help this study could not have proceeded. We have received helpful comments from Patty Beeson. This work has been supported by grants from the Alfred P. Sloan Foundation, the National Science Foundation, and the Office of Naval Research.

This analysis reveals that the PBS algorithm is different in an important way from the algorithm around which the American medical market is organized and the algorithms around which some of the most successful and long-lived of the medical markets in the United Kingdom are organized. Those other algorithms have the property that they produce matchings that are *stable*, in the sense that, for markets of the kind considered here, they are in the core of the market, when agents state their true preferences.¹ The PBS algorithm does not have this property: the matchings it produces are stable in the preliminary market in which the algorithm operates, but they are not in the core of the market as a whole. Furthermore, for many configurations of preferences, the algorithm fails to produce a matching at all.

Nevertheless, when we examine data from 21 rushes on four campuses,² we observe only one such failure. A large part of the explanation appears after further examination of the data, which makes clear that the submitted preferences are unlikely to correspond to the true full preferences of the students. Instead, the observed pattern of preferences corresponds to what we would expect to see if the students respond strategically to the incentives induced by the matching procedure. When students do respond this way, the PBS procedure will not fail, and the resulting matching will be stable.

I. A Brief History

The first Greek-letter sorority was founded in 1870 (James Brown, 1920). A sorority may be present on campuses throughout the United States, and each

sorority location is called a chapter. In the literature of fraternities and sororities, from which we will sometimes quote, "fraternity" is used to mean either the all-male or the all-female social organizations, while "sorority" refers to the all-female organizations. Many sororities have joined a national organization, the National Panhellenic Conference (NPC), which consisted of 26 member sororities as of 1985. On each campus, all NPC sorority chapters are members of a College Panhellenic Council, the local governing body that determines rushing regulations.

Brown (1920 p. 14) described the early competition for members:

In the early days of the fraternities only seniors were admitted to membership, but the sharp rivalry for desirable men soon pushed the contest into the junior class, and so on down, until at some colleges it scarcely stops at the academy. The general rule is, however, that members shall be drawn from the four undergraduate classes. ... As the colleges usually open about the middle of September, the campaign for freshmen is then commenced and lasts until Christmas, when each chapter has secured its most desirable candidates. Where there is great rivalry, however, initiations take place all year round.

Earlier appointment dates were not the only evidence of competition:

Membership in two fraternities has been a source of trouble and vexation. It is almost universally forbidden. When it occurs between two chapters of different fraternities located at the same college, and a student leaves one and joins the other, it is termed "lifting," and such disloyalty is usually followed by expulsion. ... All of the fraternities now forbid this, although many years ago it was not uncommon. [Brown, 1920 pp. 15-16]

An early attempt to resolve these problems occurred in 1891, when the first meeting of sororities, in what was then called the

¹These markets involve many-to-one matching, since each student joins at most one sorority. In markets of many-to-many matching, stable matchings need not be in the core (see Roth, 1991).

²The data available to us come from campuses in which, loosely speaking, there is a "buyers' market" for sorority positions. We will discuss differences expected in "seller's markets."

Inter-Sorority Conference, was called to discuss intersorority cooperation. Although resolutions were passed decrying the practice of "lifting" and calling for the "abolition of the practice of pledging and initiating preparatory students" (National Panhellenic Council, 1983 p. 5), this had little effect. Similar sentiments were expressed in subsequent years, to equally little effect, and by 1928 the NPC was ready to turn to a centralized system of matching, and the first mention of the preferential bidding system appears.³ Francis Shephardson (1930 p. 8) reviews the events leading up to this:

The constant rivalry among chapters and the multiplication of fraternities have led in many cases to an indiscriminate scramble for members at the beginning of each year. Both fraternities and the colleges have perceived the danger of this sort of "rushing," as the contest for members is called, and are giving the subject thoughtful consideration. The deferred pledging of students until a fixed date and the deferred initiation of pledged members until they have completed a prescribed portion of their college course or secured a predetermined grade are both becoming common. Such procedure is in striking contrast with earlier custom in some of the larger Western and Southern colleges where, the preparatory schools being intimately connected with the colleges, "preps" were not only pledged, but initiated before they entered the college proper, or with the reprehensible custom which prevails in some places, where pledge pins are given out to boys in the high school or even in the grammar grades.

The preferential bidding system has since been incorporated into the recruiting activities of sororities, as described next.

II. The Organization of Recruiting Activities

The activities of a sorority seeking new members are called rush.⁴ There are two types of rush: formal rush and continuous open bidding. The NPC recommends "one formal rush period per year, held in the early fall, as close as possible to the start of the academic year, and conducted in as short a period of time as possible" (National Panhellenic Conference, 1979 p. 29).

Women participating in formal rush ("rushees") attend a sequence of parties designed to enable rushees and sororities to "narrow their choices gradually" (National Panhellenic Conference, 1979 p. 35). The first parties are "open houses" in which all sororities issue invitations to all rushees. In subsequent rounds, sororities issue invitations selectively. "Panhellenic strongly urges each sorority to re-invite ... only those rushees they are seriously considering for membership. This will enable both the rushee and the sororities to know 'how they stand' early in the formal rush period" (National Panhellenic Conference, 1979 p. 46). In each round, the number of sororities a rushee can attend is reduced. A rushee who receives more invitations than the number of parties permitted in a given round must decline, or "regret," the excess invitations. In the last round of invitational parties, the "preference parties," a rushee is usually permitted to attend only two or three parties. "Panhellenic strongly urges each sorority to invite *only* those rushees to the preference party to whom they will definitely issue a bid" (National Panhellenic Conference, 1979 p. 46).

After the last preference party, rushees indicate their sorority preferences on a card, which they sign. (A rushee who lists only a single sorority is said to have *suicided*.) Sororities similarly submit a preference or-

³See *National Panhellenic Review* (1985 p. 5) for a dated list of motions passed.

⁴The process described next is the recommended procedure appearing in the "*How To*" *Manual for College Panhellenics* (National Panhellenic Conference, 1979). While these rush procedures are not required, the essential features have been incorporated in each of the campuses we contacted.

dering of rushees. Once all preferences have been submitted, the PBS algorithm matches rushees to sororities.

Each sorority is eligible to be matched to up to *quota* (q) rushees during formal rush, where *quota* is "the number of rushees accepting at least one invitation to the first round of invitational parties, divided by the number of participating fraternities"⁵ (National Panhellenic Conference, 1979 p. 37).

Following the completion of the PBS algorithm, there is one more step in the formal rush process, which officially exists in two slightly different forms (and which in practice seems to vary somewhat more from campus to campus). Under the "quota-only" procedure, any sorority that has been assigned some number p of rushees by the PBS algorithm with $p < q$ is allowed to extend one additional set of at most $q - p$ bids to unmatched rushees. Under the "quota-plus" procedure, any sorority that has not been assigned q new members under the PBS algorithm or whose total membership $m + p$ (including the p new members) is below the total allowable chapter size, T (which is the same for all sororities on a given campus), is allowed to extend one additional set of at most $\max\{q - p, T - (m + p)\}$ bids to unmatched rushees. Rushees who were unmatched by the PBS algorithm are free to accept at most one of the bids they receive or to decline all such bids.

The results are announced on "pledge day," marking the end of formal rush. A rushee who enters formal rush by signing a preference card but who subsequently declines to join a sorority to which she has been matched, is not permitted to join another sorority for one year.

Continuous open bidding begins immediately after the close of formal rush. During continuous open bidding, any sorority that

has not received q new members or that has received q new members but is nevertheless below the total allowable chapter size is allowed to recruit additional members by simply extending them invitations to join. At this stage, sororities are not restricted to make a single set of bids but may recruit continuously until their membership reaches T (or, in the case of sororities whose initial membership m was greater than $T - q$, until they have recruited q new members).

This recruitment and matching process resembles those of the centralized medical labor markets (Roth, 1984a, 1991) mentioned in the Introduction: an information gathering period is followed by a centralized matching algorithm, which is followed by a decentralized "after-market." In the case of the medical labor markets, analysis of the matching algorithms is critical to understanding the matching process as a whole. We turn now to a detailed description of the PBS algorithm.

Rushees submit a "preference card" listing the sororities they would be willing to join, in order of preference. Sororities submit a "bid list" of rushees whom they would be willing to have as members. While a rushee can join no more than one sorority, every sorority is able to extend at least q (quota) invitations for new members through the formal rush process. Beyond the first-quota names, sororities list rushees in order of preference. These preference lists are used by the PBS algorithm to assign rushees to sororities. The instructions in Table 1 are from the manual "*How To*" for *College Panhellenics* (National Panhellenic Conference, 1979 pp. 41–2). These instructions are incomplete and contain ambiguous phrases, such as "This process is repeated as long as there is any possibility of a rushee receiving a bid from the fraternity of her first choice" and "When it becomes apparent that a rushee will not receive a bid from the fraternity of her first choice..." The NPC does have a pamphlet explaining the instructions of the PBS algorithm via an example to be conducted in a workshop. This example still does not handle some of the contingencies that may arise during an actual PBS execution.

⁵If this number is not an integer, it is rounded either up or down at the discretion of the individual supervising the rush. Quota can be rounded down without leaving some students unmatched, since rushees sometimes drop out of the formal rush process after the first round of invitational parties.

TABLE 1—SORORITY RUSHING INSTRUCTIONS:
THE PREFERENTIAL-BIDDING-SYSTEM ALGORITHM

Bid Lists:

1. At a specified time, each fraternity files with the Panhellenic Executive a list of women it wishes to bid.
 - a. Lists are in duplicate; one copy is used in bid matching, the other is returned to the chapter when the bid matching is completed.
 - b. The fraternity bid list should be on paper ruled into three columns:

Left-hand column—List in alphabetical order of fraternity's first choices up to the limit of quota.

Right-hand column—List in order of preference the fraternity's additional choices which may number as many as the chapter wishes to submit.

Center column—Is left blank, as this is the column in which the matched bids are entered.

As a bid is matched, the rushee's name is crossed off every fraternity's first or second list. Her name is entered in the center column of the fraternity list of the group to which she is being pledged.
2. Along with its bid lists, each fraternity brings to Panhellenic enough formal bids (in envelopes) for each woman to be pledged. These formal bids are to be addressed after bid matching is completed.

Procedure for Matching Bids:

1. Persons matching bids include the reader, the tabulator, and one alumna handling the bid list from her fraternity. Undergraduates are not to participate in bid matching.
2. Before bid matching begins, names of all rushees who chose not to sign a preference card should be crossed off all preference lists, and those lists adjusted to fill the space of these women.
3. Mechanics:
 - a. After alphabetizing the preference cards, the reader calls the rushee's name and her first choice. If the fraternity of her first choice has given her a bid on its first bid list, it is a matched bid, and all others should cross her from their list. If the rushee's name is not on the fraternity's first bid list, her preference card is temporarily laid aside. Names of rushees who list only one preference and are unmatched at the end of the first reading should be crossed off all other bid lists and their cards laid aside.
 - b. Each time a name is crossed off a fraternity's first bid list, if openings in the fraternity's pledge quota remain, a name from the fraternity's second bid list is added, in the listed order, to the bottom of the unmatched names remaining on the first list. *The number of unmatched names on the adjusted first bid list and the number of those pledged must always equal quota* (unless a chapter has run out of names to add from its second bid list).
 - c. The cards laid aside in step "a" are read again according to the first choice of the rushee. This process is repeated as long as there is any possibility of the rushee receiving a bid from the fraternity of her first choice.
 - d. Those cards remaining are those of rushees whose names are on the second bid list of the fraternities of their first choice.
 - e. When it becomes apparent a rushee will not receive a bid from the fraternity of her first choice, a rushee's second choice is then matched, if possible, in the above manner.
 - f. Any remaining cards are then read according to the rushee's third choice and the same procedure followed.
 - g. The tabulator reads the results, and all bid lists are reviewed for accuracy.
 - h. Unmatched bids—If a rushee's preference card has failed to match for a bid, the Panhellenic Executive may contact the rushee and ask if she will accept a bid from a fraternity not previously listed among her choices, if this other fraternity has her name on one of their bid lists. Any rushee not bid by any of her preference choices is eligible at any future time for rushing and pledging by any fraternity.

Unfilled Quotas—If a fraternity has failed to fill its quota through this bid matching in formal rush, it may be contacted by the Panhellenic Executive to ask if the fraternity wishes to extend a bid to anyone not originally on its bid lists.

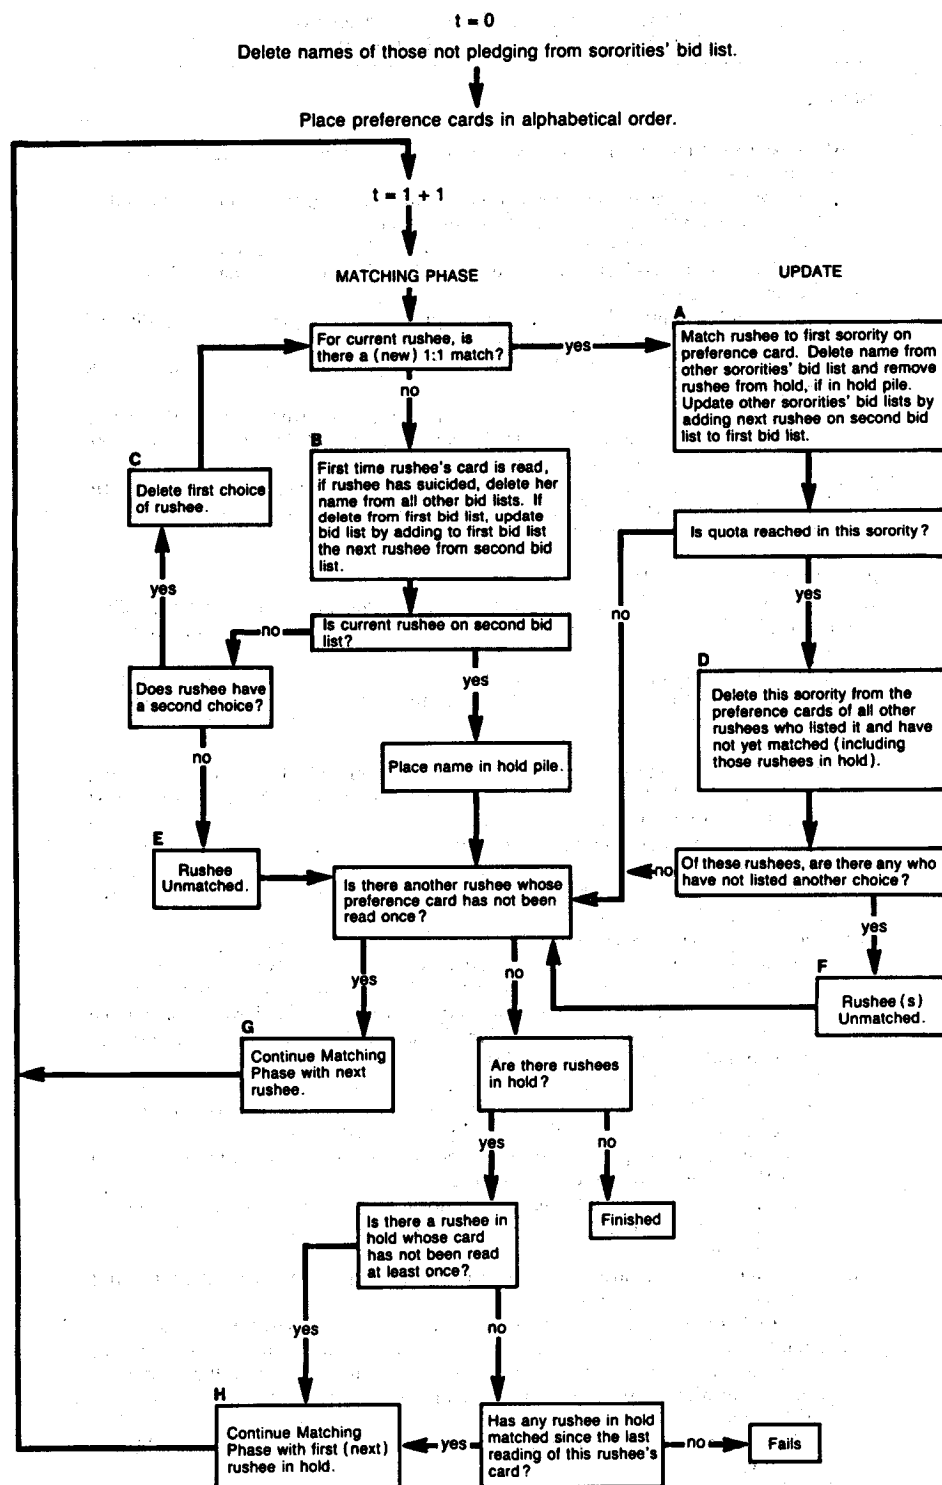


FIGURE 1. PBS-ALGORITHM FLOW CHART

When the instructions given for the PBS algorithm do not indicate what should be done with those rushees whose cards have been "laid aside," we will say that the algorithm "fails."⁶ When we examine the data, it will be seen that, in practice, the PBS algorithm very seldom fails. Indeed, the individuals in charge of administering the algorithm on each of the campuses from which our data are drawn were all initially unaware of the possibility of this kind of failure.⁷

Figure 1 is a flow chart of the PBS algorithm. No such flow chart is found in the sorority literature: this was compiled from both the literature mentioned above and interviews with individuals charged with supervising the matching process on some of the campuses contacted. The original bid list (before any rushees who have not signed a preference card have been deleted) is employed at step $t = 0$. Each time a rushee's preference card is read, t increases by 1.

The final step of the formal rush procedure, during which one set of additional bids may be made (see item h in Table 1 for one variation) has been omitted from the flowchart. Unlike the operation of the algorithm, such additional bids need input from the participants in addition to their initial preference lists. We will consider any such additional bids when we analyze the aftermath of the algorithm.

III. A Formal Model

The first elements of our formal model are two finite and disjoint sets of sororities

⁶For example, consider the case of two rushees and two sororities with $q = 1$. If rushee r_1 ranks sorority S_2 before S_1 , rushee r_2 ranks S_1 before S_2 , sorority S_1 ranks r_1 before r_2 , and S_2 ranks r_2 before r_1 , then both rushees will remain in hold, and the algorithm will fail.

⁷When subsequently presented with examples contrived so as to cause the algorithm to fail, these individuals suggested a variety of ad hoc procedures for restarting the algorithm and completing the matching procedure. Therefore, in the flow chart (Fig. 1), the box labelled "fails" can be viewed as a point in the algorithm in which the implementation would be different on different campuses.

and of rushees, $S = \{S_1, \dots, S_n\}$ and $R = \{r_1, \dots, r_m\}$, respectively. Each rushee has preferences over the sororities, and each sorority has preferences over the rushees. We will assume that these preferences are complete and transitive, with $P(S) = r_1, r_2, S, r_3, \dots$ denoting that sorority S prefers to enroll r_1 rather than r_2 , that it prefers to enroll either one of them rather than leave a position unfilled, and that all other rushees are unacceptable, in the sense that S prefers to leave a position unfilled rather than filling it with, say, rushee r_3 . Similarly, $P(r) = S_2, S_1, S_3, r, \dots$ represents the preferences of rushee r , indicating for example that the only positions the rushee would accept are those offered by S_2, S_1 , and S_3 , in that order. Sorority S is *acceptable* to rushee r if r prefers to be matched to S rather than to remain unmatched, and rushee r is *acceptable* to sorority S if S prefers to have r as a member rather than to leave a position unfilled.

The number of positions each sorority can fill during formal rush is q (quota). An outcome of the PBS algorithm is a matching of rushees to sororities, such that each rushee is matched to at most one sorority and each sorority is matched to at most q rushees. After formal rush (i.e., during the continuous open bidding which follows), each sorority S_k may have a different quota q_k , depending on its membership prior to the start of formal rush and on how many new members it has been matched to during formal rush.

A rushee who is not matched to any sorority will be modeled as "matched to herself," and a sorority with some number of unfilled positions will be matched to itself in each of those positions. A rushee is matched to a given sorority if and only if the sorority is matched to that rushee. To give a formal definition, first define, for any set X , an *unordered family of elements* of X to be a collection of elements, not necessarily distinct. Thus, an element of X may appear more than once, which distinguishes an unordered family from a subset of X .

Definition: A matching μ is a function from the set $S \cup R$ into the set of unordered

families of elements of $S \cup R$ such that:

- 1) $|\mu(r)| = 1$ for every rushee r and $\mu(r) = r$ if $\mu(r) \notin S$;
- 2) $|\mu(S)| = q$ for every sorority S , and if the number of rushees in $\mu(S)$, say p , is less than q , then $\mu(S)$ contains $q - p$ copies of S ;
- 3) $\mu(r) = S$ if and only if r is in $\mu(S)$.

Thus, $\mu(r_1) = S$ denotes that rushee r_1 is enrolled at sorority S at the matching μ , and $\mu(S) = \{r_1, r_3, S, S\}$ denotes that sorority S , with $q = 4$, enrolls rushees r_1 and r_3 and has two positions unfilled. (When sororities have different quotas, q_k replaces q for each sorority S_k .)

Each rushee's preferences over alternative matchings correspond exactly to her preferences over her own assignments at the two matchings. While we have described sororities' preferences over rushees, when q is greater than 1, each sorority must be able to compare groups of rushees in order to compare alternative matchings, and we have yet to describe the preferences of sororities over groups of rushees. (Until we have described sororities' preferences over matchings, our model will not be a well-defined game.)

The simplest assumption connecting sororities' preferences over groups of rushees to their preferences over individual rushees is one insuring that, for example, if $\mu(S)$ assigns sorority S its third- and fourth-choice rushees, and $\mu'(S)$ assigns it its second- and fourth-choice rushees, then sorority S prefers $\mu'(S)$ to $\mu(S)$. Specifically, let $P^*(S)$ denote the preference relation of sorority S over all assignments $\mu(S)$ it could receive at some matching μ . A sorority S 's preferences $P^*(S)$ will be called "responsive" to its preferences $P(S)$ over individual rushees if, for any two assignments that differ in only one rushee, it prefers the assignment containing the more preferred rushee, as described formally in the following:

Definition: The preference relation $P^*(S)$ over groups of rushees is responsive [to the

preferences $P(S)$ over individual rushees] if, whenever $\mu'(S) = \mu(S) \cup \{r_k\} \setminus \{\sigma\}$ for σ in $\mu(S)$ and r_k not in $\mu(S)$, then S prefers $\mu'(S)$ to $\mu(S)$ [under $P^*(S)$] if and only if S prefers r_k to σ [under $P(S)$]. (Subtraction of sets is denoted by \setminus .)

Note that S may be indifferent between distinct assignments $\mu(S)$ and $\mu'(S)$ even if S has strict preferences over individual rushees.

A matching μ is *individually irrational* if $\mu(r) = S$ for some rushee r and sorority S with either r unacceptable to S or S unacceptable to r . Such a matching is *blocked* by the unhappy agent. This terminology reflects that the rules allow every agent to withhold her (or its) consent from such a match. Similarly, a sorority S and rushee r will be said to block a matching μ together if they are not matched to one another at μ but would both prefer to be matched to one another rather than to (one of) their present assignments. That is, μ is blocked by the sorority-rushee pair (S, r) if $\mu(r) \neq S$ and if r prefers S to $\mu(r)$ and S prefers r to σ for some σ in $\mu(S)$. (Note that σ may equal either some rushee r' in $\mu(S)$ or, if at least one of sorority S 's positions is unfilled at $\mu(S)$, σ may equal S .) Matchings blocked by an individual or by a pair of agents are unstable in the sense that there are agents with the incentive and the power to disrupt such matchings.

Definition: A matching μ is stable if it is not blocked by any individual agent or any sorority-rushee pair.

It is not obvious that this definition will be adequate, since we might need to consider coalitions consisting of sororities and several rushees (all of whom might be able to enroll at the sorority). However, it can be shown that considering larger coalitions would not change the set of stable outcomes, which equals the core of the game with respect to weak domination (see Roth, 1985b; Roth and Marilda Sotomayor, 1990).

Sets S of sororities and R of rushees, together with a vector P of preferences, one

for each agent, constitute a *matching market*.⁸ In what follows, we will assume for simplicity that all preferences over individuals are strict (i.e., that no sorority is indifferent between two acceptable rushees and no rushee is indifferent between two acceptable sororities). (We do not assume that sororities may not be indifferent between different groups of rushees.)

IV. Analysis of the Algorithm

We begin with a model of the market up to the conclusion of the PBS algorithm: in this part of the market, each sorority may admit q new members (q is the same for all sororities). We will refer to this as the market with quota q .

Some notation will help describe the working of the algorithm. Denote by $x'(r_i) = S_j$ that rushee r_i was matched to sorority S_j during step t , where a *step* is the working of the algorithm associated with a reading of a single rushee's preference card. Denote by $x'(r_i) = r_i$ that rushee r_i was assigned as unmatched during step t . At step t when a rushee r_i 's preference card is read, if r_i neither matches to a sorority nor is assigned as unmatched, then her preference card will be placed in "hold" and will be reread after all other rushees who have yet to be assigned (as unmatched or matched to a sorority) have had their preference cards read. Finally, denote by $x(r) = S$ that rushee r was matched to sorority S at some step of the algorithm, and similarly denote by $x(r) = r$ that rushee r was assigned to be unmatched; define $x(S)$ to be the set of all rushees assigned to S (i.e., $x(S) = \{r | x(r) = S\}$). Note that x is not a matching, because it is not defined for all rushees, but only for

those not left in hold when the algorithm ends, and because $|x(S)|$ may be less than q (the remaining positions are not filled with copies of S).

Sororities indicate preferences by listing rushees on a first bid list of no more than q names and a second bid list. Denote by $r_i \in Q_t(S_k)$ that rushee r_i is listed on the first bid list of sorority S_k at step t in the algorithm. That rushee r_j is listed on the second bid list of S_k at step t in the algorithm is denoted by $r_j \in Q_t^+(S_k)$. For each sorority, the bid list at step $t=0$ is the original bid list.

Consider next what the algorithm does when confronted with preferences for which it is well-defined (i.e., for which it does not fail to produce a matching).

THEOREM 1: *If no rushees are left in "hold" at the end of the PBS algorithm, its outcome is stable in the market with quota q .*

Theorem 1 is proved in the Appendix.

Furthermore, the PBS algorithm has the property that all its assignments are *inevitable*, in the sense that all rushees who match to sororities by the PBS algorithm must match to the same sorority at every stable outcome and rushees assigned as unmatched by the algorithm must be unmatched at every stable outcome in the market with quota q . That is, we have the following result (proved in the Appendix).

THEOREM 2: *The preferential-bidding-system algorithm only makes inevitable assignments in the market with quota q .*

An immediate consequence of the theorem is the following.

COROLLARY: *The preferential bidding system assigns all rushees only when there exists a unique stable outcome in the market with quota q .⁹*

⁹The converse is not true; it is possible to construct examples in which the algorithm fails to produce a matching even though there is a unique stable matching.

⁸This matching market is an example of what is sometimes called a "college admissions" model (see David Gale and Lloyd Shapley, 1962). When quotas all equal 1, the model is symmetric between both sides of the market and is called the "marriage model." For many years, it was thought that the college admissions model was essentially equivalent to the marriage model. That this is not the case was shown in Roth (1985a). The model presented here is the college admissions model as reformulated in Roth (1985a).

The corollary confronts us squarely with a puzzle. Typically there may be many stable outcomes to this kind of two-sided matching market, but the PBS algorithm is rarely observed to fail. To see how these two observations may be resolved, we will examine data from a number of rushes in Section V; but first consider the operation of the PBS algorithm as part of the larger market in which sororities may be able to admit more than q new members, even though they are not allowed to fill more than q positions through the algorithm. Specifically, consider a market in which the membership of each sorority S_k before formal rush is m_k , in which quota for the PBS algorithm is q , and in which the total allowable size for any sorority is T . Then, the number of positions sorority S_k may be able to fill either through formal rush or the informal rush which follows is $q_k = \max\{q, T - m_k\}$. That is, every sorority has the right to fill up to q positions (whether or not this will bring membership above T), and any sorority that has not filled q positions or does not have T members at the end of formal rush is able to continue recruiting new members. We have the following result.

THEOREM 3: *In the market with quotas q_k , matchings produced by the PBS algorithm with quota q may not be stable.*

PROOF:

Suppose the first $q + 1$ rushees on the bid list of some sorority S_k with $q_k > q$ all list S_k as their first choice. Further suppose that the $(q + 1)$ th rushee on S_k 's bid list, r_{q+1} , lists sorority S' as her second choice and that S' lists r_{q+1} among its first q rushees. Then, if the PBS algorithm with quota q results in a matching μ , Theorem 1 implies that $\mu(r_{q+1}) = S'$; but μ is unstable in the market with quotas q_k , since in that market S_k has a vacant position, and μ is blocked by (S_k, r_{q+1}) .

There is ample reason (both empirical and theoretical) to believe that instabilities give agents strong incentives to circumvent the procedures that produce them. Therefore, Theorem 3 raises a further question

about how the PBS algorithm has survived for so long. The empirical observations reported next will shed some light on this.

V. Some Empirical Observations

Preference cards and bid lists from formal rush were solicited from 12 campuses. These are regarded as highly confidential, and only four of the campuses agreed to make this material available, and then only under the condition that neither the names of sororities and rushees nor the campuses themselves would appear in any report. The data from 21 recent PBS-algorithm assignments taken from the four campuses are summarized in Tables 2, 3, and 4.¹⁰

Campus A is a rural college with approximately 1,500 full-time students; B is an urban university with a full-time undergraduate enrollment of about 4,500; C is a university in a rural setting with approximately 10,400 full-time undergraduates; and D is an urban university with roughly 9,400 full-time undergraduates. These four campuses are not a representative sample. All are located in the northeastern United States, and each had many sororities whose membership was sufficiently below their maximum capacity (their "total") so as to pose only loose constraints on the number of bids they could issue after formal rush.

As this latter factor will play a role in our subsequent analysis, the number of sororities that have "constrained" and "unconstrained" totals is noted below each table. Operationally, a sorority was said to be constrained only if the number of rushees on its second bid list who listed that sorority as their first choice was greater than the number of positions the sorority had available after formal rush (see Section VI). Other-

¹⁰In all but one case, the reported statistics are based upon the original preference lists (these were not available in 1986 on campus C: the statistics for that year were compiled by an administrator with direct access to the data). Some of the campuses retained old records and had many past PBS assignments available. Others only kept the most recent PBS assignments.

TABLE 2—CAMPUS A DATA

Statistic	1982 ^a	1983	Spring 1984	Fall 1984 ^a	1987
Number of rushees	93	84	72	92	68
Number of suicides	47	38	49	54	47
Matched	38 (41)	33	34	51 (46)	38
Unmatched	9 (6)	5	15	3 (8)	9
(Not listed)	0	5	4	0	0
Number with two choices	46	44	23	38	21
Matched	44 (43)	40	22	38 (37)	21
First choice	34 (35)	38	18	37 (35)	18
Second choice	10 (8)	2	4	1 (2)	3
Unmatched	2 (3)	4	1	0 (1)	0
(Not listed)	9	11	3	0	0
Number with three choices	0	2	0	0	0
Matched ^b	0	2	0	0	0
Unmatched	0	0	0	0	0
(Not listed)	0	4	0	0	0
Total matched (percentage)	88.2 (90.3)	89.3	77.8	95.7 (90.2)	86.8
First choice	77.4 (81.7)	86.9	72.2	95.6 (88.0)	82.4
Second choice	10.8 (8.6)	2.4	5.6	1.1 (2.2)	4.4
Total unmatched (percentage)	11.8 (9.7)	10.7	22.2	3.3 (9.8)	13.2
Suicides (percentage)	50.5	45.2	68.1	58.7	69.1

Notes: The maximum chapter size (*T*) was 55. Of the five sororities on campus A, two were constrained and three were unconstrained. There were two periods of formal rush in 1984. In 1984, the timing of formal rush was changed from spring (1982–1984) to fall. The 1985 fall formal rush results were unavailable. In 1986, formal rush was again changed to the spring.

^aIn 1982, an error occurred in the execution of the PBS algorithm. In fall of 1984, quota was incorrectly determined to be 25 when it should have been 21.8 or 22. The numbers shown in parentheses are the correct statistics based upon the correct assignments. All statistical tests are based upon the statistics resulting from the actual (not the correct) assignments.

^bBoth rushees who listed three choices matched to their first choice. Note that neither rushee was listed by her second- or third-choice sorority.

wise, the sorority was said to be unconstrained. With one exception, the constrained status of each sorority has remained unchanged over the years under observation.¹¹

The most striking feature of the data is the high percentage of rushees who chose to

list only one sorority on their preference card. This is particularly striking in view of the fact that this practice ("suiciding") is explicitly discouraged.¹² Nevertheless, of the

¹²The following suggestions or guidelines were listed in an orientation booklet distributed to rushees during formal rush on one of the campuses we studied:

...if a rushee does not receive her first choice, she must be willing to accept any of the other choices she has listed. However, if she only preferences one sorority (sometimes called

¹¹In the most recent PBS assignment occurring on campus D (1987), some sororities became constrained (see notes to Table 4).

TABLE 3—CAMPUS B DATA

Statistic	Spring 1979	Fall 1980	Spring 1981	1982	1983	1984	1985	1986	1987
Number of rushees	62	70	57	82	91	76	102	96	125
Number of suicides	53	58	53	79	67	51	86	73	80
Matched	52	51	38	56	57	47	75	50	63
Unmatched	1	7	15	23	10	4	11	23	17
(Not listed)	1	4	0	9	4	3	3	11	8
Number with two choices	9	11	4	2	23	24	15	20	34
Matched	9	11	3	1	21	23	13	18	31
First choice	9	10	3	1	15	19	10	13	23
Second choice	0	1	0	0	6	4	3	5	8
Unmatched	0	0	1	1	2	1	2	2	3
(Not listed)	1	2	1	0	9	6	4	8	17
Number with three choices	0	1	0	1	1	1	1	3	11
Matched	0	1	0	1	1	1	1	2	8
First choice	0	1	0	1	0	0	1	2	7
Second choice	0	0	0	0	0	1	0	0	1
Third choice	0	0	0	0	1	0	0	0	0
Unmatched	0	0	0	0	0	0	0	0	3
(Not listed)	0	0	0	0	1	1	0	4	14
Total matched (percentage)	98.4	90	71.9	70.7	86.8	93.4	87.3	72.9	81.6
First choice	98.4	88.6	71.9	70.7	79.1	86.8	84.3	67.7	74.4
Second choice	0	1.4	0	0	6.6	6.6	2.9	5.2	7.2
Third choice	0	0	0	0	1.1	0	0	0	0
Total unmatched (percentage)	1.6	10	28.1	29.3	13.2	6.6	12.7	27.1	18.4
Suicides (percentage)	85.5	82.9	93	96.3	73.6	67.1	84.3	76	64

Notes: The maximum chapter size (T) was 50. Of the six sororities on campus B, two were constrained and four were unconstrained. This campus had two formal rush periods every year, fall and spring, until 1982. The 1982 data represent the first year that there was only one formal rush period, held in the spring. Formal rush has continued to be held in the spring since 1982. There are two missing observations: spring 1980 and fall 1981. In 1986, an error occurred in the execution of the PBS algorithm. This error had no effect upon the aggregated statistics.

21 rushes observed on four campuses, there were only three in which the number of rushees suiciding was less than 50 percent of those who submitted preference cards. Even on campuses C and D, which each have a dozen or more sororities active in formal rush, relatively few rushees list more

than two sororities on their preference cards.

In Tables 2–4, the number of “suicides” is given immediately below the number of rushees submitting preference cards (with the percentage given in the last line of the table). For comparison, the number of rushees listing two choices on their preference lists and the number with three choices are also given. For each of these, the table shows how many were matched by the algorithm, broken down into how many are matched to their first, second, and third choices. Also shown for each category of rushees is the number of times a rushee

“suiciding”) she must realize she is limiting her chances of pledging a sorority all together.

No sorority shall encourage a rushee to single preference their sorority (suicide).

TABLE 4—CAMPUS C DATA AND CAMPUS D DATA

Statistic	Campus C			Campus D			
	1984	1985	1986	1984	1985	1986 ^a	1987
Number of rushees	59	79	93	89	78	96	119
Number of suicides	35	41	44	34	54	57	57
Matched	32	36	30	22	38	44	48
Unmatched	3	5	14	12	16	13	9
(Not listed)	0	0	1	0	0	0	0
Number with two choices	24	38	49	52	16	31	44
Matched	23	38	49	47	16	29	40
First choice	20	37	46	42	15	25 (24)	30
Second choice	3	1	3	5	1	4 (5)	10
Unmatched	1	0	0	5	0	2	4
(Not listed)	0	0	0	4	0	3	4
Number with three choices	0	0	0	3	8	8	18
Matched	0	0	0	3	8	8	17
First choice	0	0	0	3	7	6	13
Second choice	0	0	0	0	1	1	4
Third choice	0	0	0	0	0	1	0
Unmatched	0	0	0	0	0	0	1
(Not listed)	0	0	0	4	8	7	18
Total matched (percentage)	93.2	93.7	85.0	80.9	79.5	84.4	88.2
First choice	88.1	92.4	81.7	75.3	76.9	78.2	76.5
Second choice	5.1	1.3	3.2	5.6	2.6	5.2	11.7
Third choice	0	0	0	0	0	1.0	0
Total unmatched (percentage)	6.8	6.3	15.1	19.1	20.5	15.6	11.8
Suicides (percentage)	59.3	51.8	47.3	38.2	69.2	59.4	47.9

Notes: The maximum chapter size (*T*) was 65 on campus C and 55 on campus D. All of the 13 sororities on campus C and all of the 12 sororities on campus D were unconstrained during 1984–1986; in 1987, three of the sororities on campus D were constrained, and nine were unconstrained. Campus C requires that a sorority list all rushees who were extended a bid to its final party somewhere on its bid list. All rushes for both campus C and D take place in the fall. Quota-plus was adopted during the 1984 formal rush on Campus D, quota-only was adopted for all other years (1985–1987). On Campus D, sorority totals became relevant for the first time in 1987. The PBS algorithm failed to assign all rushees in 1987.

^aAn error occurred in the execution of the PBS algorithm on Campus D in 1986. The numbers shown in parentheses are the correct statistics based upon the correct assignments. All statistical tests are based upon the statistics resulting from the actual (not the correct) assignments.

placed on her preference card a sorority that did not in turn list the rushee, either on the first or second bid list.

For example, Table 3 shows that, in the 1979 rush on campus B, 62 rushees signed preference cards. Of these, there were 53 rushees who listed only a single sorority, and 52 of these were matched to that sorority, while one was unmatched. In this case, we can see that this is because one rushee

who listed only one sorority was not listed on that sorority's bid list. Similarly, nine rushees listed two sororities, and all nine were matched to their first choice. One of the sororities so listed did not place one of these rushees on either of its bid lists. Over 98 percent of the rushees were matched (all to the first choice on their preference cards), in a rush in which over 85 percent listed only one sorority on their preference card.

TABLE 5—DATA ON CONTINUOUS OPEN BIDDING

Campus	Year	Number of rushees unmatched during formal rush		Number of rushees matched to first choice during continuous open bidding	
		Suicides	Nonsuicides	Suicides	Nonsuicides
C	1985	5	0	1	0
C	1986	14	0	1	0
D	1984	12	5	11	4
D	1985	16	0	14	0
D	1986	13	2	13	2
D	1987	3 (9)	2 (5)	2 (4)	2 (3)

Notes: The 1984 data on campus C were unavailable. The 1987 results for campus D are broken into two groups, depending on whether the first-choice sorority was constrained or unconstrained. The first number is the number of rushees whose first choice was unconstrained, while the number in parentheses is the total number of rushees (i.e., those whose first choice was constrained plus those whose first choice was unconstrained). For example, nine rushees listed only a single sorority, of which three listed an unconstrained sorority. Of these three, two matched to their first choice in continuous open bidding.

On this campus, the maximum membership allowed in a sorority (T), which is the same for each sorority on campus, is 50, and of the six sororities on campus, only two were near enough this number so that it could constrain their post-PBS bidding.

The maximum number of rushees each sorority can be assigned under the PBS algorithm (quota) will vary each year. Quota is the number of rushees attending the first round of invitational parties divided by the number of sororities on the campus. Quota is not shown in the tables nor can it be calculated from the information given. The "number of rushees" shown in the tables is the number signing preference cards, which may be substantially smaller than the number of rushees attending the first round of preference parties.

The PBS algorithm failed to assign all rushees (as either matched to a sorority or as unmatched) on campus D during the 1987 formal rush. This was the only failure observed. Those rushees not assigned by the PBS algorithm were assigned by the individual in charge of the execution of the PBS algorithm. (The resulting matching was stable; see Mongell [1988] for an analysis of this incident.) The statistics in Tables 2–4 indicate the assignments made by the PBS algorithm. The quota-only method was

adopted by all the campuses observed, except for one year (1984) on campus D. Under quota-only, sororities may extend additional bids to rushees assigned as unmatched by the PBS algorithm. The procedure by which these additional bids are extended varies on each campus (and sometimes from year to year). The sorority may be notified that an unmatched rushee's name appears on its bid list and asked whether it would like to extend her a bid. The sororities may be notified before the PBS execution that unmatched rushees whose names appear on their bid list will be extended bids (on the sorority's behalf) if the sorority has not reached quota during the PBS algorithm. From the available data it was observed that few sororities extended bids at this time. On some campuses, a rushee assigned as unmatched by the PBS algorithm will be called by one of the individuals involved with the PBS execution and asked whether she would be willing to join another sorority that listed her on its bid list and has not reached quota. It appears from the (limited) available evidence on this point that virtually all rushees so called have refused these bids.

The numbers of rushees assigned as unmatched by the PBS algorithm who match to their first choice during continuous open

bidding were available on campuses C and D and are shown in Table 5.

VI. Strategic Analysis

The data raise two questions. What accounts for the consistently high percentage of rushees who list only a single sorority on their preference cards? And how might this high percentage be related to the low frequency of failure of the PBS algorithm and to its long life? To address these questions requires a model of the entire rush process.

During the PBS algorithm, each sorority may gain up to q new members. Following the PBS algorithm, there is a second stage of formal rush during which one additional set of bids and acceptances or rejections may be made (with how many bids depending on whether the quota-only or quota-plus rules are adopted). Finally, following formal rush there is continuous open bidding, during which each sorority with fewer than T members (both new members and old members who have not yet graduated) may admit new members to bring its membership up to T (and each sorority that has not yet enrolled q new members may bring its new members up to q). On the campuses from which our data are drawn, T imposed such a loose constraint that most sororities could attempt to recruit all rushees who showed serious interest in them. (These are the "unconstrained" sororities noted below in Tables 2–4.) That is, on these campuses, the demand for membership in most sororities is less than the supply, in the sense formalized by the following definition.

Definition: A sorority S is unconstrained if its membership is sufficiently below its allowed total so that it can extend bids to all rushees who it finds acceptable and who have S as their first choice among all sororities that find them acceptable.

Nevertheless, the number of rushees interested in joining even an unconstrained sorority may exceed q . As we saw in the proof of Theorem 3, a rushee who lists more than one sorority on her preference card runs the risk of being matched to her

second-choice sorority during the PBS algorithm and forgoing a chance to be matched to her first-choice sorority after the formal rush. For simplicity, consider the case in which all sororities have unconstrained totals (i.e., for every sorority S_k , the number of acceptable rushees who regard S_k as their first choice is less than q_k , where $q_k = \max\{q, T - m_k\}$ is the number of positions S_k may fill by the end of open bidding).

We model the matching procedure as a multistage game. In stage 1, all sororities and rushees simultaneously state preferences and are matched by the PBS algorithm. In stage 2, the unmatched rushees are announced, as are the sororities that have not filled q positions (if the quota-only rules are used, or which have fewer than T members, if the quota-plus rules are used). Each such sorority S_k may issue invitations to up to $q - |x(S_k)|$ (or $q_k - |x(S_k)|$) unmatched rushees, and each rushee who receives invitations may accept at most one. In stage 3 and subsequent stages, all matches from previous stages become public, and any sorority S_k that has been matched to a set $y(S_k)$ of rushees in the prior stages may issue invitations to up to $q_k - |y(S_k)|$ rushees who have not been matched to sororities in earlier stages. Each rushee may accept at most one invitation and must decline all others when they are received; at any stage in which she accepts an invitation, she is matched. Starting with stage 4, no sorority may issue an invitation to a rushee to whom it has previously issued an invitation at stage 3 or later. The game ends at any stage in which no invitations are issued. Stages 1 and 2 represent formal rush, with stage 1 corresponding to the PBS algorithm and stage 2 corresponding to the quota-only (or quota-plus) system. Subsequent stages represent open bidding.

Note that we have chosen one of several possible ways to model the second stage of formal rush. Also, by dividing the bidding into stages we have imposed on the model some structure beyond what we observe in practice in open bidding. Finally, so the game will end in finitely many periods, we have imposed the rules that sororities may

not reinvoke rushees and that rushees must either accept or reject all invitations in the period they are received. These choices seem to be among the simplest that are broadly consistent with the sometimes diverse and sometimes ambiguous rules of the rush process. Because there is some irreducible arbitrariness in choosing the elements of a model, it is also important to note that the equilibrium considered below seems robust to changes in these arbitrary features of the model.

We would like to demonstrate that the observed behavior corresponds to equilibrium behavior in this market. One potential difficulty we face is that we have not fully specified what happens when the PBS algorithm fails. To show that a particular set of strategies is in equilibrium, we have to show that no agent can profitably deviate, and for this we have to show that no agent can profitably deviate even in a way that causes the algorithm to fail. As noted earlier, different individuals charged with supervising sorority rush have indicated that they would proceed differently in the circumstances we call failure (i.e., in these circumstances, the results would be different on different campuses). One approach, therefore, would be to make further assumptions about how the algorithm would proceed on each of these campuses. We take a different approach and demonstrate an equilibrium with the properties that only rushees can deviate in a way that might cause the algorithm to fail and no rushee can profit from this, no matter how failures are resolved (however failures of the PBS algorithm might be resolved on different campuses, rushees may not be matched to sororities that have not issued them invitations).

THEOREM 4: *Suppose all sororities are unconstrained. Then:*

(a) *The following strategies constitute an equilibrium in the multistage game. In the first stage, each rushee lists only her first-choice acceptable sorority (from among those that find her acceptable) and in subsequent stages accepts only an offer from this sorority. Each sorority S_k lists its true preferences in the first stage and makes offers in stage 2 to*

its most preferred $q_k - |x(S_k)|$ acceptable rushees from among those unmatched in stage 1. In stage 3, it makes offers to all of its $(q_k - |y(S_k)|)$ highest-ranked unmatched rushees, and in any subsequent stage it makes offers to its highest-ranked set of unmatched rushees who have not previously rejected it.

(b) *Furthermore, at this equilibrium the PBS algorithm never fails, and the matching that results is stable in the market with quotas q_k (it is the rushee-optimal stable matching).*

PROOF:

We prove part b first. Suppose the rushees and sororities play the strategies described. To see that the PBS algorithm does not fail, suppose to the contrary that it ends with some rushee r_i in hold. Then r_i must be on the second bid list of the sorority on the top of her preferences at the final step of the algorithm, and this sorority, S , must not have reached quota (since if it had it would have been crossed off r_i 's preference list at box D of the flow chart in Fig. 1). Therefore, there must be another rushee, r_j , not matched to S but in the first q positions of S 's final bid list. However, this cannot be, since r_j has listed only one sorority. If this is not S , then r_j would have been crossed off S 's list at box B of the flow chart, and if it is S , then (all such) r_j would be matched to S , contradicting that S has not reached quota. Therefore, the PBS algorithm leaves no rushees in hold; that is, it does not fail.

When all agents play these strategies, each rushee is eventually matched to her first choice among all acceptable sororities that find her acceptable. Since this matching is individually rational and not blocked by any sorority-rushee pair, it is stable, and since each rushee is matched to her highest-ranked achievable match, it is the rushee-optimal stable matching μ_R . This completes the proof of part b.

To prove part a, we show that no sorority or rushee can do better than to play the strategy described, so long as the other agents all do so. First consider sororities. As we saw in the proof of part b, so long as all rushees list only a single sorority on their preference cards, the algorithm will not fail, regardless of what sororities may do. Fur-

thermore, if the rushees follow the strategies indicated in the theorem, any sorority S that deviates from the strategy indicated for it will be matched to a subset of $\mu_R(S)$, rather than to all of $\mu_R(S)$. Since S has responsive preferences over groups of rushees, and since all rushees in $\mu_R(S)$ are strictly preferred to vacant positions, S prefers $\mu_R(S)$ to any strict subset of $\mu_R(S)$, and so cannot profit from any such deviation.

Now consider rushees. For each r , $\mu_R(r)$ is r 's most-preferred mutually acceptable sorority; that is, $\mu_R(r)$ is the most preferred match r can achieve at any individually rational outcome (if there are no mutually acceptable sororities, $\mu_R(r) = r$.) Therefore, if r deviates from her indicated strategy, she cannot improve her outcome even if by deviating she causes the PBS algorithm to fail, since no rushee may be matched to a sorority that has not issued her an invitation.

While the equilibrium specified in Theorem 4 is not a perfect equilibrium, the equilibrium behavior is certainly consistent with perfectness. Since the extensive-form game begins with the simultaneous submission of all parties' preferences, all equilibria are subgame perfect. However, after formal rush, all parties learn all the payoff-relevant information of the game, and the subsequent information sets all consist of single nodes, so an appropriate formulation of perfectness is backward induction to the nodes of stage 3. The off-the-equilibrium-path behavior we must consider arises if a rushee's first-choice sorority fills all its positions before issuing her an invitation. In this event, the rushee's strategy should be, from stage 3 onward, to accept the offer from her highest-ranked sorority among those that will still have positions to offer when they reach her, following the strategy for sororities given in the theorem.¹³ Note also that

the stage-2 behavior of sororities plays little role in this equilibrium (e.g., nothing would change if sororities made no offers in stage 2 but otherwise behaved as in Theorem 4).

Theorem 4 considers the case in which all sororities are unconstrained, whereas in our data this was the case only on campuses C and D; both campuses A and B had some constrained sororities, although most were unconstrained. Thus, the assumptions of the theorem do not precisely model the situation we observed any more than the equilibrium strategies it characterizes precisely mirror the data, which on every campus show significant numbers of rushees listing more than a single sorority on their preference cards in almost every year. Similarly, the equilibrium outcome has all rushees ultimately matched to the sorority they list first on their preference card, while Table 5 shows that, of campuses C and D, this approximately characterizes only the situation on campus D. However the theorem shows how the striking regularities observed in our data can arise at equilibrium. It shows how stage 2 of the formal rush procedure plays a much less important role than does the continuous open bidding which follows formal rush. Most importantly, it makes clear why the presence of unconstrained sororities may be expected to give so many rushees an incentive to list only a single sorority. Even when some sororities have tighter constraints, this incentive persists, since for example a rushee whose first-choice sorority is constrained but whose second-choice sorority is unconstrained also has no incentive to list more than her first choice when the strategies are as described.

Theorem 4 and its proof also suggest why an increase in the number of rushees who list only one choice on their preference cards will reduce the probability that the PBS

¹³If the constraints on sororities were completely relaxed (e.g., if the quotas q_k all exceeded the number of rushees), then the equilibrium in Theorem 4 would be perfect, and μ_R would be the unique stable match-

ing. However such a relaxed constraint does not describe what we observed. Similarly, we could have modeled open bidding by rules that would lead to stable outcomes at every perfect equilibrium regardless of the quotas (see Roth, 1984b; Roth and Sotomayor, 1990), but this would involve imposing particular detailed rules beyond what we observe.

algorithm will fail, even if some rushees behave differently. That is, increasing the number of rushees who submit a single choice on their preference cards may remove the cause of failure of the PBS algorithm but may never cause failure. The following proposition, stated without proof, formalizes this.

PROPOSITION: *Let P be a collection of stated preferences for a set S of sororities and a set R of rushees, and let P' be a collection that differs from P only in that some of the preference orderings in P have been truncated after their first element. Then the PBS algorithm with input P' will never fail if the PBS algorithm with input P does not.*

VII. Modeling Issues and Open Questions

Many choices must be made in modeling a complex system. We have already pointed out some of the modeling decisions we have made. Here, we discuss some aspects of sorority rush that we have not included in the formal analysis. Our motivation for discussing these explicitly is that, if such choices are not made carefully, the conclusions of the analysis may be misleading. Therefore, we want to explain briefly the reasons behind our choices. Then, we turn to some open questions.

First, although we have modeled sororities as being concerned with groups of rushees, we have modeled rushees as having preferences only over sororities and not over which other rushees join the same sorority. This seems justified both because particular sororities typically draw from the same part of the rushee population year after year (so preferences over sororities are a good proxy for preferences over other rushees) and because rushees are typically freshmen who have not yet had time to form many close friendships with other freshmen. Nevertheless, it is not unheard of for pairs of rushees, typically friends from high school, to wish to join the same sorority, and the problem facing such a pair differs from that analyzed here.

Second, our strategic analysis considered only the behavior of individual rushees and

sororities and not sorority-rushee coalitions. There may be an additional reason why some rushees list only a single sorority, since in some circumstances it may be in the interest of an unconstrained sorority to encourage certain rushees to do so, although this is regarded as one of the more serious violations of the rules. Briefly, certain rushees (called "legacies") may have close relations with a given sorority even before the beginning of rush, by virtue of having a family member who is a member or alumna of that sorority. If this rushee r lists only that sorority S on her preference card, then sorority S can plan to list rushee r somewhere on its second bid list, and can count on enrolling her during open bidding. This permits the sorority to rank higher on its bid list other rushees, who may list more than one sorority and who might therefore be matched during the PBS algorithm to another sorority if sorority S submitted its true preferences. We have been unable to gather data on how widespread this phenomenon might be, both because it is explicitly forbidden by the authorities concerned with rush and because it is not easily distinguishable from the simpler reasons for listing only a single sorority already described.¹⁴ That is, the reasons this might be a viable agreement between a sorority and a rushee are not substantially different from the reasons that individual rushees, acting on their own, might choose to submit a single preference.

Third, we have not analyzed the several rounds of parties described in Section II, which precede the submission of preferences by sororities and rushees. Our sense is that, on the campuses we have observed, because most sororities are unconstrained or only loosely constrained, the strategic considerations that arise in deciding which rushees to invite and which parties to accept have at most secondary importance, and the primary role of the parties on these

¹⁴ We are indebted to Patty Beeson for pointing out to us that some sororities have rules that any legacy who attends the final preference party must be listed on the first bid list.

campuses is to help rushees and sororities form their preferences and signal them to one another.

This brings us naturally to the next modeling issue. We have analyzed the game as a game of complete information, in which sororities and rushees know one another's preferences. Essentially we are assuming that, in the course of the preference parties, these preferences are fully communicated. This is obviously only an approximation of reality, but a rough idea of how adequate this approximation might be can be gotten from Tables 2-4, which show how many times a sorority was listed on the preference card of a rushee who did not appear anywhere on its bid list. Each such incident is likely to be a case in which the complete-information assumption is not met.¹⁵ The relatively low frequency of this suggests that the complete-information assumption is a rough approximation of what we observed. Note that, while the assumption of complete information is certainly less than fully satisfactory, serious new problems would arise in attempting to model the game as one of incomplete information, since the results of such an analysis would be sensitive to the assumptions that would have to be made about participants' prior probability distributions.¹⁶

Finally, our analysis has treated each sorority as an individual agent and not as a collection of individual members. Given that each sorority is required to submit a single

bid list, this may not raise problems, but we note that we have not investigated any aspect of how sororities arrive at their bid lists. In this regard, Roth and Sotomayor (1989) show there is a surprising coincidence of preferences over stable matchings among agents with different responsive preferences over groups, provided they have the same preference over individuals. Thus, for many purposes, the relevant differences among sorority members will be precisely those that go into determining the preferences over individuals (i.e., the bid list) and not more complex issues regarding the makeup of the entire entering group of new members.

We now consider briefly one of the major open empirical questions raised by this work: on campuses having mostly constrained sororities, how will rush differ from what we have observed on campuses with mostly unconstrained sororities? We conjecture that there will be at least two important (and related) differences. As we have seen, the high percentage of rushees listing only one sorority on their preference cards in formal rush is related to the fact that this (unconstrained) sorority can issue further invitations during open bidding.¹⁷ This will not be so on campuses in which most sororities cannot accept new members after the end of formal rush, and so on these campuses we expect to see a very much smaller percentage of single preferences.

This brings us to the second major difference we expect. Recall that the PBS algorithm as delineated in the literature of the National Panhellenic Council is incompletely specified: for some configurations of preferences, it does not indicate how some rushees should be dealt with. As we saw, the very low frequency of this kind of failure in our data can be attributed to the high percentage of rushees who submit single preferences. If the percentage of single preferences is much lower on campuses with

¹⁵Of the four campuses observed, only the sororities on campus C are required by their college panhellenic to list every rushee invited to the final preference party somewhere on their bid list. Thus, the invitations to the final round of parties are particularly effective signals of sorority preferences on campus C (but even on this campus, there was one case of an unlisted rushee in 1986; see Table 4).

¹⁶Although perhaps some headway could be made due to the fact that the preferences of rushees for sororities and those of sororities for rushees may follow certain identifiable patterns (e.g., on a given campus some sororities may be known as athletic, others as wealthy, etc.). A discussion of equilibria when agents have incomplete information about other agents' preferences is found in Roth (1989) and Roth and Sotomayor (1990).

¹⁷At least in theory. The differences between the various campuses in our data set (see the notes in Tables 2-4) preclude meaningful statistical comparisons across campuses (see Mongell, 1988).

constrained sororities, it seems likely that the frequency of failure (i.e., the percentage of rushes in which the submitted preferences fall outside of those that can be fully processed by the standard instructions for the PBS algorithm) will be higher. (Indeed, the one failure that occurred in our data was on campus D in 1987, the first year some of the sororities on that campus became constrained.) This suggests that supplementary rules will be adopted on these campuses to determine what the algorithm should do in such cases. Since it appears that these rules will have to be developed separately on each campus, there may be more variation in the formal rush procedures found on such campuses (as well as in the strategic behavior of rushees and sororities).

Finally, on campuses with many constrained sororities, it seems likely that the initial rounds of preference parties would involve nontrivial strategic decisions. (For example, if you can only accept two final invitations, it might sometimes be advisable to decline an invitation from your first-choice sorority, in order to signal your interest to a lower-ranked choice which has a greater chance of giving you a high ranking on its preference list.) As described in Section II, a good deal of formal structure is involved in these parties, much of which is common from campus to campus, which would help in modeling the strategic aspects of these decisions.¹⁸

A second kind of open question is: what caused the unraveling of recruitment dates that occurred in this market prior to the introduction of the PBS, and what caused this unraveling to stop? Sorority rush may not be the two-sided matching market that will best illuminate these issues,¹⁹ but be-

cause this phenomenon occurs in other two-sided matching markets, the unraveling observed in sorority rush appears to be an example of a much more general phenomenon (see Roth 1984a, 1991).

VIII. Conclusions

One way to summarize this story is to say it is about the difficulty of central planning. The agreement hammered out in the 1920's among sororities to reduce the competitiveness of recruitment implemented a plan designed to give all sororities the ability to recruit the same number q of new members, before any sorority had a chance to recruit more. When sororities have capacity greater than q , which is the case on the campuses we have observed, this could in general lead to unstable outcomes (Theorem 3). That is, there will in general be rushees and sororities who share an incentive to circumvent this constraint. As we have seen, the agents in the market have adapted their behavior to do so: rushees list only their first-choice sororities in the part of the procedure constrained by the quota q , and sororities approach desirable rushees after this constraint has been lifted. Ironically, this adaptation contributes to the smooth operation of what would otherwise be an incompletely specified procedure (Theorem 4).

From a more general perspective, this study reaffirms the importance of examining systems of rules from the point of view of what incentives they give participants. If we had looked only at the formal rules of the PBS algorithm and analyzed it as if agents all submitted their true preferences, we would not have been able to explain what we were seeing. However, the data on submitted preferences clearly made it unlikely that rushees were submitting their true full

¹⁸For example, of the campuses we observed, all but campus B permit each rushee to attend only two final parties during the last round of the invitational parties. (Campus B permits rushees to attend up to three final preference parties.)

¹⁹Since sororities are subject to some sanctions (both from the national organization and from campus authorities), they may be able to enforce an agreement on

recruiting behavior simply, once it has been reached. Also, with the increased mobility of college students, there may not be much room to unravel recruiting much before the beginning of the freshman year (i.e., the unraveling may have stopped only because it has no further to go).

preferences, and the subsequent analysis of their incentives suggested that they were responding rationally to the incentives of the system.

We reemphasize that we study sororities not merely because of their intrinsic interest, nor merely to show that game-theoretic analysis can shed light on behavior that might not always be thought of as economic. The point of studying particular matching mechanisms is that they add to our understanding of how centralized matching mechanisms work in general (and there seems to be a surprising number of these). By understanding how centralized mechanisms work in practice, we can also hope to learn things that will be useful in the study of decentralized two-sided matching markets (see Roth and John H. Vande Vate, 1990, 1991). The advantage of beginning with centralized markets is that it is easier to determine when they reach stable outcomes and when they do not. The theoretical progress in studying labor and other markets as two-sided matching models (see references in Roth and Sotomayor [1990]) suggests that this kind of empirical research may be fruitful. The conclusions of the present study should lend further weight to the hypothesis that the stability or instability of the matchings that result from such a market are crucial to understanding the market's evolution.

If game theory is to become as important a part of empirical economics as it has become a part of economic theory, we must explore the kinds of empirical research that will allow us to test and refine game theory. Since game theory is the part of economic theory most particularly concerned with the "rules of the game," this will inevitably involve looking into the particular rules by which markets are organized. At the same time, as was emphasized in Roth (1991), notions like the stability of outcomes, which can be formulated somewhat independently of the particular rules of the game, are useful in comparing different sets of rules and in indicating when agents may have an incentive to change the rules or circumvent them.

APPENDIX

The following results from the literature will be of use.²⁰

THEOREM A1: *A stable matching exists for every matching market.*

Definition: For a given matching market (S, R, P) , a stable matching μ is S-optimal if every sorority likes it at least as well as any other stable matching. Similarly, a stable matching ν is R-optimal if every rushee likes it at least as well as any other stable matching.

Definition: A sorority S and a rushee r are achievable for each other in a matching market (S, R, P) if S and r are paired at some stable matching.

In a matching market in which all rushees and sororities have strict preferences over individuals and in which sororities have responsive preferences over groups, each sorority and rushee can strictly order its or her achievable mates, and so there can be at most one S-optimal stable matching and one R-optimal stable matching.

THEOREM A2: *When all sororities have strict preferences over individual rushees and all rushees have strict preferences over sororities, there always exists an S-optimal stable matching, μ_S , and an R-optimal stable matching, μ_R .*

THEOREM A3: *When all agents have strict preferences, the S-optimal stable matching is the worst stable matching for all the rushees; similarly, the R-optimal stable matching is the worst for all the sororities.*

²⁰Theorems A1 and A2 were proved in Gale and Shapley (1962), and Theorem A3 was proved in Donald Knuth (1976) for the marriage model. They hold not only for the marriage model, but also for the college admissions model considered here (see Roth and Sotomayor, 1990). Theorem A4 is from Roth (1984a); for a stronger result motivated by the American medical labor market, see Roth (1986).

We will also use the following result.

THEOREM A4: *When all preferences over individuals are strict, the set of rushees matched to sororities, and sorority positions filled, is the same at every stable matching.*

Theorem 1 follows immediately from the following proposition, which will also be useful in the proof of the next theorem.

PROPOSITION 1: *If r_i is not in hold when the algorithm stops, then any sorority S_k that r_i prefers to $x(r_i)$ does not prefer r_i to any element of $x(S_k)$.*

PROOF:

Consider a rushee r_i who is not in hold when the algorithm stops. Let t be the step at which this rushee is given an assignment (i.e., either matched to a sorority or left unmatched). Then $x'(r_i) = S_i$ for some sorority S_i , or $x'(r_i) = r_i$. In the former case (box A in Fig. 1), S_i is at the head of rushee r_i 's preference card at step t . (Note that such an assignment must be individually rational for both r_i and S_i .) In the latter case (box E or F), either the rushee's current preference card is empty (box F) or it contains only a sorority S_j that did not list rushee r_i (box E).

There are two points in the algorithm at which a sorority S_k can be deleted from the rushee r_i 's preference card: these are boxes C and D in Figure 1. If the deletion occurs at box C, the sorority has not listed rushee r_i on its bid list and so does not prefer rushee r_i to any element of $x(S_k)$. If the deletion occurs at box D, then the sorority has filled its quota by matching to q rushees at the top of its bid list during some step $k < t$. Since rushees are deleted from the bid lists of sororities only at boxes A and B, rushee r_i has not been deleted from sorority S_k 's bid list prior to step k , so S_k prefers all the rushees in $x(S_k)$ to r_i .

The next proposition states that, even when the PBS algorithm fails to assign all rushees, the resulting partial matching could be extended to a stable matching in the market with quota q .

PROPOSITION 2: *There exists a stable matching μ in the market with quota q such that $\mu(r) = x(r)$ for every rushee r who is matched by the PBS algorithm.*

PROOF:

Consider the "residual matching market" which arises in the market with quota q after the PBS algorithm has ended in failure, with some rushees left in hold. Define the agents in this market to consist of those rushees left in hold, together with the set of sororities. The preferences of the sororities are as in their bid lists in the original market, except that all rushees who have been matched by the PBS algorithm are deleted. The residual quota of a sorority S_k in the residual market is given by $q'_k \equiv q - |x(S_k)|$. That is, in the residual matching problem we have just defined, each sorority may fill no more positions than were left unfilled by the PBS algorithm. (The preferences of the rushees in this residual matching problem can be thought of either as the same as given by their preference cards in the original problem or as having deleted those sororities that now have a quota of 0.)

Let x' be a stable matching for the residual matching problem. Let μ be the matching for the original matching problem with quota q , such that for each rushee matched by the PBS algorithm, $\mu(r) = x(r)$ and for each rushee left in hold, $\mu(r) = x'(r)$. We will show that μ is stable in the market with quota q .

Suppose not. Then, since μ is individually rational, there is a sorority rushee pair (S, r) such that r prefers S to $\mu(r)$, and S prefers r to some σ in $\mu(S)$. It cannot be that r was assigned by the PBS algorithm, since if that were the case we would have $\mu(r) = x(r)$, which implies that, at the step of the algorithm at which r was matched to $x(r)$, sorority S had already filled its quota, so $\mu(S) = x(S)$, and by Proposition 1, (S, r) is not a blocking pair. Therefore, r was left in hold at the end of the PBS algorithm. Then, σ cannot be equal either to S (an unmatched position) or to a rushee r' in the residual matching problem, since this would contradict the stability of x' in that market. Therefore, σ must equal some rushee r'

with $x(r') = S$. Note that r cannot have listed only one sorority on her preference card, and so the fact that r' was matched to S while r was still unmatched implies $|x(S)| < q$ (since rushees can be deleted from sorority bid lists only at points A and B in the flow chart in Fig. 1). Therefore, r is one of the q' highest rushees in S 's preferences in the residual market. Thus, x' cannot be stable in the residual market, which is the contradiction needed to conclude that no blocking pair (S, r) exists, and so μ must be stable.

LEMMA: *If S_k is r_i 's most preferred achievable match and r_i is among S_k 's q most preferred achievable matches, then (r_i, S_k) is an inevitable match in the market with quota q .*

PROOF:

Theorem A3 implies that S_k is both the best and the worst achievable sorority for r_i . Since preferences are strict, it is the only achievable match.

PROOF OF THEOREM 2:

The proof is by induction. Bear in mind that, by Proposition 2, if r_i is matched to S_p at some step of the algorithm, r_i and S_p are achievable for one another.

If, when $t = 1$, $r_i \in Q_1(S_j)$ and if, of the sororities that listed r_i , S_j is the highest-ranked sorority on r_i 's preference card, then $x^1(r_i) = S_j$, and this first match is inevitable. (In what follows, the ranking on the rushee's preference card refers to the original ranking when $t = 0$.) Thus, if a match is made at $t = 1$, it is inevitable. Now suppose that all matches up to time $t = k$ are inevitable; it will be shown that any match made at $t = k + 1$ is inevitable.

Consider the following cases:

- (i) No match occurs at step $k + 1$.

Then all matches up to step $k + 1$ are inevitable by the inductive assumption.

- (ii) A match $x^{k+1}(r_j) = S_p$ occurs at step $k + 1$ (i.e., rushee r_j is matched to sorority S_p at that step).

Then it must be that $r_j \in Q_{k+1}(S_p)$ and that, of the sororities that listed r_j and have not reached quota, S_p is the highest-ranked sorority on r_j 's preference card. First suppose that S_p is the highest-ranked sorority on r_j 's (original) preference card. Then the fact that $x^{k+1}(r_j) = S_p$ implies that S_p has not reached quota prior to step $k + 1$. (a) If $r_j \in Q_0(S_p)$, then (r_j, S_p) is inevitable (by the Lemma). (b) If $r_j \notin Q_0(S_p)$ but $r_j \in Q_0^+(S_p)$, then there exist other rushees that S_p prefers to r_j . Since $r_j \in Q_{k+1}(S_p)$, this implies that all of these rushees have either suicided or matched elsewhere (since rushees are deleted from sorority bid lists only at boxes A and B of the PBS-algorithm flow chart in Fig. 1). Any rushee who suicided cannot be matched to S_p at any individually rational outcome. Any rushee who was matched to some sorority other than S_p prior to step $k + 1$ is matched to that sorority at any stable outcome, by the inductive assumption. Therefore, none of these rushees can be matched to S_p at any stable outcome, and (r_j, S_p) is inevitable, by the Lemma.

Now suppose S_p is not the highest-ranked sorority on r_j 's (original) preference card of those that listed her somewhere on their bid lists. Let S_n be a sorority that listed r_j and was ranked before S_p on her preference card. Then the fact that $x^{k+1}(r_j) = S_p$ implies that S_n has reached quota at some step $t < k + 1$ (since sororities are deleted from rushee bid lists only at boxes C and D of the flow chart in Fig. 1). By the inductive hypothesis, S_n is not achievable for rushee r_j . Therefore, by the Lemma, r_j and S_p are inevitable.

- (iii) A match $x^{k+1}(r_j) = r_j$ occurs at step $k + 1$ (i.e., rushee r_j is assigned to be unmatched at that step).

Proposition 2 and Theorem A4 together imply in this case that $\mu(r) = r$ at every stable matching.

REFERENCES

- Brown, James T., ed., *Baird's Manual of American College Fraternities*, 9th Ed.,

- New York: James T. Brown, 1920.
- Gale, David and Shapley, Lloyd, "College Admissions and the Stability of Marriage," *American Mathematical Monthly*, 1962, 69, 9-15.
- Knuth, Donald E., *Marriages Stables*, Montreal: Les Presses de l'Université de Montréal, 1976.
- Mongell, Susan, *Sorority Rush as a Two-Sided Matching Mechanism: A Game-Theoretic Analysis*, Ph.D. Diss., Department of Economics, University of Pittsburgh, 1988.
- Roth, Alvin E., (1984a) "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory," *Journal of Political Economy*, December 1984, 92, 991-1016.
- _____, (1984b) "Misrepresentation and Stability in the Marriage Problem," *Journal of Economic Theory*, December 1984, 34, 383-7.
- _____, (1985a) "The College Admissions Problem Is Not Equivalent to the Marriage Problem," *Journal of Economic Theory*, August 1985, 36, 277-88.
- _____, (1985b) "Common and Conflicting Interests in Two-Sided Matching Markets," *European Economic Review* (special issue: Market Competition, Conflict, and Collusion), February 1985, 27, 75-96.
- _____, "On the Allocation of Residents to Rural Hospitals: A General Property of Two-Sided Matching Markets," *Econometrica*, March 1986, 54, 425-7.
- _____, "Two Sided Matching with Incomplete Information about Others' Preferences," *Games and Economic Behavior*, June 1989, 1, 191-209.
- _____, "A Natural Experiment in the Organization of Entry-Level Labor Markets: Regional Markets for New Physicians and Surgeons in the United Kingdom," *American Economic Review*, June 1991, 81, 415-40.
- _____, and Sotomayor, Marilda, "The College Admissions Problem Revisited," *Econometrica*, May 1989, 57, 559-70.
- _____, and _____, *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Econometric Society Monographs, Cambridge: Cambridge University Press, 1990.
- _____, and Vande Vate, John H., "Random Paths to Stability in Two-Sided Matching," *Econometrica*, 1990, 58, 1475-80.
- _____, and _____, "Incentives in Two-Sided Matching with Random Stable Mechanisms," *Economic Theory*, 1991, 1, 31-44.
- Shepardson, Francis W., ed., *Baird's Manual of American College Fraternities*, 12th Ed., Menasha, WI: Collegiate Press, 1930.
- National Panhellenic Conference, "How To" for College Panhellenics, 10th Ed., unpublished internal manual, 1979.
- _____, NPC: An Historical Record of Achievement, Maury Boyd, 1983 (unpublished facsimile of 1957 version published by Leland Publishers, St. Paul, MN).
- National Panhellenic Review, unpublished minutes, 1985.

Copyright of American Economic Review is the property of American Economic Association and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.