

THE ECONOMIST AS ENGINEER: GAME THEORY,
EXPERIMENTATION, AND COMPUTATION
AS TOOLS FOR DESIGN ECONOMICS¹

BY ALVIN E. ROTH²

Economists have lately been called upon not only to analyze markets, but to design them. Market design involves a responsibility for detail, a need to deal with all of a market's complications, not just its principle features. Designers therefore cannot work only with the simple conceptual models used for theoretical insights into the general working of markets. Instead, market design calls for an engineering approach. Drawing primarily on the design of the entry level labor market for American doctors (the National Resident Matching Program), and of the auctions of radio spectrum conducted by the Federal Communications Commission, this paper makes the case that experimental and computational economics are natural complements to game theory in the work of design. The paper also argues that some of the challenges facing both markets involve dealing with related kinds of complementarities, and that this suggests an agenda for future theoretical research.

KEYWORDS: Market design, game theory, experimental economics, computational economics.

1. INTRODUCTION: DESIGN ECONOMICS

THE ECONOMIC ENVIRONMENT EVOLVES, but it is also designed. Entrepreneurs and managers, legislators and regulators, lawyers and judges, all get involved in the design of economic institutions. But in the 1990's, economists, particularly game theorists, started to take a very substantial role in design, especially in the design of markets. These developments suggest the shape of an emerging discipline of *design economics*, the part of economics intended to further the design and maintenance of markets and other economic institutions.³

¹ This paper is dedicated to Bob Wilson, the Dean of Design.

² Fisher-Schultz Lecture. This paper has a checkered history. Versions of it were delivered as the Leatherbee Lecture at the Harvard Business School on April 21, 1999, as the Fisher-Schultz Lecture at the European meeting of the Econometric Society in Santiago de Compostela, on August 31, 1999, and as the Pazner Lecture at the University of Tel Aviv, on February 29, 2000. I have also struggled with some of these issues in my Clarendon Lectures at Oxford in April 1998, in Roth and Sotomayor (1990), Roth (1991), Roth and Peranson (1999), Roth (2000), and in the market design sections of my webpage at <http://www.economics.harvard.edu/~aroth/alroth.html>, and in the class on market design I co-taught with Paul Milgrom at Harvard in the Spring of 2001. This paper has benefited from the many helpful comments I received on those earlier occasions, and from sympathetic readings by my colleagues and the editor and referees.

³ Economists have also become involved in the design of economic environments such as incentive systems within firms, negotiating platforms, contracts, etc., but in the present article I shall draw my examples from market design.

Game theory, the part of economics that studies the “rules of the game,” provides a framework with which to address design. But design involves a responsibility for detail; this creates a need to deal with complications. Dealing with complications requires not only careful attention to the institutional details of a particular market, it also requires new tools, to supplement the traditional analytical toolbox of the theorist. The first thesis of this paper is that, *in the service of design, experimental and computational economics are natural complements to game theory.*

Another kind of challenge is professional rather than technical, and has to do with whether and how it will become customary to report design efforts in the economics literature. The recent prominence of economists in design arose from several events, external to the profession, that created a need for designs of unusual markets. Not only did this give economists a chance to employ what we know, it gave us a chance to learn practical lessons about design. Whether economists will often be in a position to give highly practical advice on design depends in part on whether we report what we learn, and what we do, in sufficient detail to allow scientific knowledge about design to accumulate. The second theme of the present paper is that, for this purpose, *we need to foster a still unfamiliar kind of design literature in economics, whose focus will be different than traditional game theory and theoretical mechanism design.*⁴

If the literature of design economics does mature in this way, it will also help shape and enrich the underlying economic theory. The third goal of the present paper will be to show how *recent work on design has posed some new questions for economic theory, and started to suggest some answers.*

To see how these issues hang together, it may help to consider briefly the relationship between physics and engineering, and between biology, medicine, and surgery.

Consider the design of suspension bridges. The simple theoretical model in which the only force is gravity, and beams are perfectly rigid, is elegant and general. But bridge design also concerns metallurgy and soil mechanics, and the sideways forces of water and wind. Many questions concerning these complications can't be answered analytically, but must be explored using physical or computational models. These complications, and how they interact with the parts of the physics captured by the simple model, are the domain of the engineering literature. Engineering is often less elegant than the simple underlying physics, but it allows bridges designed on the same basic model to be built longer and stronger over time, as the complexities and how to deal with them become better understood.

⁴ In this connection, Baldwin and Bhattacharyya (1991), speaking of the 1984 auction of Conrail's assets and of asset sales in general, write “the literature has yet to explain the diverse methods of sale we observe, and provides little guidance in choosing among alternatives. Furthermore, those who actually make the selection almost never discuss the experience, and thus, from a practical standpoint, we know very little about why one method is chosen over another.” (For connections between engineering design and the design of the firms and industries that produce them, see Baldwin and Clark (2000).)

It was not a foregone conclusion that bridge building would have a scientific component; the earliest bridges were no doubt built without much formal analysis, and experience was accumulated only insofar as it could be passed from builder to apprentice, or learned from observation. But the close relationship of physics and engineering goes back at least as far as Archimedes.

Surgery and its relation to medicine and biology provide a more cautionary tale. The famous oath that Hippocrates formulated for physicians in the fifth century BC includes a specific prohibition against doing surgery. Two millennia later, in medieval Europe, surgery was still not considered the province of medical specialists. In England, responsibility for surgery was vested in the Worshipful Company of Barbers, which underwent a number of changes over the centuries and was not fully separated from surgery until 1745.⁵ This was not an unreasonable assignment of responsibilities, as barbering required expertise in keeping blades sharp and in stanching the flow of blood, both also essential for surgery. But the modern connections between surgery and medicine, and between medicine and biology, have certainly improved the quality of the surgery, and also of medicine and biology.

The analogy to market design in the age of the internet should be clear. As marketplaces proliferate on the web, a great deal of market design is going to be done by computer programmers, among others, since they possess some of the essential expertise. Economists will have an opportunity to learn a lot from the markets that result, just as we will learn from our own work designing unusual markets. But if we want this knowledge to accumulate, if we want market design to be better informed and more reliable in the future, we need to promote a scientific literature of design economics. Today this literature is in its infancy. My guess is that if we nurture it to maturity, its relation with current economics will be something like the relationship of engineering and physics, or of medicine and biology.

1.1. *Design in the 1990's:*⁶

The 1990's were a formative decade for design economics because economists were presented with opportunities to take responsibility for the design of detailed

⁵ As of this writing, the headquarters of the barbers' organization in London is still called Barber-Surgeons' Hall. For a brief history, see "The history of the company," at <http://www.barberscompany.org/>.

⁶ Economists' interest in market design long predates the 90's; one has only to think of Vickery's 1961 article on auctions, or Gale and Shapley's 1962 article on matching mechanisms. The study of auctions blossomed into a substantial theoretical literature, including discussions of the design of revenue maximizing auctions under simple assumptions, and an empirically oriented literature studying the outcomes of particular auctions. (Notable early theoretical papers are Myerson (1981), Milgrom and Weber (1982b), Maskin and Riley (1984), Bulow and Roberts (1989); see Wilson (1992) for a broad overview of the auction literature.) The study of matching mechanisms also produced a substantial theoretical literature, accompanied by empirically oriented papers studying the evolution of labor markets, beginning with a study of the market for American physicians (Roth (1984)). (See Roth and Sotomayor (1990) for an overview of much of the theory, and Roth and Xing (1994) for a discussion

rules for complex markets, and their suggestions were quickly implemented in large, operating markets. This in turn produced an opportunity to evaluate the new designs. Two notable design efforts were:

- the design of labor clearinghouses such as the one through which American doctors get their first jobs; and
- the design of auctions through which the U.S. Federal Communications Commission sells the rights to transmit on different parts of the radio spectrum.

The importance of good design is nowhere better illustrated than in a third set of markets in which economists have played a role, but in which politicians and regulators also continue to be deeply involved, namely markets for electric power. (See Wilson (2002) for an account of the most detailed design work.) Economists participated in the design of only parts of these markets, while other parts remained subject to regulation. An unworkable hybrid resulted in California, where utility companies were brought to the verge of bankruptcy by the rising prices in the unregulated wholesale market, which far exceeded the regulated prices at which electricity could be sold to consumers.

While these markets are very different from each other, their design efforts had some striking similarities that distinguish them from more conventional work in economics.⁷

On a technical level, each of these markets presented designers with the problem of dealing with goods that could be *complements* for one another, not just substitutes. In labor markets, how desirable a particular job is for one member of a two-career couple may depend on how good a job the other member of the couple can get in the same city. Similar complementarities arise when employers have preferences over groups of workers. In spectrum auctions, a given bandwidth of spectrum may be more valuable to a company if the company can also

of the evolution of labor markets and other matching processes.) Mechanism design in general, in the spirit laid out in Hurwicz (1973) has become a recognized subject in the theoretical literature, and even boasts a specialized journal, the *Review of Economic Design*. The strategic issues associated with motivating market participants to reveal privately held information have been explored from a design orientation, as in the work of Groves and Ledyard (1977) on mechanisms to promote the efficient provision of public goods (see Green and Laffont (1979) for an early overview). And the interest in design is not only theoretical; experimenters have explored alternative designs in simple settings, often motivated by general concerns like public goods provision (see Ledyard (1995) for a survey), or by forms of market organization such as auctions (see Kagel (1995)), and sometimes by particular allocation problems, arising e.g. from airports or space stations (e.g. Rassenti, Smith, and Bulfin (1982), Grether, Isaac, and Plott (1981), and Ledyard, Porter, and Wessen (2000), JEE). All of this helped lay the groundwork for the assumption of design responsibilities that took place in the 90's.

⁷ A wider ranging essay might also include the design of the large scale privatizations of state assets that began in Eastern Europe and the former Soviet Union following the fall of the Berlin Wall in November 1989, although I am not sure to what extent economists played direct roles in the designs of the various privatization plans. One of the most interestingly designed privatizations was the (then) Czechoslovak privatization in 1992, meant to identify prices (in terms of vouchers issued to all citizens) at which the market would clear, via a multi-round auction in which prices for over- and under-demanded firms were readjusted from round to round. The results of that auction have been studied in several papers; see, e.g., Svejnar and Singer (1994), Filer and Hanousek (1999), and Gupta, Ham, and Svejnar (2000). See also the admirable discussion of creating markets for macroeconomic risks in Shiller (1993).

obtain geographically contiguous licenses, so that it can offer a broader area of service, or if it can win licenses for adjacent radio frequencies, so that it can transmit more data. And in electricity markets, power generation is quite distinct from power transmission, but both must be consumed together. And power can be generated more cheaply from an already operating generator than from a cold start, so there are complementarities over time, and between power generation and various “ancillary services” like reserve capacity needed to keep the transmission network operating. These complementarities play a strong role in the resulting market designs.

Aside from technical issues, these design efforts also shared some features that seem to be characteristic of the context in which markets are designed. First, design is often required to be *fast*. In each of these three cases, only about a year elapsed between the commissioning of a new market design and its delivery. Second, design need not be an entirely *a priori* craft; much can be learned from the *history* of related markets, and sometimes there is an opportunity and a need to tinker with new designs, based on early experience. Finally, at least some of the work of design reflects the fact that the adoption of a design is at least partly a *political* process.

In what follows, I will attempt to develop these themes. I will do so in particular in connection with the (re)design of the entry level labor market for American doctors. That design effort, which I led, is the focus of this paper. It was able to profitably employ theory, historical observation, experimentation, and computation. Section 3 will then consider, briefly, the design of the FCC auctions, emphasizing the points of similarity in the work of design. The most recent FCC designs, of auctions in which bids can be made for packages of items, are closely related to some of the issues that arise in the design of labor market matching mechanisms. Section 4 concludes with an overview, and an exhortation.

2. THE ENTRY-LEVEL LABOR MARKET FOR AMERICAN DOCTORS

The entry-level position for an American doctor is called a residency (and was formerly called an internship). A good residency substantially influences the career path of a young physician, and residents provide much of the labor force of hospitals, so this is an important market for both doctors and hospitals. In the 1940's, the fierce competition for people and positions led to a kind of market failure (to be described below) that turns out to also have occurred in quite a few entry-level professional labor markets (Roth and Xing (1994)). This market failure was resolved in the early 1950's by the organization of a very successful clearinghouse to facilitate the matching of doctors to residency programs (Roth (1984)). Today this clearinghouse is called the National Resident Matching Program (NRMP).

Over the years, the medical profession underwent profound changes, some of them with important consequences for the medical labor market. In the Fall of 1995, amidst a crisis of confidence in the market, I was retained by the Board of Directors of the NRMP to direct the design of a new clearinghouse

algorithm for the medical match. This design, reported in detail in Roth and Peranson (1999) was completed in 1996, and adopted in 1997 as the new NRMP algorithm. Since then, over 20,000 doctors a year have been matched to entry level positions in the general medical labor market using the new algorithm, as well as a smaller number of more senior positions in about thirty medical specialty labor markets also organized by the NRMP.

The Roth-Peranson design has also been adopted by entry level labor markets in other professions since its adoption by American physicians. (Other markets that have adopted it to date are, in the United States, Postdoctoral Dental Residencies, Osteopathic Internships, Osteopathic Orthopedic Surgery Residencies, Pharmacy Practice Residencies, and Clinical Psychology Internships,⁸ and, in Canada, Articling Positions with Law Firms in Ontario, Articling Positions with Law Firms in Alberta, and Medical Residencies.)⁹

To understand the design problem, we need to understand the kinds of market failure that sometimes lead to the adoption of clearinghouses, and the manner in which clearinghouses themselves succeed and fail. For this it will help to start with a brief history of the American medical market.

2.1. *A Brief History of the Market*

The medical internship, as it was then called, came into being around 1900, and soon became the entry level position for American physicians. The labor market for such positions was decentralized, with students typically seeking positions around the time they graduated from medical school. But competition among hospitals for good students gradually led to earlier and earlier dates of appointment, despite repeated attempts to halt the process. By the 1940's students were being appointed to jobs two full years before graduation from medical school, i.e. hiring was two years before the start of employment. This meant that students were being hired before much information about their medical school performance was available to potential employers, and before students themselves had much exposure to clinical medicine to help them decide on their own career preferences.¹⁰

A serious attempt to reform the market and eliminate this source of inefficiency in matching was made in 1945, when the medical schools banded together and agreed to embargo student transcripts and letters of reference until an agreed date. This effectively controlled the unravelling of appointment dates, and as this became clear, the dates at which contracts were to be signed was moved back into the senior year of medical school.

But new problems developed between the time offers were issued and the deadline before which they had to be accepted. Briefly, students who had been

⁸ See Roth and Xing (1997) for a description of the clinical psychology market in the years when it operated as a telephone market, before it adopted a centralized match.

⁹ A related mechanism is now used to match new Reform rabbis to their first congregation (Bodin and Pankin (1999)).

¹⁰ This kind of unraveling of transaction dates is a common problem; for descriptions of it in some contemporary markets see Avery, Jolls, Posner, and Roth (2001) on the market for appellate court law clerks, and Avery, Fairbanks, and Zeckhauser (2001) on early admissions to college.

offered a job at one hospital, but told that they were on the waiting list for a job that they preferred, wished to wait as long as possible before accepting or rejecting the offer they were holding. This meant that waiting lists did not progress very fast, and that there were lots of last minute decisions, with decisions made in haste often regretted and sometimes not honored.

After several years of unsuccessful tinkering with the rules governing offers and deadlines, the chaotic recontracting that was being experienced became intolerable. It was decided to organize a centralized clearinghouse, modeled on regional clearinghouses that already existed in Philadelphia and Boston.¹¹ Students and internship programs would arrange interviews in a decentralized way, as before, but instead of engaging in a chaotic process of exploding offers, acceptances, and rejections, students would submit a rank order list of the jobs for which they had interviewed, and employers (internship directors) would submit a rank order list of the students they had interviewed. An algorithm would be devised to process the lists submitted in this way and produce a recommended matching of students to hospitals.

After a brief trial of an algorithm that was replaced because it had unacceptable incentive properties, an algorithm was adopted that proved very successful. Starting in 1951, and lasting into the 1970's, over 95% of positions were filled through the match. Small changes were made in the clearinghouse rules from time to time. There was some dropoff in this percentage in the 1970's, most interestingly among the growing number of married couples, graduating together from medical school, who wished to find two positions in the same city. When I first studied the market (see Roth (1984)) there had been an attempt to accommodate married couples that had been largely unsuccessful. But a subsequent further change of rules allowing couples to rank pairs of positions successfully attracted couples to once again participate in the match.

In the mid 1990's, the market experienced a serious crisis of confidence. (The various national medical student organizations issued resolutions, and eventually Ralph Nader's organization Public Citizen got involved.) There was a great deal of talk about whether the market served students' interests, and whether students would be well advised to "game the system" or even to circumvent the market entirely. It was in this context that I was asked to direct the design of a new algorithm, and to compare different ways of organizing the match.

Now, if an economist is going to act as a doctor to a medical market, he can do worse than to consult medical authorities on the standard of care. Hippocrates' [circa 400BC] advice on this subject applies to design economists too:

"The physician must be able to tell the antecedents, know the present, and foretell the future—must mediate these things, and have two special objects in view with regard to disease, namely, to do good or to do no harm."

¹¹ I am indebted to Dr. Cody Webb for bringing these regional clearinghouses to my attention. For Philadelphia, see Hatfield (1935). For Boston, see Mullin and Stalaker (1952), who suggest (p. 200) that elements of the national clearinghouse may have been taken directly from the Boston plan.

For the medical marketplace, this means that before replacing a clearinghouse that had effectively halted the coordination failures that preceded its introduction in the 1950's, it is important to know what causes clearinghouses to be successful. Fortunately, this can be addressed empirically, as there are both successful and unsuccessful clearinghouses in various labor markets. To discuss these, we will first need a simple theoretical model. We start with a somewhat *too* simple model—think of this as the simple model of a bridge with perfectly rigid beams—and then we'll complicate it as needed when we talk about the complications of the medical market.

2.2. A (Too) Simple Model of Matching¹²

There are disjoint sets of firms and workers, $F = \{f_1, \dots, f_n\}$ and $W = \{w_1, \dots, w_p\}$. (I will refer interchangeably to firms and workers, and to hospitals and students, or applicants, when speaking of the medical market.) Each worker seeks one job, and each firm f_i seeks (up to) q_i workers. A matching is a subset of $F \times W$, i.e. a set of matched pairs, such that any worker appears in no more than one pair, and any firm f_i appears in no more than q_i pairs. A matching is identified with a correspondence $\mu: F \cup W \rightarrow F \cup W$ such that $\mu(w) = f$ and $w \in \mu(f)$ if and only if (f, w) is a matched pair; and if no matched pair contains w , then $\mu(w) = w$ (i.e. if a worker is not matched to a firm, then she is matched to herself).¹³

Each agent has complete and transitive preferences over the “acceptable” agents on the other side of the market, i.e. over those that it prefers rather than remaining unmatched (or leaving a position empty) and waiting for the less desirable post-match market, called “the scramble.”¹⁴ The preferences of firms and workers will be given as lists of the acceptable agents on the other side of the market. For example, the preferences of a worker w_i are given by $P(w_i) = f_2, f_4, \dots$, indicating that she prefers firm f_2 to f_4 [$f_2 >_{w_i} f_4$], etc.

Because firms may be matched to groups of workers, firm f_i 's preference list $P(f_i) = w_i, w_j, \dots, w_k$ of acceptable workers does not fully define its preferences over groups of workers it might employ. For this simple model, it is enough to specify that a firm's preferences for groups are “responsive” to their preferences over individual workers (Roth (1985)), in the sense that, for any set of workers $S \subset W$ with $|S| < q_i$, and any workers w and w' in W/S , $S \cup w >_{f_i} S \cup w'$ if and only if $w >_{f_i} w'$, and $S \cup w >_{f_i} S$ if and only if w is acceptable to f_i . That is, if

¹² This is essentially the model of Gale and Shapley (1962), as extended in Roth (1985) for the case when firms employ multiple workers and hence have preferences over groups of workers.

¹³ It is more usual to define a matching directly in terms of the correspondence, but it was easier to describe the couples algorithm using this formulation, adapted from Blum, Roth, and Rothblum (1997).

¹⁴ By modeling workers' preferences as over firms rather than over positions, we are taking all positions offered by the same firm to be identical or (in practice) treating the different positions that might be offered by a firm as being offered by different sub-firms. Also, wages and other personal terms are set prior to the match, and are reflected in the preferences.

a firm prefers worker w to w' , this means that the firm would always prefer to add w instead of w' , to any group of other workers, and always prefers to add an acceptable worker when a space is available.

A matching μ is *blocked by an individual k* if $\mu(k)$ is unacceptable to k , and it is *blocked by a pair of agents (f, w)* if they each prefer each other to their mates at μ , i.e. if

$$[w >_f w' \text{ for some } w' \text{ in } \mu(f) \text{ or } w \text{ is acceptable to } f \text{ and } |\mu(f)| < q_f],$$

and

$$f >_w \mu(w).$$

A matching x is *stable* if it isn't blocked by any individual or pair of agents. When preferences are responsive, the set of stable matchings equals the core (defined by weak domination) of the game whose rules are that any worker and firm may be matched, if and only if they both agree.

Gale and Shapley (1962) showed that the set of stable matchings in this simple model is never empty, by observing that algorithms like the one below always produce a stable outcome. For our present purpose, the algorithm should be read as a way to process preference lists that workers and firms have submitted to a centralized clearinghouse. But the steps of the algorithm are written as if workers and firms were going through a decentralized process of application and eventual acceptance or rejection. This should help make clear why algorithms of this sort have been independently invented over the years in a number of markets. (It was shown in Roth (1984) that the algorithm adopted for the medical match in 1951 is equivalent to an algorithm like the one below, but with offers of positions being initiated by the hospitals rather than applications for positions being initiated by the students.)

Deferred Acceptance Algorithm, with workers applying (roughly the Gale-Shapley (1962) version):

- 1.a. Each worker applies to his or her first choice firm.
- 1.b. Each firm f (with q_f positions) rejects any unacceptable applications and, if more than q_f acceptable applications are received, "holds" the q_f most preferred, and rejects the rest.
- ⋮
- k.a. Any worker whose application was rejected at step $k - 1$ makes a new application to its most preferred acceptable firm that hasn't yet rejected it (i.e. to which it hasn't yet applied).
- k.b. Each firm f holds its (up to) q_f most preferred acceptable applications to date, and rejects the rest.

STOP when no further applications are made, and match each firm to the applicants whose offers it is holding.

Call the matching that results from this worker-proposing algorithm μ_W . To see that it is stable, note first that no unacceptable matches are even temporarily held,

so the only possible cause of instability would be a blocking pair. But suppose worker w prefers firm f to her outcome $\mu_w(w)$. Then she must have applied to firm f before the final step of the algorithm, and been rejected. Hence firm f does not prefer worker w to (any of) its workers, and so (w, f) is not a blocking pair.

This proves the first of the following theorems, all of which apply to the simple market modeled above. We will see that the theoretical picture changes when we consider the complexities of the medical market.

Theorems Concerning Simple Matching Markets

THEOREM 1: *The set of stable matchings is always nonempty (Gale and Shapley (1962)).*

THEOREM 2: *The deferred acceptance algorithm with workers proposing produces a “worker optimal” stable match, that matches each worker to the most preferred firm to which she can be matched at a stable matching. The parallel “firm proposing” algorithm produces a “firm optimal” stable matching that gives to each firm f_i the (up to) q_i most preferred workers to which it can be matched at a stable matching. The optimal stable matching for one side of the market is the least preferred stable matching for the other side (Gale and Shapley (1962); Roth and Sotomayor (1989)).*

THEOREM 3: *The same applicants are matched and the same positions are filled at every stable matching. Furthermore, a firm that does not fill all its positions at some stable matching will be matched to the same applicants at every stable matching (McVitie and Wilson (1970); Roth (1984, 1986)).*

THEOREM 4: *When the worker proposing algorithm is used, but not when the firm proposing algorithm is used, it is a dominant strategy for each worker to state her true preferences. (There exists no algorithm that always produces a stable matching in terms of the stated preferences and that makes it a dominant strategy for all agents to state their true preferences.) Furthermore, when the firm proposing algorithm is used, the only applicants who can do better than to submit their true preferences are those who would have received a different match from the applicant proposing algorithm (Roth (1982, 1985), Roth and Sotomayor (1990)).*

2.3. *The Importance of Stability*

The theoretical motivation for concentrating on the set of stable outcomes is that, if the market outcome is unstable, there is an agent or pair of agents who have the incentive to circumvent the match. Even in a large market, it is not hard to ascertain if an outcome is unstable, because the market can do a lot of parallel processing. Consider a worker, for example, who has received an offer

TABLE I
STABLE AND UNSTABLE (CENTRALIZED) MECHANISMS

Market	Stable	Still in use (halted unraveling)
American medical markets		
NRMP	yes	yes (new design in '98)
Medical Specialties	yes	yes (about 30 markets)
British Regional Medical Markets		
Edinburgh ('69)	yes	yes
Cardiff	yes	yes
Birmingham	no	no
Edinburgh ('67)	no	no
Newcastle	no	no
Sheffield	no	no
Cambridge	no	yes
London Hospital	no	yes
Other healthcare markets		
Dental Residencies	yes	yes
Osteopaths (<'94)	no	no
Osteopaths (≥'94)	yes	yes
Pharmacists	yes	yes
Other markets and matching processes		
Canadian Lawyers	yes	yes (except in British Columbia since 1996)
Sororities	yes (at equilibrium)	yes

from her third choice firm. She only needs to make two phone calls to find out if she is part of a blocking pair.

The empirical evidence offers a good deal of support to this intuition. Table I lists a number of markets that have at one point in their history adopted centralized clearinghouses (see Roth (1990, 1991), Roth and Xing (1994, 1997), and Mongell and Roth (1991)). In addition, it indicates whether they produce matchings that are stable with respect to the submitted preferences. (The question of whether they are stable with respect to the actual preferences will be discussed below.) The table further lists whether these clearinghouses were successful (at halting unraveling) and are still in use, or whether they have failed and were abandoned.

The table suggests that producing a stable matching is an important criterion for a successful clearinghouse. Stable mechanisms have mostly (but not always) succeeded, and unstable mechanisms have mostly (but not always) failed. The situation is complicated by the many differences among the markets in the table other than the stability or instability of their clearinghouse algorithm. The set of markets that come closest to providing a crisp comparison are the different regional markets for new physicians and surgeons in Britain (Roth (1990, 1991)). Of these, the two that employ stable mechanisms have succeeded, while all but two of those that do not employ stable mechanisms have failed. But even here, there are differences between the markets—e.g. differences between Newcastle

and Edinburgh—other than the organization of their clearinghouses. It could therefore be possible that the success of a stable clearinghouse in Edinburgh and the failure of an unstable one in Newcastle were for reasons other than how these clearinghouses were designed.

2.3.1. *Experimental Evidence*

Laboratory experiments can help clarify the impact of different clearinghouse designs. In a controlled environment, we can examine the effect of different matching algorithms while holding everything else constant. In this spirit, Kagel and Roth (2000) report an experiment that compares the stable, deferred acceptance market mechanisms used in Edinburgh and Cardiff with the kind of unstable, “priority matching” mechanism used in Birmingham, Newcastle, and Sheffield. Unver (2000b, c) reports a follow-up experiment that additionally compares these two classes of algorithms with the linear programming based algorithms used in Cambridge and at the London Hospital.

The successful algorithms adopted first in Edinburgh (in 1969) and then in Cardiff are essentially firm-proposing deferred acceptance algorithms. An alternative kind of algorithm that was widely tried in England proper (in Birmingham, Newcastle, and Sheffield), but always soon abandoned, defined a “priority” for each firm-worker pair as a function of their mutual rankings. Such an algorithm matches all priority 1 couples and removes them from the market, then repeats for priority 2 matches, priority 3 matches, etc. For example, in Newcastle, the priorities for firm-worker rankings were organized by the *product* of the rankings, (initially) as follows: 1-1, 2-1, 1-2, 1-3, 3-1, 4-1, 2-2, 1-4, 5-1. . . . That is, the first priority is to match firms and workers who each list one another as their first choice, the second priority is to match a firm and worker such that the firm gets its second choice while the worker gets his first choice, etc.

This can produce unstable matchings—e.g. if a desirable firm and worker rank each other fourth, they will have such a low priority ($4 \times 4 = 16$) that if they fail to match to one of their first three choices, it is unlikely that they will match to each other. (E.g. the firm might match to its fifteenth choice worker, if that worker has ranked it first.)

The Kagel and Roth experiment created a simple laboratory environment in which subjects would initially gain experience with a decentralized matching market with sufficient competition and congestion to promote unraveling of appointments.¹⁵ The subjects would have the opportunity to make early matches, but at a cost. Once subjects had time to adapt to this market, one of the two centralized matching mechanisms would be made available for those subjects who did not make early matches. The only difference between the two conditions of the experiment was that one employed the priority matching algorithm used unsuccessfully in Newcastle, and the other the stable matching algorithm used successfully in Edinburgh. The idea was not to reproduce the Edinburgh and Newcastle

¹⁵ The role of congestion was clearly seen in the unusually fast, but nevertheless congested telephone market operated by clinical psychologists, prior to their adoption of a centralized clearinghouse (Roth and Xing (1997)).

markets, but rather to see whether the algorithms employed in those labor clearinghouses had the same effect in a simple environment in which any difference could be unambiguously interpreted as being due to the algorithm.

Each experimental market consisted of 12 subjects, half of whom were assigned the role of firms and half of workers. Subjects were paid according to with whom they matched, and when (with payoffs for a successful match ranging from \$4 to \$16 per market). Each matching market consisted of three periods, called (in order) periods -2 , -1 , and 0 , to denote that the profit to a worker or firm from any match was reduced by \$2 if the match was made in period -2 , by \$1 if the match was made in period -1 , and by \$0 if the match was made in (the final) period 0 . A firm or worker who had not matched by the end of period 0 earned nothing.

In each matching market, firms could hire one worker, and likewise each worker could accept only one job. Firms were restricted to one offer in each period. Workers who received multiple offers in any period could accept at most one. Contracts were binding; a firm whose offer was accepted, and a worker who accepted an offer could not make other matches in a later period.

Each experimental session began with 10 consecutive decentralized matching markets. After the tenth decentralized matching market it was announced that, in each subsequent market, periods -2 and -1 would proceed as before, but henceforth period 0 would employ a centralized matching algorithm. For the centralized matching algorithm, subjects were instructed that if they were not matched prior to round 0 , they were to “submit a rank order list of their possible matches, and the centralized matching mechanism (which is a computer program) will determine the final matching, based on the submitted rank order lists.” In each experimental session we then conducted an additional 15 matching markets, with periods -2 and -1 employing the same decentralized procedures described above, but with one of the two centralized matching algorithms in place in period 0 . Half of the experimental sessions employed the stable, Edinburgh algorithm, and half employed the unstable priority matching algorithm used in Newcastle.

An experiment like this yields a rich set of data, but for our present purposes, one figure, of aggregate behavior, will suffice. Figure 1 shows the average costs paid by all subjects in a given market for making early matches, over time. (If all 12 subjects made matches at period -2 this cost would be \$24, if all matched at period -1 it would be \$12, while if no subjects made early matches this cost would be 0.) The figure shows that after the subjects have gained some experience with the decentralized market, there is a good deal of actual early matching (reflecting a great deal of attempted early matching; the realized cost is around \$8 in periods 6–10.) In periods 11–15, i.e. the first 5 periods after a centralized clearinghouse has been introduced, the amount of early matching is about the same for both mechanisms. But by periods 21–25, the Newcastle algorithm has not reduced early contracts, while the stable algorithm has. Thus the experimental results qualitatively reproduce the field results, under conditions in which

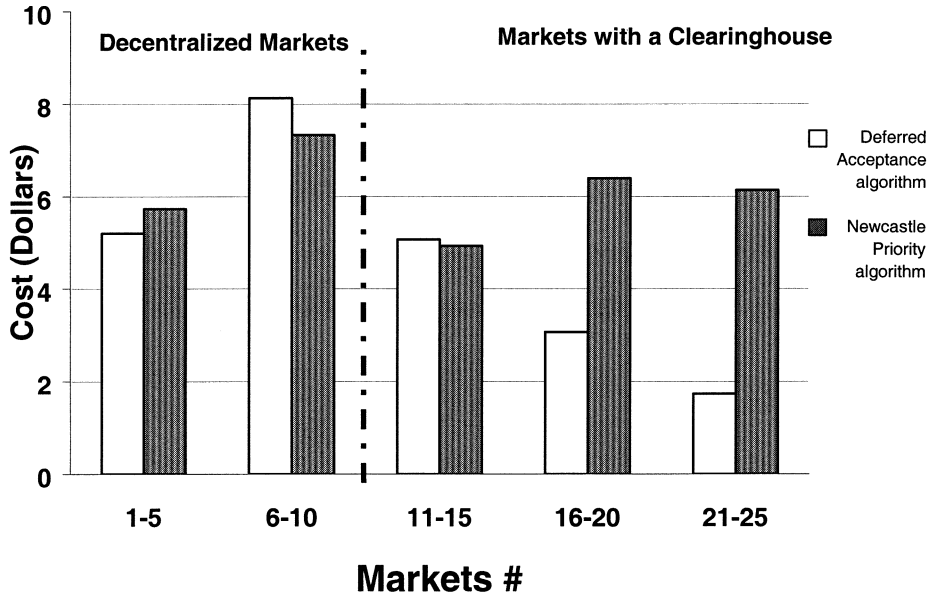


FIGURE 1.— Average costs of early markets, over time. In the first ten markets, #1–10, only the decentralized matching technology is available, and the participants match early, despite the cost. Starting with Market 11, and continuing through market 25, participants who wait until the last market period can use a centralized matching mechanism. In one cell of the experiment this was a stable, deferred acceptance algorithm of the kind used in Edinburgh, and in the other it was the unstable priority algorithm used in Newcastle. With the stable mechanism, costs of going early fell over time, indicating that more matches were being made at the efficient time.

the differences in outcomes can be unambiguously attributed to the matching algorithms. This adds confidence that the differences observed in the field are also due to the different clearinghouse designs (and not to some other difference between Edinburgh and Newcastle).¹⁶

¹⁶ The subsequent experiment by Unver (2000b, c) adds support to the hypothesis that the long life of the linear programming mechanisms used at the London Hospital and in Cambridge has less to do with desirable features of those algorithms than with special features of those markets (these are the two smallest markets, and each involves only the graduates of a single medical school and jobs in the associated teaching hospital (Roth (1991))). In Unver's experiments, the linear programming mechanism performed no better than the priority matching mechanism when the mechanisms were employed in otherwise identical markets. Unver goes on to reproduce the main laboratory observations in a computational comparison of the mechanisms, in which each of the experimental conditions are played by genetic algorithms rather than by human subjects. This further clarifies the manner in which the mechanisms differ in the rewards they give to different strategies. And the comparisons with a prior computational study of these markets (Unver (2000a)) are illuminating in that they suggest that some kinds of strategic behavior come more naturally than others to human subjects; that

2.4. *Complications in the American Medical Market*

As the above discussion suggests, there is ample evidence that stability of the outcome is a critical element of clearinghouse design. The evidence proved sufficiently compelling to the various stakeholders in the American medical market so that the design discussion was always framed in terms of what kinds of stable matching mechanisms would perform best from various points of view. But this turns out to be a deceptively simple question, because the complications of the medical market radically change the theoretical properties of any matching mechanism. To state the matter starkly, none of the conclusions of Theorems 1–4 apply to the medical match, not even that a stable matching always exists. Under the conditions that actually prevail in the match, counterexamples can be constructed to all the conclusions of Theorems 1–4.¹⁷

What makes the NRMP different from a simple market is that it has complications of two kinds: complications that cause two positions to be linked to one another, and complications that involve the preferences of employers for variable numbers of workers. In the first category are *couples*, who need a pair of positions, and individual applicants who also need two positions, because they match to positions for second year graduates, and then need to find a prerequisite first year position. In the second category are requests by residency programs to have an even or an odd number of matches, and reversions of unfilled positions from one residency program to another, typically in the same hospital.¹⁸

For the present discussion, I will concentrate on the implications for design of the fact that there are couples in the market. In the 1990's there were somewhat more than 20,000 applicants for residencies participating in the match each year, and of these approximately 1,000 were in couples, seeking two jobs together. Table II gives statistics for the five years used most heavily for computational experiments during the design.

In the early 1980's, there were already couples in the market, and attempts had been made to modify the clearinghouse to accommodate them. However these attempts had not met with much success; many couples circumvented the match and made arrangements with hospitals on their own. It will help clarify the design problem to briefly consider that earlier attempt.

Prior to 1983, couples participating in the match (after being certified by their dean as a legitimate couple) were required to specify one of their members as the "leading member." They would then submit a rank ordering of positions for each

study showed that when genetic algorithms are allowed to operate on a wide strategy set, they might learn to manipulate the unstable mechanisms more thoroughly than the human subjects in the experiment did.

¹⁷ Except of course the impossibility result in the second sentence of Theorem 4.

¹⁸ This situation arises when the director of a program for second year graduates, e.g. in neurology, who require first year training, e.g. in internal medicine, makes an arrangement with the internal medicine director that the neurology residents will spend their first year with the other first year medicine residents. So internal medicine will now seek to hire fewer residents than it would otherwise. However, if the neurology program doesn't fill as many positions as anticipated, these empty positions have to be "reverted" to the medicine program, so that it will have the residents it needs.

TABLE II
NUMBER OF APPLICANTS, AND COUPLES BY YEAR

Year:	1987	1993	1994	1995	1996
Number of applicants submitting preference lists	20071	20916	22353	22937	24749
Number of applicants participating as part of a couple (i.e. twice the number of couples)	694	854	892	998	1008

member of the couple; i.e. a couple submitted two preference lists. The leading member was then matched to a position in the usual way, the preference list of the other member of the couple was edited to remove distant positions, and the second member was then matched if possible to a position in the same vicinity as the leading member. The algorithm applied to these lists was essentially a hospital proposing deferred acceptance algorithm, and so the resulting outcome was stable with respect to the individual lists submitted. That is, the resulting matching had no instabilities with respect to the individual worker-firm pairs. However it is easy to see why such an outcome would often in fact have an instability involving a couple and perhaps two employers, once we recognize that instabilities involving couples may look slightly different than those involving single workers.

Consider a couple whose first choice is to have two particular jobs in, say, Boston, and whose second choice is to have two particular jobs in New York. The leading member might be matched to his or her first choice job in Boston, while the other member might be matched to some undesirable job in Boston. Since the fundamental law of marriage is that you can't be happier than your spouse, an instability could now exist. If their preferred New York residency programs ranked them higher than students matched to positions in those programs, the couple, on calling the hospitals in New York, would find those hospitals glad to have them. (The students originally matched to those New York positions might be told that, due to an unexpected budget shortfall, their positions no longer existed.)

To make this precise, we can augment our simple model to accommodate couples who have preferences over pairs of positions, and can submit a rank order list that reflects this.

2.4.1. *A More Complex Market: Matching with Couples*

This model is the same as the simple model above, except the set of workers is replaced by a set of applicants that includes both individuals and couples.

Denote the set of applicants by $A = A1 \cup C$, where $A1$ is the set of (single) applicants who seek no more than one position, and C is the set of couples. A member of C is a couple $\{a_i, a_j\}$ such that a_i is in the set $A2$ (e.g. of husbands) and a_j is in the set $A3$, and the sets of applicants $A1$, $A2$, and $A3$ together make up the entire population of individual applicants, denoted $A' = A1 \cup A2 \cup A3$.

The reason for denoting the set of applicants both as A and as A' is that from the point of view of the firms, the members of a couple $c = \{a_i, a_j\}$ are two distinct applicants who seek distinct positions (typically in different residency programs), while from the point of view of the couple they are one agent with a single preference ordering of pairs of positions. That is, each couple $c = \{a_i, a_j\}$ in C has preferences over ordered pairs of positions, i.e. an ordered list of elements of $F \times F$. The first element of this list is some (r_i, r_j) in $F \times F$ which is the couples' first choice pair of residency programs for a_i and a_j respectively, and so forth. Applicants in the set $A1$ have preferences over residency programs, and residency programs (firms) have preferences over the individuals in A' , just as in the simple model. A matching is a set of pairs in $F \times A'$.

Each single applicant, each couple, and each residency program submits to the centralized clearinghouse a Rank Order List (ROL) that is their stated preference ordering of acceptable alternatives.

As in the simple model, a matching μ is blocked by a single applicant (in the set $A1$), or by a residency program, if μ matches that agent to some individual or residency program not on its ROL. A matching is blocked by an individual couple (a_i, a_j) if they are matched to a pair (r_i, r_j) not on their ROL. No individual or couple blocks a matching at which he or it is unmatched.

A residency program r and a single applicant a in $A1$ together block a matching μ precisely as in the simple market, if they are not matched to one another and would both prefer to be.

A couple $c = (a_1, a_2)$ in A and residency programs r_1 and r_2 in F block a matching μ if the couple prefers (r_1, r_2) to $\mu(c)$, and if either r_1 and r_2 each would prefer to be matched to the corresponding member of the couple, or if one of them would prefer, and the other already is matched to the corresponding couple member. That is, c and (r_1, r_2) block μ if:

1. $(r_1, r_2) >_c \mu(c)$; and if either
2. $\{(a_1 \notin \mu(r_1), \text{ and } a_1 >_{r_1} a_i \text{ for some } a_i \in \mu(r_1) \text{ or } a_1 \text{ is acceptable to } r_1 \text{ and } |\mu(r_1)| < q_1\}$ and either $a_2 \in \mu(r_2)$ or $\{(a_2 \notin \mu(r_2), a_2 >_{r_2} a_j \text{ for some } a_j \in \mu(r_2) \text{ or } a_2 \text{ is acceptable to } r_2 \text{ and } |\mu(r_2)| < q_2\}$

or

3. $\{a_2 \notin \mu(r_2), \text{ and } a_2 >_{r_2} a_j \text{ for some } a_j \in \mu(r_2) \text{ or } a_2 \text{ is acceptable to } r_2 \text{ and } |\mu(r_2)| < q_2\}$ and $a_1 \in \mu(r_1)$.

A matching is *stable* if it is not blocked by any individual agent or by a pair of agents consisting of an individual and a residency program, or by a couple together with one or two residency programs.

It isn't hard to see why the presence of couples will cause design problems. Consider the deferred acceptance algorithm discussed above for the market without couples. It is a "one pass" algorithm: in the worker proposing version, each worker moves down her preference list only once. The reason this produces a stable matching is that no firm ever regrets any rejections it issues, since it only does so when it has a better worker in hand. So there is never a need for a worker to reapply to a firm that has already rejected her.

Now consider the case in which some pair of workers is applying for positions as a couple (a, b) , by submitting a preference list of pairs of firms, and suppose that at some point in the algorithm their applications are being held by the pair of firms (f, g) , and that, in order to hold b 's application, firm g had to reject some other worker c . Suppose at the next step of the algorithm, firm f , holding a 's application, gets an offer it prefers, and rejects worker a . Suppose further that couple (a, b) 's next choice, after positions at firms f and g , is positions at two other firms f' and g' . Then in order to move down the couple's preference list, worker b has to be *withdrawn* from firm g . This creates a potential instability involving firm g (which now regrets having rejected worker c), and worker c . So an algorithm like the deferred acceptance procedure will no longer be able to operate in one pass through worker c 's preferences, since to do so would miss this kind of instability.

This is not a difficulty linked to a particular kind of algorithm. In Roth (1984) I showed that, when couples are present, the set of stable matchings may be empty, and hence no algorithm can be guaranteed to converge to stability. Parenthetically, the difference between the treatment in that paper and in the design problem illustrates what I mean when I say that design carries a responsibility for detail. In Roth (1984) I found it sufficient to note that allowing couples to submit preference lists over pairs of positions would improve the chance of reaching a stable outcome, but that this might not completely solve the problem because stable matchings might sometimes not exist. While that was a reasonable observation to make as a disinterested observer, as a designer I had to think about an algorithm that will perform well, even in the absence of a theorem establishing that it will always produce a stable matching.

It seemed likely that no "one pass" algorithm would be feasible, and that therefore any algorithm would need to check for and resolve instabilities that might be present at intermediate stages. With this in mind, we took as the basis for the conceptual design of the new algorithm the class of algorithms explored in Roth and Vande Vate (1990), which, starting from any matching, seek to resolve instabilities one at a time. That paper showed how, in a simple market, instabilities could be sequenced in such a way that the process would always converge to a stable matching. The idea is that, starting from an unstable matching, a new matching can be created by "satisfying" one of its blocking pairs (f, w) ; i.e. creating a new matching in which f and w are matched with one another, perhaps leaving a worker previously matched to f unmatched, and a position previously occupied by w vacant.¹⁹

When the algorithm starts, all positions are empty, and the algorithm begins, like the applicant-proposing deferred acceptance algorithm, by satisfying blocking pairs involving unmatched applicants. Then, as potential instabilities develop due to the presence of couples, these are resolved one at a time. Each step of the algorithm begins by selecting an applicant, either an individual or a couple, and

¹⁹ The paper was motivated by the observation by Knuth (1976) that the process of satisfying a sequence of blocking pairs could sometimes cycle, and return to matchings already considered.

letting that applicant start at the top of the applicant's preference list, and work down until the most preferred firm or firms willing to hold the application is reached. Applicants who are displaced in this process continue working their way down their preference lists, and when this causes a member of a couple to be withdrawn from a residency program, that residency program is put on the "hospital stack" to be checked later for potential instabilities. It is this process that causes the algorithm to take more than one pass through some applicants' preferences, as applicants who form blocking pairs with a hospital will be put back on the "applicant stack" to be considered again, starting with their most preferred programs. Throughout, the algorithm is applicant-proposing in the sense that, whenever there are multiple blocking pairs involving the same applicant, the blocking pair that is satisfied next is the one most preferred by the applicant.

In a simple market, without couples, the order in which applicants are placed in (and hence drawn from) the applicant stack can be shown not to matter, because the algorithm would produce the applicant-optimal stable matching no matter in what order the applicants were processed. However the order can matter when couples are present. Consequently, we performed computational experiments before making sequencing choices. These computational experiments focused on two issues:

- Do sequencing differences cause substantial or predictable changes in the match result (e.g., do applicants or programs selected first do better or worse than their counterparts selected later)?
- Does the sequence of processing affect the likelihood that an algorithm will produce a stable matching?

Computational experiments to test the effect of sequencing were conducted using data from three NRMP matches: 1993, 1994, and 1995. The results were that sequencing effects existed, but were unsystematic, and on the order of 1 in 10,000 matches. In the majority of years and algorithm sequences examined, the match was unaffected by changes in sequencing of algorithm operations, and in the majority of the remaining cases only 2 applicants received different matches. However sequencing decisions did influence the speed of convergence to a stable matching. Because sequencing decisions had no systematic effect on outcomes, it was decided to design the algorithm to promote rapid convergence to stability.

Based on these computational experiments, the applicant proposing algorithm for the NRMP was designed so that all single applicants are admitted to the algorithm for processing before any couples are admitted. This reduces the number of times that the algorithm encounters cycles and produced the fastest convergence.

Of course in examples for which no stable matching exists, like that of Roth (1984), no procedure for detecting and resolving cycles will produce a stable matching. But one result of the computational experiments conducted during the design of the algorithm is that the procedure never failed to converge to a stable matching. So there is reason to believe that the incidence of examples with no stable matchings may be rare.

TABLE III
 COMPUTATIONAL EXPLORATION OF THE DIFFERENCE BETWEEN HOSPITAL AND
 APPLICANT PROPOSING ALGORITHMS

Applicants	1987	1993	1994	1995	1996
Number of Applicants Affected	20	16	20	14	21
Applicant Proposing Result Preferred	12	16	11	14	12
Program Proposing Result Preferred	8	0	9	0	9
New Matched	0	0	0	0	1
New Unmatched	1	0	0	0	0

2.4.2. Comparison of the Applicant and Program Proposing Algorithms

Once alternative worker-proposing and firm-proposing algorithms could be compared, it was possible to examine the scope for (designer) discretion in choosing a stable matching. Unexpectedly, it turned out that this scope is very limited. The requirement that a matching be stable determines 99.9% of the matches; only one applicant in a thousand is affected by the choice of algorithm. Table III illustrates this on the data from the same matches shown in Table II. Recall that in each of those matches, there were more than 20,000 applicants. But as Table III shows, only about 20 applicants a year would have received different jobs from an applicant proposing algorithm than from a hospital proposing algorithm.

The table confirms that some of the properties of the set of stable matches are different in fact as well as in theory when couples and the other complications of the medical market are present. In a simple match, without couples or other complications, all of the applicants would have preferred the applicant proposing match, and no applicants who were matched or unmatched at the outcome of one algorithm would change employment status at the outcome of the other.²⁰

On the other hand, Table III also illustrates the very small magnitude at which these differences from the simple model are exhibited. Of the more than 100,000 applicants involved in the data from which Table III is drawn, only two who were unmatched at one stable matching were matched at another. (Note also that even this tiny difference is unsystematic; it doesn't suggest that one or the other of the algorithms produces a higher level of employment.)

If this were a simple market, the small number of applicants whose matching is changed when we switch from hospitals proposing to applicants proposing would imply that there was also little room for strategic behavior when it

²⁰ Table III reflects results from the actual market, not the simplified market with couples as its only complication. It is easiest to see why the welfare comparisons from the simple model do not carry over to the medical market by considering the case of an individual who needs two jobs, e.g. a second year job in his desired specialty and a first year job that provides the necessary preparation. If he does better in the applicant proposing algorithm because of an improvement in his specialty position, he now requires a corresponding first year position, from which he displaces another applicant, who consequently does worse.

comes time to state rank order lists. Theorem 4 doesn't guarantee that this will be the case in the complex market. However the method of proof of Theorem 4 allowed a set of computational experiments on the submitted preference lists to be designed that would determine an upper bound on the number of applicants who could potentially have profited from changing their preference lists, based on the preferences submitted in previous years. *Under the assumption that the preference lists submitted in previous years represent the true preferences*, these computational experiments confirmed that the numbers of applicants who could have potentially profited by submitting different (shorter) preference lists are almost the same as those in Table III. Similarly, the number of residency programs that could potentially profit by changing their preferences or their stated capacities (cf. Sonmez (1997)) is comparably small (see Roth and Peranson (1999) for the design and results of these computational experiments).

However the assumption that the submitted preference lists are a good proxy for the true preferences needs careful investigation. If instead of reflecting the true preferences, the submitted preferences instead reflect misrepresentations of the preferences induced by experience with the existing clearinghouse, then it could be that a new algorithm would over time elicit quite different preferences, and affect many more applicants than the above calculations suggest. To state the competing hypotheses starkly, it could be that the set of stable matchings is small because the market is large, or it could be that the set of stable matchings is in fact large, but appears small because participants have successfully gamed the system. (At equilibrium, the set of stable matchings would appear to be small in terms of the submitted preferences.) Some further computation was required to resolve this issue.

Roth and Peranson report a set of theoretical computations on simple matching models with randomly generated preferences, to determine how the size of the market influences the size of the set of stable matchings, measured by how many applicants receive different matches at different stable matchings. For these computations we considered markets with no match complications, so that Theorems 1–4 apply. In particular, firm and worker optimal stable matches exist, and we can compute the number of applicants who receive different matches at different stable matchings by simply counting the applicants who receive different matches at the optimal stable matchings for each side. (The proof is simple: if an applicant's best and worst stable matching are the same, then he is matched to the same firm at every stable matching.)

When preferences are highly correlated, the set of stable matchings is small regardless of the size of the market. But when preferences are uncorrelated, the core quickly grows quite large, if as the market grows large, the number of firms on an applicant's preference list grows correspondingly large. (By the time there are 500 firms and workers, over 90% of workers receive different matches at different stable matchings). However, in the medical market, applicants and residency programs only list one another on their preferences if they have completed an interview, and so no applicant has a very long list of programs (and the vast majority have fewer than 15 programs on their submitted preference list). When




FIGURE 2.— Small core of large markets, with k fixed as n grows. $C(n)/n$ is the proportion of workers who receive different matches at different stable matchings, in a simple market with n workers and n firms, when each worker applies to k firms, each firm ranks all workers who apply, and preferences are uncorrelated (from Roth and Peranson (1999)). Note that for any fixed k , the set of stable matchings grows small as n grows large.

this restriction is added, the set of stable matchings shrinks rapidly as the size of the market grows, and essentially reproduces the results obtained in the computational investigation of the medical market (see Figure 2). In these computations, there is no question what are the true preferences, and the fact that the set of stable matchings is small confirms that there are effectively no opportunities for firms or applicants to profitably manipulate their preferences. That is, in the simulations, we know the true preferences, and we see that the opportunities for profitable strategic misrepresentation are as small as they are in the field data.

The fact that only one in a thousand applicants in the match could even potentially profit from misrepresenting his or her preferences, together with the fact that no one can tell if they are in this tenth of a percent of the population, and that those who are not can only hurt themselves by misrepresenting their preferences, makes it easy to advise applicants that the incentives for straightforward reporting of rank order lists are clear.

Another observation from the matches we have analyzed is that we have never yet observed in the field data a year in which no stable matching could be found. This is despite the fact that, when couples and other complementarities exist, the set of stable matchings can be empty. It appears that, when the percentage of couples is not too high, as the market becomes large and the set of stable

matchings becomes small, it also becomes less likely to be empty. Preliminary computational simulations support this conjecture, but at the moment it remains a conjecture.

The computational results suggest an agenda for theoretical work quite different than might have been expected to follow from the prior theory. We cannot hope to find, say, restrictions on agents' preferences that will be met in the market and will allow us to generalize the conclusions of Theorems 1–4. The conclusions of those theorems don't in fact generalize to the medical market. But the conclusions of the theorems are in some sense close to being correct; the number of individual firms and workers for whom the conclusions are violated is small. The computational results suggest that there may be theorems that explain why it becomes increasingly unlikely that the set of stable matchings will be either large or empty, as the market grows large.

The availability of computation meant that the design effort could proceed without waiting for theoretical resolution of outstanding problems. Computation was used in several quite different ways in the course of the design and evaluation of the new medical labor market. We relied on computation in three places:

- Computational experiments were used in the algorithm design.
- Computational explorations of the data from previous years were used to study the effect of different algorithms.
- Theoretical computation, on simple markets to which existing theoretical results apply, was used to understand the effect of market size.

Before moving on, a word is in order about the political context in which this new design was adopted as the matching algorithm for the American medical match. The crisis of confidence that sparked the initial demand for a new market design also led to a heightened sensitivity about the conduct of the design effort. Early in the process I fielded visits from the American Medical Students Association, as well as numerous conference calls with members of the board of directors of the National Resident Matching Program, whose members represent a variety of institutional and student interests. While I worked on the design, I maintained a web page on which I posted my interim reports.

My status was that of an outside expert hired to design a new algorithm that would be able to handle all the match complications, and to evaluate the scope for favoring one side of the market over the other (i.e. applicants and residency programs) while achieving a stable matching. The responsibility for deciding whether to adopt the new design was retained by the NRMP, in consultation with its various constituencies. Once my design and evaluation were complete, I conferred at length with the various interested parties (including travelling to present the results to representatives of the various organizations of residency directors). Although it was widely anticipated that the results of the study would provoke bitter disagreement, the fact that the set of stable matchings proved to be so small was widely understood to mean that making the match as favorable as possible to applicants would not create any systematic problems for any segment of residency programs. Consequently my reports were received without provoking much controversy, and the NRMP board voted in May of 1997 to adopt the

new algorithm for the match starting in 1998. The first few years of operation of the new match design seem to have been extremely smooth, and while many issues related to the medical labor market continue to be of lively concern, the organization of the match no longer appears to be a source of significant controversy.²¹

We turn next to consider briefly an ongoing design process that is embedded in a much more formal political process.

3. THE FCC SPECTRUM AUCTIONS²²

In 1993 the US Congress amended the Communications Act of 1934 to direct the Federal Communications Commission (FCC) to design and conduct auctions to efficiently allocate radio spectrum licenses. The first auction was held in July of 1994. In the interim, the FCC hired John McMillan (then of UCSD) to advise their staff, and instituted a series of hearings at which the major potential bidders could offer proposals and comments on auction design. Many of the interested parties also hired game theorists to help formulate their proposals.

How spectrum licenses were previously allocated in the U.S. didn't give much guidance, since spectrum had been given away free for most of the 20th century. Until 1981, spectrum licenses were allocated through a political process called "comparative hearings." After 1981, licenses were allocated by lottery. Both procedures led to lots of rent-seeking behavior and bureaucratic complications, and to very substantial delays.

However, spectrum licenses had been auctioned overseas, and the experience of these related markets taught some important lessons. In Australia, satellite-television licenses had been sold in a sealed-bid, first-price auction with rules that merely specified that, if the winning bid were withdrawn after the auction,

²¹ Because design involves detail, I have chosen to tell one story in depth rather than many without detail, but I can't resist putting in some pointers to the elegant theoretical and experimental work on matching individuals with indivisible objects such as student housing in Abdulkadiroglu and Sonmez (1998, 1999, 2000), Balinski and Sonmez (1999), and Chen and Sonmez (2002). That work builds on the observation by Shapley and Scarf (1974) that in a simple "housing market" model of allocation of indivisible goods there is a nonempty core reachable by an algorithm of "top trading cycles" proposed by David Gale, which yields a unique allocation when preferences are strict (Roth and Postlewaite (1977)), and that the mechanism that chooses this allocation is strategy proof (Roth (1982b)). Note also the practical work of designing decentralized web-based job matching services such as reported in Nakamura, Percy, Davenport, Fraser, and Piper (1998). In this latter context, note that complementarities will be endemic to job markets, since they arise, e.g., even from budget constraints (Mongell and Roth (1986)). For a contemporary problem in labor market design, see the study of the market for law clerks for the US Circuit Courts of appeals in Avery, Jolls, Posner, and Roth (2001), and for some experimental tests of proposed design solutions see Haruvy, Roth, and Unver (2002).

²² Detailed accounts of some of the events reviewed in this section can be found in FCC (1997, 2000) and on the FCC auction website at <http://www.fcc.gov/wtb/auctions/Welcome.html>, in Cramton (1995), McAfee and McMillan (1996), McMillan (1994), and Milgrom (2000). I concentrate here on a few issues, and ignore others, such as the mandate to design the auctions in a way that made special provisions for certain affirmative action concerns (on which see Ayres and Cramton (1996)).

the next highest bid would become the winning bid. This procedure was gamed in spectacular fashion by a newcomer to the industry. After the auction closed, it was found that not only had this bidder submitted the highest bid, but it had also submitted the next highest bid, and the next, and the next. By withdrawing each high bid in turn, it eventually purchased the two licenses up for auction at massively lower prices than its initially winning bids (McMillan (1994)). With this experience in mind, the withdrawal rules adopted by the FCC required up-front deposits by bidders wishing to participate in the auction, and established that a high bidder who withdrew his bid would be liable for the difference between the withdrawn winning bid and the actual selling price.

Much of the discussion of auction design focused on the questions of how to promote efficient allocation, by eliciting bidding that would reveal the value of the licenses to the bidders, and allocate licenses where they were most valued. Because a spectrum license for new communication services has a large but uncertain value, one hazard facing bidders, called the “winner’s curse,” is that a bidder who overestimates the value of the license is more likely to submit the winning bid, and runs the risk of bidding more than the license will prove to be worth. Knowing that other bidders have comparable estimates would reduce the chance that the bidder’s own estimate was mistaken. Not knowing how much other users think the license is worth, therefore, means that each bidder has to treat his own estimate with great caution. In a sealed bid auction, this would mean that a bidder would be wise to bid substantially less than his estimated value for the license.²³

To allow bidders to get a sense of what the other bidders think the licenses are worth (“price discovery”), it was decided not to have a sealed bid auction, but to let bidders observe each others’ bids in a multi-round ascending bid auction. It was further decided that this should be a first price auction, in which the winning bidders pay the full price they have bid.²⁴

A chief concern was how to deal with the potential complementarities that might influence bidders’ valuations of groups of licenses. To be concrete, suppose

²³ Our current understanding of the winner’s curse is itself a case study of the interaction between field data, theory, and experimental economics. It was discussed in the context of auctions for drilling rights for oil (Capen, Clapp, and Cambell (1971)), it was modelled theoretically in “common value” auctions in which agents understand the danger and how to discount appropriately (cf. Wilson (1969), Milgrom and Weber (1982a,b), Klemperer (1998)) and it has been studied experimentally in the laboratory where profit-making behavior is sometimes learned only slowly and painfully (cf. Kagel and Levin (1986, 1999)).

²⁴ McMillan (1994) writes that based on the experience of a spectrum auction in New Zealand, it was judged politically unwise to have a second price auction that would be subject to the criticism that the government was selling licenses to firms at a price less than what it knew they were willing to pay. Bidders too may be reluctant to reveal their true reservation prices as freely as they might do in truly isolated second price auctions (cf. Rothkopf, Tiesberg, and Kahn (1990)), so that not all the potential advantages of second price auctions may be realizable in practice. However, for a generalization of second price auctions to multiple goods, see the patent by Ausubel (2000). (The fact that auction methods are now patented speaks volumes about the changing nature of the design business.)

a bidder values licenses A and B at 100 (million) each if he can get both, but otherwise only at 50 each. How should they be auctioned so that if their highest value is achieved only if they are owned together, the bidder can afford to bid aggressively on them, without too much risk that he will win only one of them, but pay more than it is worth on its own?

Some complementarities could be dealt with by appropriate definition of the licenses themselves. For example, the first auction, in 1994, was for Narrowband Personal Communication Services (PCS), two-way paging services in which a central transmitter relays a message to a personal device, which can then transmit a return message. The central transmitter is powerful, and can transmit on a noisy frequency, but the personal device is low powered, and must transmit on a quiet frequency. So efficient use of this technology calls for the pairing of two complementary frequencies. Rather than rely on the auction to aggregate efficient pairs, the Narrowband PCS licenses were each defined to be for an appropriate pair of frequencies. That is, from the outset, the rights that were being auctioned were for a complementary pair of frequencies (see Cramton (1995)).²⁵

However not all potential complementarities can be clearly defined by the technology, and so a major design question was how to structure the auction to allow bidders to take into account any important complementarities in their valuations. The idea of auctioning spectrum licenses one at a time was rejected for this reason, as was the idea of simultaneously beginning separate auctions for many licenses, but letting each one end independently, when no bidder wished to raise his bid on that license. Instead, in proposals put forward by Preston McAfee of the University of Texas (representing Airtouch Communications) and by Paul Milgrom and Bob Wilson of Stanford (representing Pacific Bell), it was suggested that all the licenses being sold at a given time be auctioned simultaneously, and that none of the auctions should end until they all did. That is, in this proposal, which was ultimately adopted, the market for every license would remain open until no bidder wished to raise his bid on any license. (The impact of the academic commentators, especially Milgrom and Wilson, on all aspects of the FCC design, is evident in the FCC documents, for example the Notice of Proposed Rulemaking (FCC, 1993) and the Final Report and Order (FCC, 1994). At the same time, the tremendously detailed rules needed in order to produce a working auction, and to guard against foreseeable contingencies, represent many person-months of work by the FCC staff.)

There was concern that the auction rules should not give bidders an incentive to avoid making serious bids until the end, in an effort to benefit from the information contained in other bidders' bids, without revealing any of their own.

²⁵ Of course no auction design can prevent the winners' curse in an auction involving a new technology like paging, for which there are wildly different forecasts. One of the licenses I have followed was won at the auction in 1994 by KDM Messaging/ATT for \$80,000,000, but was sold to TSR Wireless for \$20,000,000 in 1999, and after TSR declared Chapter 7 bankruptcy in December 2000, this license was sold at the liquidation auction for \$3,585,266 (*In re* TSR Wireless, LLC, Nos. 00-41857 & 0041858 (Bankr. D. N.J. Mar. 28, 2001) at [4]). (Of course, a sharp price decline is not itself a proof that the winner's curse was the cause.)

This would inhibit the price discovery that the multi-round design is intended to promote, and might also simply make the auctions long and cumbersome.

To prevent bidders from concealing their intentions by delaying their bids, the spectrum auctions imposed an *activity rule*, proposed by Milgrom and Wilson. Under this rule, bidders had to maintain their eligibility to bid on a given volume of licenses (measured in population of the area covered) by being the high bidder, or by raising their bids by specified minimum percentages, on a sufficient volume of licenses. Bidders who did not remain active in this way on a sufficient volume of licenses would see their eligibility decline to reflect the volume of licenses they were actively bidding on. That is, this activity rule specified that a bidder who did not remain active would not be able to suddenly become active on a high volume of licenses late in the auction.

As of this writing more than two dozen spectrum auctions have been run under a simultaneous, multiple round auction design, with numerous small modifications made along the way based on early experience.²⁶ Overall, these auctions appear to be working smoothly, and they have raised many billions of dollars of revenue.

The largest design change to date is presently scheduled to be implemented in the 2002 auction of 700 MHz spectrum licenses suitable for high speed internet access. Instead of requiring bidders to submit bids for individual licenses, and to assemble the packages they want by bidding simultaneously on multiple licenses, the new rules allow bidders to explicitly bid for *packages* of licenses. This is meant to solve the “exposure problem” of simultaneous bidding, in which bidders with strong complementarities are exposed to the risk of winning only part of the package they are trying to assemble, and of having bid more for that part than it is worth to them on its own. Under package bidding, a bidder will either win the whole package or none of it.

To determine the winning bids, when bidders bid on different packages of their own devising, the auction has to determine the revenue-maximizing set of packages. This will make the auction more complicated in the sense that a bid for a particular package that is not part of the revenue maximizing set of packages when it is made, may later become part of the revenue maximizing set, as higher bids for other packages come in.

²⁶ One of the more difficult problems encountered to date has to do with who owns the license if, after submitting a winning bid, a firm files for bankruptcy. In the C block broadband PCS auction, the rules allowed generous installment payment plans, “which led to all the major bidders defaulting and declaring bankruptcy” (Cramton (2000)). The FCC no longer allows such payment plans, but resolving the bankruptcies, and putting this spectrum to use may depend on further legislation or on rulings from the bankruptcy court. Making collusion difficult has also been a concern. In early auctions, bidders signaled their intentions using the last digits of their (six digit) bids to communicate (three digit) license identification numbers. (See Cramton and Schwartz (2000) for some detailed examples in which threats were communicated in this way.) In 1997 the FCC changed the bidding format, and allowed only bids in preset increments, to prevent this. Bid withdrawals were also used for signaling in early auctions, and these too have been limited. See the FCC’s anticollusion page at www.fcc.gov/wtb/auctions/collusio/collusio.html.

Interestingly, discussion of package bidding began even prior to the first auction in 1994, because of concern that complementarities might be of great importance.²⁷ Two kinds of potential difficulties with package bidding loomed large in these discussions. The first class of anticipated difficulties concerned combinatorial properties of such an auction. Because of the large number of possible packages (there are $2^n - 1$ subsets of n licenses), it might not be possible to run a package bidding auction properly, either because of the computational difficulty of computing the revenue maximizing set of packages, or because of difficulty in eliciting and presenting the bids in a way that would make the progress of the auction easily comprehensible to the bidders.

A second class of difficulties concerned the incentives that bidders might face in a package bidding auction. With so many potential packages to bid on, activity rules might lose their force, and bidders who wished to delay revealing the packages they were interested in and the prices they were willing to pay might be able to maintain apparent activity by “parking” their bids on packages unlikely to be part of the winning set of packages.

Also, package bidding can create a free-rider problem among small regional bidders in a way that gives an artificial advantage to large bidders who want packages providing comprehensive national coverage. Consider the problem facing a regional firm interested in a single license. Its bid will be part of the revenue maximizing set of packages only if its bid, together with those of the other regional bidders, sum to more than the bid on the national package comprising licenses in all the regions. But this means that the success of the regional bids depends mostly on the bids of the other regional bidders—there is little incentive for each regional bidder to bid aggressively. (And, since each winning bidder must pay the full amount of the winning bid, there is ample incentive to stick with a low bid and hope that the other regional bidders will raise their bids enough to raise the sum of the regional bids above the bid for the national package.) Consequently, a large bidder who seeks to win a package of national scope may not have to bid aggressively, and may win the national package even in the case that the regional licenses have a higher value separately. In the auction design discussion, this came to be known as the “threshold problem,” in the sense that the national bid establishes a threshold that the sum of the regional bids must surpass.²⁸

²⁷ In fact, discussion of package bidding preceded the discussion of spectrum licenses entirely, having come up in connection with other complementarities. Rassenti, Smith, and Bulfin (1982) report an experiment showing that efficiencies can be realized by allowing package bidding in an environment in which complementarities were motivated by the need of airlines for complementary takeoff and landing slots. (See also the experiments of Grether, Isaac, and Plott (1981), who report experiments with a nonauction mechanism for allocating packages of airline slots by committee decision.) For early theoretical discussions of package bidding, see Vickrey (1961) and Bernheim and Whinston (1986), and, in a two-sided matching context in which firms bid for groups of workers, Kelso and Crawford (1982).

²⁸ A more careful historian than I will have to see if the economists lined up on each side of the package bidding question were representing firms with predictable interests in the matter. (In addition, designers can have professional incentives to see their designs adopted, and these needn't always be perfectly aligned with their client's interests.) But politics is part of design, and we're going to have to learn to deal with the politics of design.

Each of these questions raised by the complementarities in the market eluded analytical solutions, but lent themselves to computational and experimental exploration. Computational studies were conducted to understand how difficult it would be to compute a winning set of bids from a set of submitted package bids. While the worst case scenarios make it computationally intractable to compute the winning set of bids, when most bidders bid on only small subsets of packages the average problem is not hard, and can be solved with commercial integer programming packages (see, e.g., Kelly and Steinberg (2000), Rothkopf, Pekec, and Harstad (1998), or see deVries and Vohra (2000) for an overview).

Ease of use, and threshold problems, were addressed at least in principle by experiments (cf. Cybernomics (2000), Plott (1997), Ledyard, Porter, and Rangel (1997)). Because the package bidding mechanisms being considered are all novel, there isn't a source of field data that would be informative for design. So, unlike in the design of the medical match, the role of experiments here wasn't to complement field data, but to add an empirical component to the discussion in the absence of field data. In addition, the impetus for package bidding was motivated by the strong sense that complementarities existed, but without any data or models to predict their distributions. So experiments were constructed not to test specific hypotheses related to the spectrum market, but rather as "proof of concept" demonstrations that package bidding could achieve efficiencies that might be missed in single item auctions. (Plott (1997) describes the role that experiments played at various parts in the FCC's process of soliciting comments and advice.) The experiments show that package bidding can indeed achieve some of the hoped for efficiencies in simple environments. While the implications of these results for the proposed spectrum auctions remained a subject of lively controversy among the potential bidders and their advisors, experiments provided an empirical dimension to the debate.

The Cybernomics experiment compared the simultaneous ascending auction to a package bidding auction under various configurations of values. In the experiment, as complementarities became more important in bidder valuations of packages, the efficiency of the simultaneous ascending auction diminished, while the package bidding auction was largely unaffected. The package bidding auctions also took many more rounds to complete. Partly motivated by this latter observation, Ausubel and Milgrom (2001) analyzed the performance of a simple package bidding auction in which bidders are assumed to bid "straightforwardly," in the sense that at each round they make whatever is the minimum allowable bid at that time on the package that would maximize their profits at current prices.

They observe that when bidders bid straightforwardly, the auction operates as a deferred acceptance algorithm, with results quite similar to those found in the matching literature. In particular they deduce some analogs to the matching results in Theorems 1, 2, and 4, discussed earlier. The outcome of the package bidding auction they analyze with straightforward bidding will be the point in the core that is optimal for the buyers (i.e., it minimizes seller revenue), and if goods are substitutes, straightforward bidding will be a dominant strategy for buyers.

As in the case of matching theory, complementarities and budget constraints can create complications.²⁹

Following a conference on combinatorial bidding held in May, 2000, at which a wide range of views were solicited, the FCC (2000, 2001) rules for package bidding strike a cautious compromise.³⁰ Based largely on a proposal by Milgrom (2000), the auction allows bidders to formulate bids for no more than a dozen packages. The idea is to reduce the combinatorial complexity, as well as the opportunities to evade the activity rules by making early bids on less valuable packages. It is worth noting the similar ways in which complementarities were addressed in the spectrum auctions and in the labor market clearinghouse designs. As in the package bidding auctions, the complementarities involving couples on the job market were addressed by allowing the couples to submit “package bids” for pairs of jobs.³¹

Overall, looking at the FCC auction design process from 1993 to the present, we see an ongoing design discussion in which decisions were made, then modified and revisited in light both of experience and of further discussion. While the FCC retained design responsibility, and FCC staff made the final design decisions, and hammered out the detailed procedures, the FCC solicited and implemented suggestions from economists at every stage of the process.³² By and large, the FCC has mostly chosen to gradually adapt its initial design, rather than to make

²⁹ Ausubel and Milgrom thus establish a close parallel between the auction and matching literatures, with some of the closest points of contact in the matching literature being the “package bidding” by firms for groups of workers in Kelso and Crawford (1982) and Roth and Sotomayor (1990, Chapters 5 and 6), and Mongell and Roth (1986) (regarding budget constraints).

³⁰ The conference papers are available on the FCC website at <http://www.fcc.gov/wtb/auctions/combin/papers.html>.

³¹ Other complementarities in the medical match, not discussed in detail here, were addressed similarly. Applicants needing a second year position and a complementary first year position were invited to submit “supplementary lists” of first year positions for each such second year position, and the supplementary list was activated only when a second year position was obtained. In this way applicants needing two jobs could express their preferences for a package consisting of a pair of positions.

³² The broad participation by economists in the design discussion is made clear in the FCC call for comments (2000). Footnote 3 reads:

“The progress the Bureau has made in designing and testing a combinatorial bidding system has been made possible only by the extraordinary work done by a number of people. The procedures we are proposing are based largely on a paper presented by Stanford University Professor Paul R. Milgrom at the Conference on Combinatorial Bidding jointly sponsored by the Federal Communications Commission, the Stanford Institute for Economic Policy Research, and the National Science Foundation, that took place on May 5–7, 2000 at the Aspen Institute’s Wye River Conference Center. Paul R. Milgrom, FCC-SIEPR-NSF, Wye Woods Conference: Lessons plus a Simple Proposal (May 2000). This paper builds on ideas from many of the people who attended the conference. Some of the proposals we are considering are also importantly based on the reports by Professor Charles R. Plott of the California Institute of Technology, Charles River Associates Incorporated, Market Design, Inc., and Computerized Market Systems, Inc. that were produced pursuant to contract with the FCC, and the Cybernomics, Inc. reports by Jeffrey Banks, David Porter, Stephen Rassenti, and Vernon Smith that were also written pursuant to contract with the FCC. In addition, Professor John Ledyard of the California Institute of Technology has rendered invaluable assistance. These papers and reports, as well as the other papers presented at the conference, can be found at the Commission’s website on the conference, <http://conbin.fcc.gov>.”

radical changes. It was in this spirit that the FCC adopted the Milgrom proposal of limiting bidders to bid on no more than 12 packages of licenses. (It seems likely that this limit will be lifted in future auctions as more confidence develops in the design and operation of combinatorial auctions).

Aside from presenting economists with an opportunity to think about these design issues, the completed auctions will present an opportunity to investigate the magnitudes of some of the effects that have played a role in the design debate, including the magnitude of the complementarities between licenses. The rather different experiences of various European spectrum auctions (see Klemperer (2002b)) also present fertile grounds for further investigation.³³

A little further afield, the growing number of auctions on the internet provide a new opportunity to investigate the effects of different auction designs. Some of the concerns that occupied the designers of spectrum auctions, such as the need for rules to promote price discovery by preventing bidders from withholding their bids until the end, can be investigated by comparing different rules used for ending internet auctions. With this in mind, Roth and Ockenfels (2001) compare bidding behavior on eBay and Amazon. Auctions conducted by eBay have a fixed deadline for ending the auction, while Amazon auctions, which have otherwise similar rules, cannot close until both the scheduled end time has been reached *and* ten minutes have passed with no bidding. This difference in rules has a dramatic effect on the distribution of bids over time, with bids on eBay concentrated very near the end, while bids on Amazon are not. In terms of price discovery, early bids are less informative about final selling price on eBay than on Amazon. Of course there are other differences between the markets found on eBay and

³³ Different concerns may influence the design of related auctions, as in the recent work by Ken Binmore and Paul Klemperer on spectrum auctions in the UK. A prime concern in the U.K. market had to do with making sure that there were more licenses for sale than incumbent broadcasters, to encourage competition by newcomers. Binmore and Klemperer (2002) report that initially, the plan was to auction four licenses, and there were already four incumbent firms. It was feared that an ascending auction on the American model would allow the incumbents to collusively divide up the market, without facing much competition from new entrants, who would be deterred from bidding aggressively by a well justified fear of the winner's curse. To ameliorate these concerns, a hybrid auction was considered whose final stage would be a sealed bid auction, and experiments on this hybrid appeared promising. However the government finally was persuaded to auction an additional license, and an ascending auction on the American model was adopted, with the licenses themselves defined in a way intended to make new entrants competitive. The auction was successful in that it raised unprecedented revenue. In Germany, in contrast, the auctions were criticized by economists for their design on the grounds that new entrants might be deterred: see Jehiel and Moldovanu (2001) and Klemperer (2002a). However these auctions raised even more revenue than the English auctions (although less per capita). See Grimm, Riedel, and Wolfstetter (2001) and Wolfstetter (2001) for a different take on the German auctions. See Klemperer (2002b) for a discussion of all the European 3G spectrum auctions, most of which employed an ascending design, but many of which had disappointing results. In general, the European spectrum auctions have made clear that, quite apart from the details of auction design, the outcome depends on some of the underlying industrial organization. For example, several European auctions have been preceded by mergers of potential bidders. And to the extent that a single European market for telecommunications will emerge, it will have been affected not merely by the individual spectrum auctions conducted by national governments, but by how these auctions and their outcomes interact.

Amazon, and so this is an issue that also repays study in the controlled conditions of the laboratory. Ariely, Ockenfels, and Roth (2001) show that the difference in rules for ending the auction has the same effect in the lab as in the field. And in the laboratory, this difference in ending rules is the only difference between the auctions studied, since the same goods are sold under both rules, and (the same number of) bidders are randomly assigned to each kind of auction, rather than self selecting themselves between auctions, as in the field data. That is, here too experiments and field studies work in harness to show that in eBay-style auctions with fixed deadlines, many bidders reserve much of their activity to the closing seconds of the auction, in contrast to auctions that are automatically extended when late bids arrive.³⁴

4. CONCLUDING REMARKS

When I was asked in 1995 to redesign the labor market clearinghouse for American physicians, I had to confront the fact that none of the available theory—e.g. none of the theorems in my book on the subject (Roth and Sotomayor (1990))—could be directly applied to the complex market. All the theorems were about simpler, static models, in which there were no complementarities. Like Theorems 1–4 in this paper, the theorems all took the form “in a simple model, the following things always happen.” The only theoretical parts of the book that applied directly to the medical market were the *counterexamples*—and they all warned that, in more complicated markets, problems could sometimes arise. What was missing in the theory, but needed in design, was a sense of magnitudes; how often those problems would arise, and how big their consequences would be.

It turned out that the simple theory offered a surprisingly good guide to the design, and approximated the properties of the large, complex markets fairly well. Field and laboratory data showed that the static idea of stability went a long way towards predicting which kinds of clearinghouse could halt the dynamics of unraveling. And computation showed that many of the departures from the simple theory were small, and that some of the most severe problems that the counterexamples anticipated, such as the possibility that no stable matching would exist, were rare in large markets. Computation also revealed that large markets could achieve even nicer incentive properties than anticipated by the simple theory. That is, by unanticipated good luck, some of the knotty problems posed by couples, and other complementarities, could be solved without losing the most attractive design options that the simple theory suggested.

When I speak of unanticipated good luck, I mean that these computational and experimental results, while suggested by the theory of simple matching markets, are not explained by that theory. These results point to a need for new theory, to

³⁴ Here too the empirical observations motivate new theory: cf. Ockenfels and Roth (2001). In between field studies and laboratory experiments are field experiments in which some controlled variation is introduced into a natural, uncontrolled population. See, e.g., Lucking-Reilly (1999), who has been a pioneer in using field experiments to study internet auctions, by auctioning the same goods by different auction rules.

explain and further explore the behavior of large labor markets with couples and linked jobs. We also need new theory to explore the dynamics of market failures like unraveling, and their cures. Also, some of the nice properties of the medical market turn out to be related to the fact that each applicant interviews for only a small fraction of the jobs on the market (recall Figure 2). This seems likely to be an important variable, not only for this market. As electronic communication increases (standardized application forms, tele-conferencing, etc.) it may well be that the fraction of jobs to which an applicant can apply, and exchange meaningful information, will increase, in many markets. The results discussed here suggest that this may have important consequences, but we will need better theory than we have today to know what to expect.

Some of these same, general design themes also emerged in the design of the radio spectrum auctions. The simple theory there was somewhat less tightly connected to the final design decisions. McAfee and McMillan (1996) summarize the early role of theory in various aspects of the debate, and conclude as follows:

“A lesson from this experience of theorists in policy-making is that the real value of the theory is in developing intuition. The role of theory, in any policy application, is to show how people behave in various circumstances, and to identify the trade-offs involved in altering those circumstances. What the theorists found to be the most useful in designing the auction and advising the bidders was not complicated models that try to capture a lot of reality at the cost of relying on special functional forms. Such theorizing fails to develop intuition, as it confounds the effects of the functional forms with the essential elements of the model. A focused model that isolates a particular effect and assumes few or no special functional forms is more helpful in building understanding” (p. 172).

Here, too, computation and experimentation played a role in filling the gaps between theory and design, particularly in connection with the forthcoming auctions that allow package bidding. The simple theory organized the discussion, and the design effort has opened the door on a whole new realm of auction theory that needs to be developed. Early indications are that there are further formal connections that can be profitably made between the auction and matching literatures.

Some of this theory is starting to make inroads on the question of how complementarities make market design difficult. Milgrom (2000) shows, in an auction context, that the real threat to the existence of equilibria is associated with goods that are complements to some bidders but not to others. This is also characteristic of the kinds of complementarities found in labor markets with couples—the positions that a couple regards as complements are not typically regarded as complements by others in the labor force.

The largest lesson in all this is that design is important because markets don't always grow like weeds—some of them are hothouse orchids. Time and place have to be established, related goods need to be assembled, or related markets linked so that complementarities can be handled, incentive problems have to be overcome, etc. If game theory is going to be as important a part of design economics as it is a part of economic theory, we'll have to develop tools not just to develop conceptual insights from simple models, but also to understand how

to deal with these complications of real markets. These complications come in two kinds:

- complications in the strategic environment, and consequently in the possible outcomes, and the strategies available to the players; and
- complications in the behavior of real economic agents (who may not be the simple maximizers we are accustomed to studying in formal game theory, even in simple environments).

Computational methods will help us analyze games that may be too complex to solve analytically. Laboratory experiments will help inform us about how people will behave when confronted with these environments, both when they are inexperienced and as they gain experience. We'll need to learn from further experience how and when these tools can best be employed in different kinds of design efforts. And, since design involves anticipating how people will behave in novel environments, these concerns will make us want to deepen our understanding of *learning* in strategic environments. Successful designs must not only perform well in the long term, they must not crash and burn in the short term, and so models of learning with the ability to predict behavior in new environments will be a valuable addition to the designer's toolbox, alongside the more familiar, equilibrium analyses of long term behavior.³⁵

In closing, we can start to see the shape of the economics of design. A decade ago, as part of its centenary celebrations, the *Economic Journal* asked a number of economists to write an essay about what the next hundred years might bring in various areas of economics. The last paragraph of my essay (Roth (1991)), read as follows:

“... in the long term, the real test of our success will be not merely how well we understand the general principles that govern economic interactions, but how well we can bring this knowledge to bear on practical questions of microeconomic engineering. . . Just as chemical engineers are called upon not merely to understand the principles that govern chemical plants, but to design them, and just as physicians aim not merely to understand the biological causes of disease, but their treatment and prevention, a measure of the success of microeconomics will be the extent to which it becomes the source of practical advice, solidly grounded in well tested theory, on designing the institutions through which we interact with one another.”

A decade makes a difference. The steps we have taken make the difficulties in constructing an engineering science of design economics even more apparent. But there are grounds to feel optimistic that, in the matter of design, game theory and experimental and computational economics can be harnessed together for the work at hand.

Dept. of Economics, Harvard University, Littauer Center, Cambridge, MA 02138-3001, U.S.A.; al_roth@harvard.edu; <http://www.economics.harvard.edu/~aroth/alroth.html>

Manuscript received February, 2001; final revision received August, 2001.

³⁵ See, e.g., Roth and Erev (1995) and Erev and Roth (1998) for one train of thought in that direction.

REFERENCES

- ABDULKADIROGLU, A., AND T. SONMEZ (1998): "Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems," *Econometrica*, 66, 689–701.
- (1999): "House Allocation with Existing Tenants," *Journal of Economic Theory*, 88, 233–260.
- (2000): "School Choice: A Solution to the Student Assignment Problem," Mimeo, Columbia University.
- ARIELY, D., A. OCKENFELS, AND A. E. ROTH (2001): "An Experimental Analysis of Late-Bidding in Internet Auctions," in preparation, Harvard University.
- AUSUBEL, L. M. (2000): "System and Method for an Efficient Dynamic Auction for Multiple Objects," U.S. Patent Number 6,026,383, issued 15 Feb 2000.
- AUSUBEL, L. M., AND P. MILGROM (2001): "Ascending Auctions with Package Bidding," Working Paper, University of Maryland.
- AVERY, C., A. FAIRBANKS, AND R. ZECKHAUSER (2001): "What Worms for the Early Bird? Early Admissions at Selective Colleges," Working Paper, Harvard University.
- AVERY, C., C. JOLLS, R. POSNER, AND A. E. ROTH (2001): "The Market for Federal Judicial Law Clerks," *University of Chicago Law Review*, 68, 793–902.
- AYRES, I., AND P. CRAMTON (1996): "Deficit Reduction Through Diversity: How Affirmative Action at the FCC Increased Auction Competition," *Stanford Law Review*, 48, 761–815.
- BALDWIN, C. Y., AND S. BHATTACHARYYA (1991): "Choosing the Method of Sale: A Clinical Study of Conrail," *Journal of Financial Economics*, 30, 69–98.
- BALDWIN, C. Y., AND K. B. CLARK (2000): *Design Rules: The Power of Modularity*, Volume 1. Cambridge, MA: MIT Press.
- BALINSKI, M., AND T. SONMEZ (1999): "A Tale of Two Mechanisms: Student Placement," *Journal of Economic Theory*, 84, 73–94.
- BERNHEIM, B. D., AND M. WHINSTON (1986): "Menu Auctions, Resource Allocation and Economic Influence," *Quarterly Journal of Economics*, 101, 1–31.
- BINMORE, K., AND P. KLEMPERER (2002): "The Biggest Auction Ever: The Sale of the British 3G Telecom Licenses," *Economic Journal*, forthcoming.
- BLUM, Y., A. E. ROTH, AND U. G. ROTHBLUM (1997): "Vacancy Chains and Equilibration in Senior-Level Labor Markets," *Journal of Economic Theory*, 76, 362–411.
- BODIN, L., AND (RABBI) A. PANKEN (1999): "High Tech for a Higher Authority: The Placement of Graduating Rabbis From Hebrew Union College–Jewish Institute of Religion," Working Paper, Hebrew Union College–Jewish Institute of Religion.
- BULOW, J., AND J. ROBERTS (1989): "The Simple Economics of Optimal Auctions," *Journal of Political Economy*, 97, 1060–1090.
- CAPEN, E. C., R. V. CLAPP, AND W. M. CABELL (1971): "Competitive Bidding in High-Risk Situations," *Journal of Petroleum Technology*, 23, 641–653.
- CHEN, Y., AND T. SONMEZ (2002): "Improving Efficiency of On-Campus Housing: An Experimental Study," *American Economic Review*, forthcoming.
- CRAMTON, P. (1995): "Money Out of Thin Air: The Nationwide Narrowband PCS Auction," *Journal of Economics and Management Strategy*, 4, 267–343.
- (2000): "Lessons from the United States Spectrum Auctions," Prepared Testimony before the United States Senate Budget Committee, 10 February.
- CYBERNOMICS (2000): "An Experimental Comparison of the Simultaneous Multi-Round Auction and the CRA Combinatorial Auction," paper presented at the FCC Combinatorial Auction Conference, May 5–7, 2000.
- DE VRIES, S., AND R. VOHRA (2000): "Combinatorial Auctions: A Survey," Working Paper, Northwestern University.
- EREV, I., AND A. E. ROTH (1998): "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," *American Economic Review*, 88, 848–881.
- FEDERAL COMMUNICATIONS COMMISSION (1997): "The FCC Report to Congress on Spectrum Auctions," October 9, FCC 97-353.

- (2000): "Auction of Licenses in the 747-762 and 777-792 MHz Bands Scheduled for September 6, 2000: Comment Sought on Modifying the Simultaneous Multiple Round Auction Design to Allow Combinatorial (Package) Bidding," Public Notice DA00-1075, May 18.
- (2001): "Auction of Licenses in the 747-762 and 777-792 MHz Bands Scheduled for March 6, 2001: Modifications to the Calculation for Determining Minimum Acceptable Bids and the Provisions Concerning 'Last and Best Bids' and other Procedural Issues," Public Notice DA 01-12, January 5.
- FILER, R. K., AND J. HANOUSEK (1999): "Informational Content of Prices Set Using Excess Demand: The Natural Experiment of Czech Voucher Privatization," Mimeo, CERGE, Charles University.
- GALE, D., AND L. SHAPLEY (1962): "College Admissions and the Stability of Marriage," *American Mathematical Monthly*, 69, 9–15.
- GREEN, J. R., AND J.-J. LAFFONT (1979): *Incentives in Public Decision-Making*. Amsterdam: North-Holland.
- GRETHER, D. M., M. R. ISAAC, AND C. R. PLOTT (1981): "The Allocation of Landing Rights by Unanimity among Competitors," *American Economic Review*, 71, 166–171.
- GRIMM, V., F. RIEDEL, AND E. WOLFSTETTER (2001): "Low Price Equilibrium in Multi-Unit Auctions: The GSM Spectrum Auction in Germany," Working Paper, Humboldt University (Berlin).
- GROVES, T., AND J. O. LEDYARD (1977): "Optimal Allocation of Public Goods: A Solution to the Free Rider Problem," *Econometrica*, 45, 783–809.
- GUPTA, N., J. C. HAM, AND J. SVEJNAR (2000): "Priorities and Sequencing in Privatization: Theory and Evidence from the Czech Republic," Mimeo, William Davidson Institute, University of Michigan.
- HARUVY, E., A. E. ROTH, AND U. UNVER (2002): "The Dynamics of Law Clerk Matching: An Experimental and Computational Investigation of Some Proposed Market Designs," Working Paper, Harvard Business School.
- HATFIELD, J. N. (1935): "How Philadelphia Hospitals Appoint Interns," National Hospital Forum, February (2 pages, unnumbered).
- HIPPOCRATES (approx. 400 B.C.): *Of The Epidemics*, Book I, Section II, Translated by Francis Adams.
- HURWICZ, L. (1973): "The Design of Mechanisms for Resource Allocation," Richard T. Ely Lecture, *The American Economic Review*, Papers and Proceedings of the Eighty-fifth Annual Meeting of the American Economic Association, 63, 1–30.
- JEHIEL, P., AND B. MOLDOVANU (2001): "The European UMTS/IMT-2000 License Auctions," Mimeo, University of Mannheim.
- KAGEL, J. H. (1995): "Auctions: A Survey of Experimental Research," in *Handbook of Experimental Economics*, ed. by J. Kagel and A. Roth. Princeton: Princeton University Press, 501–585.
- KAGEL, J. H., AND D. LEVIN (1986): "The Winner's Curse and Public Information in Common Value Auctions," *American Economic Review*, 76, 894–920.
- (1999): "Common Value Auctions with Insider Information," *Econometrica*, 67, 1219–1238.
- KAGEL, J. H., AND A. E. ROTH (2000): "The Dynamics of Reorganization in Matching Markets: A Laboratory Experiment Motivated by a Natural Experiment," *Quarterly Journal of Economics*, 115, 201–235.
- KELLY, F., AND R. STEINBERG (2000): "A Combinatorial Auction with Multiple Winners for Universal Service," *Management Science*, 46, 586–596.
- KELSO, A. S., JR., AND V. P. CRAWFORD (1982): "Job Matching, Coalition Formation, and Gross Substitutes," *Econometrica*, 50, 1483–1504.
- KLEMPERER, P. (1998): "Auctions with Almost Common Values: The Wallet Game and Its Applications," *European Economic Review*, 42, 757–769.
- (2000): "Why Every Economist Should Learn Some Auction Theory," in *Advances in Economics and Econometrics: Invited Lecture to 8th World Congress of the Econometric Society*, ed. by M. Dewatripont, L. Hansen, and S. Turnovsky. Cambridge: Cambridge University Press.
- (2002a): "What Really Matters in Auction Design," *Journal of Economic Perspectives*, 16, 169–190.
- (2002b): "How Not to Run Auctions: The European 3G Telecom Auctions," *European Economic Review*, forthcoming.

- KNUTH, D. E. (1976): *Marriages Stables*. Montreal: Les Presses de l'Université de Montréal.
- LEDYARD, J. O. (1995): "Public Goods: A Survey of Experimental Research," in *The Handbook of Experimental Economics*, ed. by John H. Kagel and Alvin E. Roth. Princeton: Princeton University Press.
- LEDYARD, J. O., D. PORTER, AND A. RANGEL (1997): "Experiments Testing Multi Object Allocation Mechanisms," *Journal of Economics and Management Strategy*, 6, 639–675.
- LEDYARD, J., D. PORTER, AND R. WESSEN (2000): "A Market Based Mechanism for Allocating Space Shuttle Secondary Payload Priority," *Experimental Economics*, 2, 173–195.
- LUCKING-REILLY, D. (1999): "Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet," *American Economic Review*, 89, 1063–1080.
- MASKIN, E., AND J. RILEY (1984): "Optimal Auctions with Risk Averse Buyers," *Econometrica*, 52, 1473–1518.
- MCAFEE, R. P., AND J. McMILLAN (1996): "Analyzing the Airwaves Auction," *Journal of Economic Perspectives*, 10, 159–175.
- McMILLAN, J. (1994): "Selling Spectrum Rights," *Journal of Economic Perspectives*, 8, 145–162.
- McVITIE, D. G., AND L. B. WILSON (1970): "Stable Marriage Assignments for Unequal Sets," *BIT*, 10, 295–309.
- MILGROM, P. (2000): "Putting Auction Theory to Work: The Simultaneous Ascending Auction," *Journal of Political Economy*, 108, 245–272.
- MILGROM, P., AND R. WEBER (1982a): "The Value of Information in a Sealed-Bid Auction," *Journal of Mathematical Economics*, 10, 105–114.
- (1982b): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089–1122.
- MONGELL, S. J., AND A. E. ROTH (1986): "A Note on Job Matching with Budget Constraints," *Economics Letters*, 21, 135–138.
- (1991): "Sorority Rush as a Two-Sided Matching Mechanism," *American Economic Review*, 81, 441–464.
- MULLIN, F. J., AND J. M. STALNAKER (1952): "The Matching Plan for Internship Placement: A Report of the First Year's Experience," *Journal of Medical Education*, 27, 193–200.
- MYERSON, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–63.
- NAKAMURA, A., M. PERCY, P. DAVENPORT, R. FRASER, AND M. PIPER (1998): "The Genesis of CareerOwl: The Story of How SSHRC Funded University Research Led to an On-line Electronic Hiring Hall for Canadian Post Secondary Students and Alumni," <http://www.careerowl.ca/papers/TheGenesisofCareerOwl11.pdf>, accessed 6/28/00.
- OCKENFELS, A., AND A. E. ROTH (2001): "Late-Bidding in Second-Price Internet Auctions: Theory and Evidence," Working Paper, Harvard University.
- PLOTT, C. R. (1997): "Laboratory Experimental Testbeds: Application to the PCS Auction," *Journal of Economics and Management Strategy*, 6, 605–638.
- RASSENTI, S. J., V. L. SMITH, AND R. L. BULFIN (1982): "A Combinatorial Auction Mechanism for Airport Time Slot Allocation," *Rand Journal of Economics*, 13, 402–417.
- ROTH, A. E. (1982a): "The Economics of Matching: Stability and Incentives," *Mathematics of Operations Research*, 7, 617–628.
- (1982b): "Incentive Compatibility in a Market with Indivisible Goods," *Economics Letters*, 9, 127–132.
- (1984): "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory," *Journal of Political Economy*, 92, 991–1016.
- (1985): "The College Admissions Problem is not Equivalent to the Marriage Problem," *Journal of Economic Theory*, 36, 277–288.
- (1986): "On the Allocation of Residents to Rural Hospitals: A General Property of Two-Sided Matching Markets," *Econometrica*, 54, 425–427.
- (1990): "New Physicians: A Natural Experiment in Market Organization," *Science*, 250, 1524–1528.
- (1991): "A Natural Experiment in the Organization of Entry Level Labor Markets: Regional Markets for New Physicians and Surgeons in the U.K.," *American Economic Review*, 81, 415–440.

- (1996): “The NRMP as a Labor Market,” *Journal of the American Medical Association*, 275, 1054–1056.
- (2000): “Game Theory as a Tool for Market Design,” in *Game Practice: Contributions from Applied Game Theory*, ed. by Fioravante Patrone, Ignacio García-Jurado, and Stef Tijs. Dordrecht, The Netherlands: Kluwer, 7–18.
- ROTH, A. E., AND I. EREV (1995): “Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term,” *Games and Economic Behavior*, 8, 164–212.
- ROTH, A. E., AND A. OCKENFELS (2001): “Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet,” *American Economic Review*, forthcoming.
- ROTH, A. E., AND E. PERANSON (1997): “The Effects of the Change in the NRMP Matching Algorithm,” *Journal of the American Medical Association*, 278, 729–732.
- (1999): “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design,” *American Economic Review*, 89, 748–780.
- ROTH, A. E., AND ANDREW POSTLEWAITE (1977): “Weak Versus Strong Domination in a Market With Indivisible Goods,” *Journal of Mathematical Economics*, 4, 131–137.
- ROTH, A. E., AND M. SOTOMAYOR (1989): “The College Admissions Problem Revisited,” *Econometrica*, 57, 559–570.
- (1990): *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Econometric Society Monograph Series. Cambridge: Cambridge University Press.
- ROTH, A. E., AND JOHN H. VANDE VATE (1990): “Random Paths to Stability in Two-Sided Matching,” *Econometrica*, 58, 1475–1480.
- ROTH, A. E., AND X. XING (1994): “Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions,” *American Economic Review*, 84, 992–1044.
- (1997): “Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists,” *Journal of Political Economy*, 105, 284–329.
- ROTHKOPF, M. H., A. PEKEC, AND R. HARSTAD (1998): “Computationally Manageable Combinatorial Auctions,” *Management Science*, 44, 1131–1147.
- ROTHKOPF, M. H., T. J. TEISBERG, AND E. P. KAHN (1990): “Why Are Vickrey Auctions Rare?” *Journal of Political Economy*, 98, 94–109.
- SHAPLEY, L., AND H. SCARF (1974): “On Cores and Indivisibility,” *Journal of Mathematical Economics*, 1, 23–28.
- SHILLER, R. J. (1993): *Macro Markets: Creating Institutions for Managing Society's Largest Economic Risks*, Clarendon Lectures in Economics. Oxford: Oxford University Press.
- SONMEZ, T. (1997): “Manipulation via Capacities in Two-Sided Matching Markets,” *Journal of Economic Theory*, 77, 197–204.
- SVEJNAR, J., AND M. SINGER (1994): “Using Vouchers to Privatize an Economy: The Czech and Slovak Case,” *Economics of Transition*, 2, 43–69.
- UNVER, M. U. (2000a): “Backward Unraveling over Time: The Evolution of Strategic Behavior in the Entry-Level British Medical Labor Markets,” *Journal of Economic Dynamics and Control*, forthcoming.
- (2000b): “Computational and Experimental Analyses of Two-Sided Matching Markets,” Ph.D. Thesis, University of Pittsburgh.
- (2000c): “On the Survival of Some Unstable Two-sided Matching Mechanisms: A Laboratory Investigation,” Mimeo, University of Pittsburgh.
- VICKREY, W. (1961): “Counterspeculation, Auctions and Competitive Sealed Tenders,” *Journal of Finance*, 16, 8–37.
- WILSON, R. B. (1969): “Competitive Bidding with Disparate Information,” *Management Science*, 15, 446–448.
- (1992): “Strategic Analysis of Auctions,” in *Handbook of Game Theory*, ed. by R. J. Aumann and S. Hart. Amsterdam: North Holland, 228–279.
- (2002): “Architecture of Power Markets,” *Econometrica*, 70, 1299–1340.
- WOLFSTETTER, E. (2001): “The Swiss UMTS Spectrum Auction Flop: Bad Luck or Bad Design?” Working Paper, Humboldt University (Berlin).