

Market Culture: How Rules Governing Exploding Offers Affect Market Performance[†]

By MURIEL NIEDERLE AND ALVIN E. ROTH*

Many markets encounter difficulty maintaining a thick marketplace because they experience transactions made at dispersed times. To address such problems, many markets try to establish norms concerning when offers can be made, accepted, and rejected. Examining such markets suggests it is difficult to establish a thick market at an efficient time if firms can make exploding offers, and workers cannot renege on early commitments. Laboratory experiments allow us to isolate the effects of exploding offers and binding acceptances. In a simple experiment, we find inefficient early contracting when firms can make exploding offers and applicants' acceptances are binding. (JEL C91, D40, D81)

Many markets have experienced difficulty in maintaining a marketplace in which transactions are made at a time, and in a manner, that promotes efficiency. One obstacle in a number of markets has been the use of exploding offers. In such markets, firms make offers to which applicants must respond very quickly, and these markets often clear very fast, with firms as well as applicants having little opportunity to consider many alternatives (see, e.g., Roth and Xiaolin Xing (1994) for an account of a number of such “unraveled” markets).

Partly for these reasons, different markets have different rules, norms, and expectations about how and when offers will be made, accepted, and rejected. For example, in some labor markets, exploding offers are rare. Similarly, norms differ concerning the circumstances under which a candidate may honorably change her mind about an offer she has accepted.

These are issues of considerable practical importance. A case in point is the entry level labor market for American gastroenterologists, which, as we began the work reported here, was suffering from just the kind of unraveling mentioned above. Offers, often exploding offers, were made to aspiring gastroenterologists two years before

*Niederle: Department of Economics, Stanford University, Stanford, CA 94305 and National Bureau of Economic Research (NBER) (e-mail: niederle@stanford.edu); Roth: Department of Economics, Harvard Business School, Harvard University, Baker Library, Boston, MA 02163 (e-mail: aroth@hbs.edu). This work has been partially supported by grants from the National Science Foundation and (Niederle) the Sloan Foundation. Niederle thanks the Institute for Advanced Studies at Princeton and Harvard University for their hospitality. We are grateful to Uri Gneezy for important early conversations that helped shape the experiment reported here, and Debbie Proctor with whom we were lucky to collaborate on redesigning the policies governing the gastroenterology fellows market to organize a new centralized clearinghouse.

[†]To comment on this article in the online discussion forum, or to view additional materials, visit the articles page at: <http://www.aeaweb.org/articles.php?doi=10.1257/mic.1.2.199>.

employment would begin (while they were still engaged in their internal medicine residency). This had caused a number of problems for the market (see Niederle and Roth 2003b), and the gastroenterology professional organizations became interested in trying to reestablish a thick market that would seek to hire applicants only one year in advance. But many employers (fellowship program directors) were concerned that, even though they would like the market to be later and thicker, they could be harmed if they waited for an orderly market and their competitors continued to lock up the best candidates early with exploding offers. So the question arose, what policies could the gastroenterology organizations adopt to make it safe for employers to wait and participate in the market at an efficient time?

It turns out that in many markets, formal market organizations exist that seek to regulate how and when offers are made, accepted, and rejected (see Niederle and Roth forthcoming). Examining these markets together, in terms of their rules about offers and acceptances, suggests a pattern. Markets that have suffered from unraveling and have been successful in reversing it and restoring a later market have done so either by preventing firms from making exploding offers, or by establishing rules and conventions that allow candidates who accept early offers to change their minds under some circumstances.¹

This paper reports an experiment designed to investigate the interplay between the rules by which offers are made, accepted, and rejected, and the timing and efficiency of the resulting market. Our experimental results confirm the pattern suggested by our reading of the field evidence. In the laboratory, unraveled markets move toward later, more efficient markets if firms can only make open but not exploding offers. The same thing is observed even when firms can make exploding offers, as long as candidates can subsequently change their minds without too much cost if they receive more desirable offers later. However, unraveled markets remain inefficiently early when firms can make exploding offers and acceptances are binding.

One market in which a good deal of effort has been spent shaping and discussing the timing of offers and acceptances is the market for graduate students. The Council of Graduate Schools has, since the mid 1960s, attempted to establish norms concerning how graduate students are recruited. Over 350 American universities subscribe to its resolution,² which attempts to foster a market culture in which exploding offers are discouraged directly, and also indirectly, by being made less enforceable. Students are told that all universities have agreed that they need not respond to offers before April 15, but that if they are pressed to respond to some offer(s) before they have had the chance to adequately consider others, they have the option “to accept each one and then, by April 15, resign from all but one.” That is, the resolution suggests that a student who accepts an exploding offer with a deadline before April 15, but subsequently declines it before April 15, should not be thought of as behaving badly. This reduces the cost of renegeing on (and also of accepting) an exploding offer.

¹ See Christopher Avery et al. (2001, 2007); Avery, Andrew Fairbanks, and Richard Zeckhauser (2003); Guillaume R. Frechette, Roth, and M. Utku Ünver (2007); Susan Mongell and Roth 1991; Niederle, Deborah D. Proctor, and Roth (2006); Roth (1984, 1991); and Roth and Xing (1994, 1997) for descriptions of some of the markets in which unraveling has been an issue.

² CGS Resolution at www.cgsnet.org.

Similar concerns, and attempts to alter market culture, have played a large role in the organization of entry level labor markets for doctors, lawyers, and college admissions at the undergraduate level. Presently, law students who apply for appellate court clerkships are frequently given exploding offers, and are almost never reported to renege on them (Roth and Xing 1994; Avery et al. 2001, 2007).³

I. Experimental Matching Markets

The experimental markets consisted of five firms and six applicants. Firms have a fixed quality, from one to five, and applicants will eventually have a quality from one to six. The qualities of firms are simply their assigned ID number from one to five, the qualities of applicants are revealed over time. In each market a firm can hire one applicant and an applicant can work for one firm. A matched firm and applicant each earn the product of their qualities. Unmatched market participants earn zero.⁴

Each market lasts nine periods. In periods one, four, and seven each applicant receives an integer signal from one to ten (uniform independently and identically distributed). The quality of each applicant is determined in period seven through the relative ranking of the sum of their three signals. Note that while signals across periods are uncorrelated, because the final quality is determined by the sum of signals, the relative quality after two signals is heavily influenced by the quality of the first signal. That is the relative quality (the sum of signals) is correlated over time. The applicant with the highest sum receives a quality of six, the second highest a quality of five, the lowest a quality of one (ties are broken randomly).⁵ Firms see all the applicants' signals as they become available over time, but applicants only receive information about their own signals.⁶ Having three periods in which new information is revealed allows us to observe several "degrees" of inefficiency of early matching.

Each information state lasts for three periods in which firms can make offers, and applicants decide whether to accept or reject them. This helps avoid exogenously imposed congestion, which occurs when firms may run out of time to make offers they would have liked to make. In pure strategy equilibria, congestion does not occur by assumption. However, in laboratory markets, even a small amount of coordination failure would lead to congestion, and potential congestion in late periods would provide an additional reason for firms to make early offers, and an additional source of inefficiency. (In naturally occurring markets, congestion is common. See the discussion in Roth and Xing 1997; Roth 2008.) John H. Kagel and Roth (2000) report an experiment in which early matching arises in response to congestion. Three periods

³ However, not all markets that use exploding offers and binding agreements suffer from unraveling. For example, in the market for junior economists, many departments impose short deadlines, though the market as a whole has, for many years, largely opened at the winter meetings.

⁴ There are no monetary transfers between firms and applicants. They each receive the mutual gains of matching, equally.

⁵ In this environment, it is possible that, after two signals, the applicant of the highest quality, or the applicant of the lowest quality, though not both, can be deduced by the firms, but not by the applicants (who do not see others' signals).

⁶ This feature of the experimental environment is motivated by the situation in many markets, in which firms see a whole pool of applicants, but applicants may have difficulty knowing how they compare with other applicants.

of making offers should be sufficient in our environment, because once information about applicants' qualities is revealed, the assortative matching, even among a subset of firms and applicants, is apparent.⁷

We consider two types of offers that can be made by firms:

- An *exploding* offer is an offer that the applicant can only accept right away, i.e., in the same period in which it was made. If an exploding offer is not accepted immediately, it is rejected.
- An *open* offer is an offer the applicant can also hold (until period nine). That is, an applicant who receives an open offer may accept or reject it immediately, or may hold it, to accept or reject at a later period. An applicant must reject a held offer if he wishes to hold or accept another offer.⁸

In a given period, first, all firms decide what offers they will make. Each firm that is unmatched, and has no open offer being held by an applicant, may decide to make, at most, one offer. Then, each applicant learns of all offers he receives in that period before having to decide how to respond to each of them. If an applicant accepts the offer of a firm, the applicant and the firm are matched, and all market participants are informed about this. Offers are made in private, i.e., until they have been accepted, they are not announced to the other firms and workers.

We consider three environments, characterized by different rules governing offers and acceptances.

Treatment 1: Exploding and Open Offers. Each firm can decide whether to make each offer open or exploding. Once an applicant accepts an offer, the acceptance is binding, and firms cannot make subsequent offers to an applicant who has already accepted an offer.⁹

Treatment 2: Open Offers Only. Firms can only make open offers. Once an applicant accepts an offer, the acceptance is binding, and firms cannot make subsequent offers to an applicant who has already accepted an offer.

Treatment 3: Renege. In this treatment, firms can, again, decide whether to make open or exploding offers. However, an applicant who accepted an offer may still receive further offers. An applicant can renege on initial acceptances and accept a new offer at a cost of one point (that is subtracted from his final payment).¹⁰

⁷ While this avoids exogenously imposed congestion, congestion may develop endogenously, if applicants hold offers until late, or firms delay making offers. Roth and Axel Ockenfels (2002) discuss congestion arising endogenously in an auction market, in which bidders delay making bids until near the close of the auction. Markets that use centralized clearinghouses avoid congestion, and often have adopted the clearinghouse for this reason.

⁸ This is not an onerous constraint for the applicants, since they have strict, unchanging preferences over the firms. In our experimental environment, it reduces the cost to a firm of making an open offer, since it reduces the likelihood that an open offer will be held by an applicant who has no intention of taking it.

⁹ One can also think of the applicants' ability to make binding agreements as an agreement among firms to not make offers to applicants who accepted another firm's offer.

¹⁰ This fee is smaller than the minimum improvement from accepting a match with a higher quality firm, since matches pay each applicant the product of his quality and the quality of the firm to which he is matched.

All three treatments allow for a whole array of Bayesian-Nash equilibria. For example, in each treatment, there is an equilibrium in which, after the first period, all the firms are matched. For example, the following strategies constitute an equilibrium in which all firms are matched in period one.

Strategies of Firms.—Each firm i makes an open offer in period one to applicant i . A firm whose offer is rejected never makes another offer.

Strategies of Applicants.—Each applicant i , in period one, accepts an open offer from firm i and rejects any other offers (i.e., exploding offers, offers from other firms, and offers received in other periods).

These strategies constitute an equilibrium, as no firm has an incentive to deviate, given the strategies of applicants and vice versa. However, this equilibrium has the unattractive property of using weakly dominated strategies.

In each of the different conditions, when firms and applicants are risk neutral, there is a sequential equilibrium that yields late matching and the efficient outcome, however. But the possibility of late matching is less robust in the case of exploding offers and binding agreements than in the other two conditions.¹¹

In particular, since a strategy is a function that specifies an agent's actions at each of his information sets, we can (even in the renege condition) speak of an agent as adopting a strategy of "locking in" an offer at some point in any of the experimental treatments by accepting the offer, *and not later reneging on it even if a preferred offer is received*.¹² We can now state the following.

PROPOSITION: *In the open offer and the renege treatments, it is a weakly dominated strategy for applicants to lock in any but the highest quality firm (for open offers: highest quality unmatched firm) early, before period seven. In the exploding offer treatment, it is not a weakly dominated strategy for the applicant to accept an exploding offer early (before period seven), which means it is not a weakly dominated strategy to lock in a firm early.*

PROOF:

We first observe that in each information set of the exploding offer treatment (i.e., for any set of matched firms and workers) for any period $t < 7$, the strategy that calls for applicant i to accept an exploding offer from firm j is *not* a weakly dominated strategy.

Since there are more applicants than firms, there must be at least one other unmatched applicant at any information set in which i receives one or more offers (if i has received more than one offer there are at least as many other unmatched applicants as offers i has received). Suppose the strategies of the firms and other

¹¹ The proofs and statements of equilibrium strategies and beliefs are straightforward but tedious. They can be found in a technical addendum to the paper in the Web Appendix, or at www.stanford.edu/~niederle.

¹² Since an agent's strategy is a function from his information sets to his actions at those information sets, a player can adopt a strategy at the beginning of the game that determines his actions later in the game.

applicants are such that each unmatched firm of higher quality than the best offer i has received has made an offer to a different applicant, and that these applicants will accept these offers. Furthermore, firms that are rejected by i will each make their next offer to a distinct other applicant, and these offers will be accepted. Then i will not receive a better offer, and would do better to accept the best exploding offer rather than rejecting it. (This is sufficient to show that accepting an exploding offer is not a dominated strategy, but note further that the strategies of the other firms and applicants are not, themselves, absurd. Given that i 's offers come from the only firms that will remain unmatched, even highly rational rejected suitors may not renew their offer to i , if they, instead, refrain from making any other offers until period seven, when final qualities are known. The reason is that for every set of i 's signals, there is a positive chance that the other remaining applicant(s) will eventually be of higher quality than i in period seven. So if applicant i rejects its best offer in a period when all other firms will become matched, it cannot do better than to receive that offer again, and, with positive probability, it will not, so, it would have been strictly better off accepting the exploding offer earlier.)

We now note that in the open offer and renege treatments it is a weakly dominated strategy for a worker to lock in a firm before period seven if any higher quality firm remains unmatched, because in the event that such a firm made an offer later, it would be profitable to decline the initial offer. (Note that in the renege condition, even firms whose offers have been accepted can be regarded as unmatched in this sense since that acceptance could be renege upon.)

This suggests that the firms' behavior may be different in the different conditions. In the case of exploding offers and binding acceptances, (risk neutral) firms are always prepared to make early offers to applicants who have an expected quality that is higher than the one the firm receives in the stable match.¹³ In the open offer and renege condition, firms may not be prepared to make any early offers, as applicants would simply use the offer's option value, which can only reduce the firms' profits.

This difference may affect the robustness of the various equilibrium refinements to deviations from equilibrium, e.g., to random or other nonequilibrium behavior of some participants. For example, suppose there are some applicants who do not want to reject offers without a better offer in hand.

In the case of exploding offers and binding acceptances, such applicants will accept early (exploding) offers. This implies that firms will have an incentive to make early offers to applicants with an expected quality higher than the firms' equilibrium match, and there will be early transactions in such markets.

Furthermore these early matched applicants (and firms) impose a negative externality on (higher quality) firms that do not make an early offer, but that may want to hire these applicants later on. The reason is that these applicants irreversibly accepted an early offer from a lower quality firm. If this negative externality is high enough, high quality firms will be well advised to also start making early offers (to applicants whose expected value is *lower* than the quality they would receive in an assortative match in period seven), only to prevent high quality applicants from

¹³ In the late matching equilibrium, applicants reject such offers.

being captured early by low quality firms. These early offers are also accepted by rational applicants, which implies that an even bigger portion of the market moves early. So, a few applicants who do not reject early offers can affect the incentives of all firms and applicants, such that matching late is not an equilibrium anymore. By this mechanism, a few applicants can affect the timing of a much bigger portion of the market, and cause the market to transact early.¹⁴

In contrast, in the open offers and renege treatment, applicants who do not want to reject offers (without having another offer in hand) can merely hold on to the offer, and use it as an option. Such applicants would not differ from rational applicants who would also use the option value of an offer. So, in the open offers and renege treatments (and again, in contrast to the exploding offer treatment), even when faced with these kinds of applicants, firms do not benefit from making early offers, and indeed have a strict incentive *not* to make early offers. Furthermore, suppose there were a few applicants who would not only hold, but accept, early offers in the open offer treatment, or never renege upon accepting an early offer in the renege treatment, that is applicants who “lock in” an offer early. Even in this case, the incentives for firms to make early offers are considerably weaker than in the exploding offer case, since other applicants would still use firms’ early offers as an option. This makes equilibria with early matching less likely to be robust when firms can make only open offers or when acceptances are not binding.

Thus, each of the treatments has multiple equilibria, including a sequential equilibrium with efficient late matching. However, the late matching equilibrium appears to be less robust in the exploding offer (with binding agreements) treatment, than in the open offer and renege treatment. Similarly, early matching seems less robust in these two treatments than in the case of exploding offers and binding agreements.

We conducted seven sessions of the exploding offer treatment, and six sessions each of the open offer and renege treatments. Subjects participated in only one session, each of which consisted of 20 consecutive markets. Participants kept their role, firm or applicant, for the whole experiment, and for firms, also the firm ID and quality (from one to five). The experiment was conducted at the Harvard Business School, with students, using z-Tree software (Urs Fischbacher 2007). Firms 1 and 2 received an additional amount of \$5,¹⁵ and each participant received \$0.10 for each point earned. All participants received a \$10 show up fee.

II. Results of the Experiment: Unraveling, Thickness, and Efficiency

First we investigate whether different rules concerning offers and acceptances affect the timing of the market. How long do firms wait to extend offers, and, hence,

¹⁴ However, it cannot be the case that in a pure strategy equilibrium all firms match before the uncertainty is resolved, because, for example, the worst firm among the ones that make offers in the last period in which there are unmatched firms would have an incentive to not make the offer, and rather hire the better of the remaining two applicants once the uncertainty is resolved.

¹⁵ The instructions stated that some participants, already determined in advance, would receive some additional fixed payment (see instructions).




FIGURE 1

Note: For each treatment and blocks of markets, the time at which final offers were made (i.e., offers that eventually result in a match) is measured in the number of available signals about the applicants' quality.

how much information about an applicant's quality do firms have when extending offers that were eventually accepted?

A market is efficiently late if final offers (i.e., offers that were eventually accepted), are all made after period seven, once all the uncertainty about applicants' qualities is resolved. Figure 1 shows the timing of final offers for all treatments in the 20 markets. The timing is presented in terms of how many signals had been revealed before the offers were made. So, a value of one corresponds to offers made when only one signal was available (periods 1–3), two denotes offers made after two signals, i.e., offers made in periods 4–6, and three corresponds to the final quality of applicants being known, that is, offers made in periods 7–9. For the renege treatment, we only consider an offer to be final if it was accepted and not renege upon. The results are presented in blocks of five markets.

The first five markets in all treatments look similar, there is no significant difference in the average number of signals observed before firms make their final offers (that is, offers that result in a match).¹⁶ All significance tests in this paper are two-sided Mann-Whitney tests, and a data point is always the session average of the variable at hand. As such, the number of data points is 12 when we compare the treatment with open offers to the renege treatment, and 13 for any comparison that involves the exploding offer treatment (unless otherwise noted).

¹⁶ A two-sided Mann-Whitney U test on session averages (on the average number of signals observed when making a final offer in the first five markets) gives p values of 0.63 when comparing open to renege, 0.32 when comparing renege to exploding, and 0.32 when comparing open to exploding.




FIGURE 2. PERIOD IN WHICH FINAL OFFER WAS MADE

Note: For each treatment, in the last five markets (markets 16–20), the proportion of final offers that were made when one, two, or all three signals (and hence the final quality) about applicants' quality were available.

As participants gain experience, matches come to be made later in the open offer and renege treatments, but not in the exploding offer treatment. In the last five markets (markets 16–20), final offers in the exploding offer treatment are made with significantly fewer signals than in the renege treatment ($p < 0.01$) and the open offer treatment ($p < 0.01$),¹⁷ while the renege and the open treatment are not significantly different ($p = 0.63$). Thus, we see that open offers and the applicants' ability to renege help the market to defer the timing of contracts, even when the market begins with early appointments (as in the first five markets of Figure 1). In the analyses that follow, we'll show that the markets in the exploding offer condition are not only earlier than the others, but that they also exhibit the other classic symptoms of unraveling.

A. Transaction Times and Market Thickness

From now on, we focus our attention on the last five markets in each experiment, when conditions seem to have stabilized. We now investigate the timing of offers in more detail. Figure 2 shows, for each treatment in the last five markets, the percentage of offers that were made when one, two, or all three signals (three signals = final quality) about the applicants' quality were available.

Like Figure 1, Figure 2 shows that unraveling occurs when firms can make exploding offers and acceptances are binding. When firms are forced to make open offers, or when applicants can renege on their acceptance, the markets experience almost no unraveling. Furthermore, high quality firms 4 and 5 also make early final offers. In the exploding offer treatment, only about 20 percent of firms 4 and 5 make final

¹⁷ Furthermore, we can compute for each of the last five markets, in any session, the average number of signals the five firms had when making their final offer. The exploding offer market with the *highest* such number among the 35 markets (that is, the one where, on average, the five firms saw the most signals before making their final offer) has a lower number than the market with the *lowest* average number of signals used by the five firms in any of the last five markets, in any session, of either the open or the renege treatments.

offers after three signals are available (and 34 percent after one signal), compared to 88 percent and 90 percent in the open offer and renege treatment.

The fact that the open offer treatment and the exploding offer treatment are so different suggests that firms make exploding offers when they are given the opportunity to do so. Indeed, in the last five markets of the exploding offer treatment, only firm 1 makes an open offer in more than 10 percent of the markets (while firm 5 makes no open offers at all). Except for firms 3 and 4, every firm made an exploding offer in every one of the last five markets in each of the seven sessions of the exploding offer treatment (and firms 3 and 4 made exploding offers in 34 of these 35 markets). In the renege treatment, firms make somewhat more open offers, but the vast majority of firms (at least 67 percent) make an exploding offer in each of the last five markets.

Thus, when firms could make open and exploding offers, the majority of offers were exploding. Firms made use of their ability to make exploding offers to put pressure on applicants. However, this effect was more pronounced when acceptances were binding. When applicants can renege on their acceptance, the value of making an exploding offer is smaller, and firms made less use of that option.

Furthermore, in the exploding offer treatment, firms made many more early offers than in either the open or the renege treatment. When we consider the number of offers before period seven in each of the last five markets of each session, then, on average, sessions in the exploding offer treatment experienced about 44 early offers, compared to 12 in the open and renege treatment. Furthermore, the exploding offer session with the lowest number of such offers has more offers than the renege or open offer session with the most early offers. We hypothesized that differences are driven by the fact that applicants would not use weakly dominated strategies, such as locking in firms early that are not the highest firm (or the highest unmatched firm in the open offer treatment). Indeed, we find that applicants use weakly dominated strategies only rarely, about twice per session totaled over all last five markets in the open offer treatment, and five times in the renege treatment (out of about 70 and 45 instances in which they could have done so).

So far, we examined the offers and the timing of final offers that were eventually accepted, averaged across different markets. Now, we explore the timing at the level of each market. A transaction is made (and announced to the market) only when an offer is accepted. The following table shows, for each treatment, the timing of first acceptances in each of the last five markets (where we use only final acceptances that were not reneged upon for the renege treatment).

Markets with exploding offers not only experience early contracting on *average*, but individually 71 percent of markets have their first acceptance with only signal one available. All 35 markets (the last five markets of all seven sessions of the exploding offer condition) have their first acceptance before the final quality of applicants becomes available. In contrast, when firms can only make open offers, or when acceptances by applicants are not binding, 77 percent and 60 percent of the markets, respectively, experience their first acceptance only *after* all the uncertainty about applicants' quality is resolved.¹⁸

¹⁸ Even though markets with different rules concerning exploding offers experience a difference in timing of the first accepted offer, their *last* accepted offer occurs predominately in periods seven to nine. In the open and

TABLE 1

	1 Signal (1–3)	2 Signals (4–6)	3 Signals (7–9)
Exploding	0.71	0.29	0.00
Open	0.00	0.23	0.77
Renega	0.30	0.10	0.60

Note: For each treatment (in the last five markets) the proportion of markets whose first acceptance (which was not reneged upon in the renege treatment) was made when only one signal, two signals, or three signals (and hence the final quality) about applicants' quality were available.

A further piece of evidence for strategic causes of unraveling comes from examining the length of the market, i.e., the timing between the first and the last acceptance. In the exploding offer treatment, 32 out of 35 markets last for 5 or more periods (with 10 lasting exactly 5 periods). In the open offers and the renege treatment, the first final offers are made later, and the markets also last for a shorter time. In the renege treatment, 40 percent of the markets last for 5 or more periods, and for the open offer treatment, 23 percent last for 5 or more periods. The fact that the markets in the exploding offer treatment last so long shows that unraveling in this treatment is not caused simply by a desire to avoid congestion.

To put it another way, in the unraveled markets the offers are not just earlier, but also dispersed, making the market less thick in any period. Figure 3 shows the cumulative distribution of offers that were finally accepted, and makes clear that in the last three periods, when all information is available, the market is much thicker in the open and renege conditions than in the exploding offer condition with binding acceptances.

B. What Are the Costs of Unraveling?

From this point on, we eliminate from our main analysis one outlier session of the renege treatment. In that particular session there was one applicant who *never* accepted an offer. No other applicant, in any session of any treatment, behaved in this way. Such behavior may affect the timing of the market, showing that the results of Section IIA are quite robust. However, such behavior has a very strong impact on the efficiency of a market; who matches with whom. Effectively, this one session is one in which there were only five as opposed to six applicants. While we could convert efficiency and final matches to accommodate this difference, we opt for simplicity (and transparency) to disregard this session for the main analysis. In footnotes, we will show the analyses that include all renege sessions.

We have seen that the market unravels and lacks thickness when firms can make exploding offers and acceptances are binding. Now, we investigate the costs of unraveling.¹⁹ We evaluate the different treatments according to the quality of the

renege treatment, not a single market (of the last five markets) ends before period seven, and in the exploding offer treatment, 89 percent (31 out of 35) of the last five markets finish after period seven. Note that, in each treatment, the last firm to be unmatched has strong incentives to wait and see which of the two remaining applicants is of higher quality.

¹⁹ Of course, these costs would be different in different environments. One way to interpret the results in this section, therefore, is that they demonstrate that unraveling of transactions occurs even when it is quite costly in terms of the information lost.




FIGURE 3

Note: Cumulative acceptances within a market: For each treatment, for the last five markets, the cumulative proportion of final acceptances up to the end of each period.

resulting matches. How much use do firms make of the information about applicants that becomes available over time? We consider three benchmarks: assortative matching when only one signal is available, assortative matching with two signals, and assortative matching once all the uncertainty about applicants' qualities is resolved (the efficient outcome).

We calculate the value of the assortative match after the first signal by producing an assortative match between firms and applicants according to the applicants' first signal. We use the actual quality of applicants determined during the experiment to compute the value of this match.²⁰ The value of the assortative match with two signals is computed analogously. Let "1 Signal" and "2 Signals" denote the value of the assortative match after the first and second signal, respectively, and "Efficient" denote the value of the unique stable and efficient match once all signals are known. Figure 4 shows the averages across sessions and markets of $(\text{Actual Profits} - 1 \text{ Signal}) / (\text{Efficient} - 1 \text{ Signal})$, and $(\text{Actual Profits} - 2 \text{ Signals}) / (\text{Efficient} - 2 \text{ Signals})$. That is, it shows the relative gains of the outcome of the experimental market toward efficiency compared to assortative matching after one signal and after two signals.

Figure 4A shows that all treatments achieve, on average, a social surplus higher than assortative matching with one signal. Efficiency gains are significantly lower in the exploding offer treatment than in the open ($p = 0.06$) and the renege treatment ($p = 0.01$).²¹ Assortative matching based on two signals (the second signal becomes

²⁰ In case of ties in the first signal between two applicants, we take the average of the two possible outcomes.

²¹ When we include the outlier session in the renege treatment, the value of the proportion of gains from assortative matching, after one signal toward efficiency, is 0.69 (instead of 0.88). The exploding offer treatment still achieves significantly lower proportionate gains in efficiency, starting from the assortative match after one signal, than all the sessions in the renege treatment: $p = 0.03$.

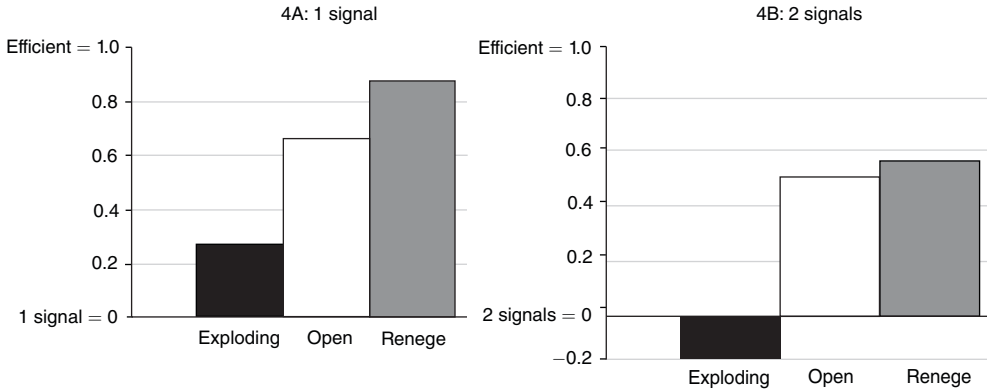


FIGURE 4. PROFITS COMPARED TO ASSORTATIVE MATCHING AFTER 1 SIGNAL AND 2 SIGNALS

Notes: 4A shows, for each treatment, the value of $(\text{Actual profits} - 1 \text{ Signal}) / (\text{Efficient} - 1 \text{ Signal})$ averaged across sessions. That is, it shows the relative gains of the actual match toward efficiency compared to assortative matching after one signal. Figure 4B shows the similar results for two signals.

available at period four) would have resulted in a higher efficiency than the exploding offer treatment, but both the open and renege treatment achieve higher efficiency levels than assortative matching after two signals. Furthermore, both the open offer and renege treatment achieve significantly higher efficiency gains than the exploding offer treatment ($p < 0.01$ and $p = 0.06$,²² respectively), while they are not significantly different from each other ($p = 0.47$). When we look at final absolute efficiency levels, the efficiency of the exploding offer treatment is significantly lower than the open offer treatment ($p = 0.03$) and the renege treatment ($p = 0.01$), while the renege and the open offer treatment are not significantly different ($p = 0.12$ using a two-sided Mann-Whitney U test with session averages).^{23, 24}

A different way to measure the functioning of a market is to count the number of “disruptive” blocking pairs. These are a firm and a worker, at least one of whom is matched, who are currently not matched to each other, but would both prefer to be so, instead of remaining with their current match (or being unmatched).²⁵ For the last 5 markets, the exploding offer treatment has, on average, in each market, 3 such

²² When we include all renege sessions, the value of the proportion of gains from assortative matching, after two signals, toward efficiency is 0.14 (as opposed to 0.56). The p -value when we include all renege sessions and compare them to the efficiency gains in the exploding offer treatment is $p = 0.25$.

²³ All treatments achieve high levels of efficiency (compared to the alternative of no firm being matched). The efficiency in the exploding offer treatment is 93 percent compared to 96 percent in the open offer treatment and 98 percent in the renege treatment. However, even a random allocation of the six applicants to the five firms achieves an efficiency of 75 percent. Average efficiency of assortative matching after one signal is 88 percent, and after two signals, it is 93 percent. We’ll see below that even small changes in efficiency can be associated with big changes in payoffs to differently ranked participants.

²⁴ When we include all renege sessions, the comparison with the exploding offer treatment has a p -value of 0.07, the comparison to the open offer treatment yields $p = 0.37$.

²⁵ These are blocking pairs that would disrupt the outcome of the market, if they had the chance. (Blocking pairs that only involve unmatched participants are much less disruptive, and in naturally occurring markets they often have a subsequent opportunity to match to one another.)




FIGURE 5. TREATMENTS AND EFFICIENT OUTCOME

Notes: The figure shows the average quality of the applicant each firm is matched to in the last five markets of each treatment. “UM” shows the quality of applicants who remain unmatched. “Efficient” shows, for each firm, the quality of the applicant in the unique stable and socially efficient match.

(firm, applicant) pairs, which is significantly higher than the about 1 such pair in the open ($p < 0.01$) and renege treatment ($p < 0.01$).^{26, 27}

C. Individual Level Consequences of Unraveling

We have seen the loss of social surplus when firms can make exploding offers that are binding, compared to when offers have to be open, or applicants can renege on their acceptance. Now, we investigate the value of the match for each applicant and each firm separately, for the last five markets of each treatment.

The Firms.—Figure 5 shows, for each firm, the average quality of the applicant they are matched to and the average quality of the applicant that remains unmatched.

The exploding offer treatment significantly lowers the payoff of the highest quality firm, firm 5, by 16 percent compared to the open ($p < 0.01$), and by 15 percent compared to the renege ($p = 0.05$) treatment. But the low quality firms, firm 2 and firm 1, achieve a significantly *higher* payoff in the exploding offer treatment compared to the open offer treatment ($p = 0.06$ and $p = 0.07$ for firm 2 and firm 1, respectively), and the renege treatment ($p < 0.01$ and $p = 0.099$).²⁸ In the exploding offer treatment, firm 2 gains 45 percent compared to the open offer, and 34

²⁶ The open offer and renege treatment do not differ significantly in the number of blocking pairs ($p = 0.22$). When we use all the renege sessions, the p -values are < 0.01 and 0.46 , when comparing it to the exploding and open treatment, respectively.

²⁷ The maximum feasible number of disruptive blocking pairs is 15 and is achieved by anti-assortative matching. Then the matched firm 5 generates five blocking pairs (four of which use matched applicants), firm 4 generates four, firm 3 generates three, firm 2 generates two, and firm 1 generates one.

²⁸ When we include all renege sessions, the p -value for firm 5, when comparing the exploding offer to the renege treatment, is $p = 0.02$, while it is $p = 0.03$ for firm 2, and $p = 0.05$ for firm 1.

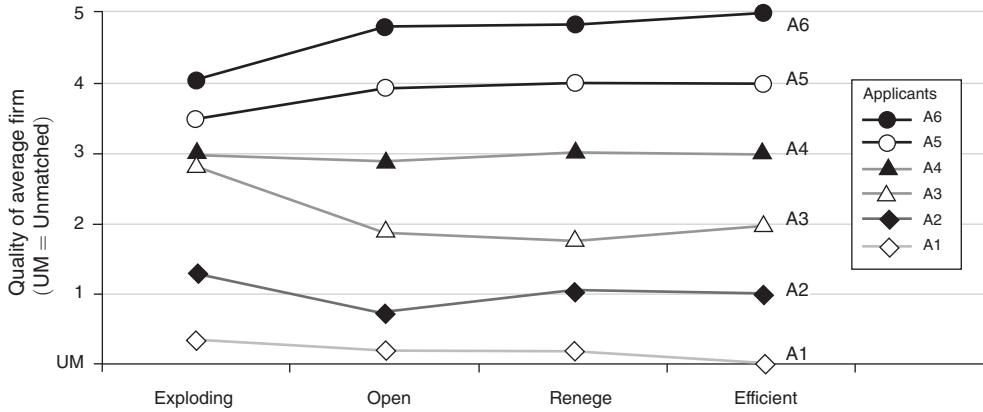


FIGURE 6. TREATMENTS AND EFFICIENT OUTCOME

Notes: The average quality of the firm the applicant is matched to in the last five markets of each treatment, where “UM” is unmatched. “Efficient” shows for each applicant the quality of the firm in the unique stable and socially efficient match.

percent compared to the renege treatment. The difference in the quality of applicants between firm 2 and firm 5 is 0.92 in the exploding offer treatment, which is significantly lower than in the open offer treatment, 3.07 ($p < 0.01$) and the renege treatment, 2.76 ($p = 0.03$).^{29, 30}

The Applicants.—Figure 6 shows, for each applicant, the average quality of the firm they are matched to in the last five markets, for each treatment.

As for the firms, it is the high quality applicants, applicants 6 and 5, that receive a significantly lower match in the exploding offer treatment than in the open offer treatment (by 10 percent and 11 percent with $p = 0.05$ and $p = 0.05$ for applicant 6 and 5, respectively) and the renege treatment (by 17 percent and 13 percent with $p = 0.02$ in both cases). And it is a medium quality applicant, applicant 3, who significantly gains from unraveling (by 49 percent and 57 percent with $p = 0.03$ and 0.04 compared to the open and renege treatment respectively).³¹ In all treatments higher quality applicants are hired by higher quality firms, on average. The difference in the quality of firms between applicant 3 and applicant 6 is 1.2 in the exploding offer treatment, which is significantly lower than in the open offer treatment, 2.9 ($p < 0.01$) and the renege treatment, 3.04 ($p < 0.01$).³²

²⁹ When we include all renege sessions, the p -value is 0.01.

³⁰ In all treatments, higher quality firms hire higher quality applicants, on average, but not in each session.

³¹ When we include all renege sessions, the p -values for applicants 6, 5, and 3 are $p = 0.01$, $p = 0.11$, and $p = 0.04$, respectively.

³² When we include all renege sessions, the p -value is less than 0.01.

The inefficient matchings in these markets are costly for the highest quality firms and applicants, while some lower quality firms and applicants tend to *gain* from early matches.³³

III. Implications for Market Design

A. Gastroenterology Fellows

From 1986 to the mid-1990s, the labor market for gastroenterology fellows (entry-level gastroenterologists) was organized through a centralized match that operated one year before employment began. The match collapsed following a shock to the market that caused fellowship programs and applicants to make early contracts before the match (C. Nicholas McKinney, Niederle, and Roth 2005). Following the collapse of the match, the market unraveled, and the hiring of fellows became increasingly early and dispersed with exploding offers, eventually moving to almost two years before employment began.³⁴ One consequence was that what had been a national market was replaced by more local markets, in which gastroenterologists were more likely to be recruited from the local pool of applicants (Niederle and Roth 2003b).

As these facts became increasingly well-known to gastroenterologists, there was a growing interest in reestablishing a match, to move the date of hiring nearer the date of employment, and to make the market thicker. However, doing so would require fellowship directors to refrain from early hiring prior to the first use of the match (as well as in subsequent years). With the memory of the collapse of the match clearly in mind, many program directors, who were themselves interested in participating in a well organized market at a later date, were concerned that their competitors would gain an advantage over them by making early exploding offers in an effort to “capture” promising candidates before the match.³⁵

In reply to this concern, we related the experience of the market for new PhDs, and shared the result of the experiment reported above. The American Gastroenterology Association subsequently published a policy statement containing a resolution modeled after that of the Council of Graduate Schools (Niederle, Proctor, and Roth 2006). A modified version of this was adopted by all four major gastroenterology professional organizations, regarding offers made before the (new) match. It states, in part:

The general spirit of this resolution is that each applicant should have an opportunity to consider all programs before making a decision and be able

³³ For this reason, it is sometimes difficult to achieve consensus in markets that suffer from unraveling about what steps, if any, to take to address the problem. See Niederle and Roth (2005) for a discussion of the situation facing gastroenterologists as they started to consider whether to organize a centralized clearinghouse.

³⁴ For the effects of the loss of the match see Niederle and Roth (2004, 2005); Niederle, Proctor, and Roth (2006). Niederle and Roth (2003a) show that other parts of the contract (apart from who matches to whom), such as wages, seem not to have been affected by the loss of the match.

³⁵ In June 2005, our colleague, Debbie Proctor, the gastroenterologist who took the lead in reorganizing the match, sent us an e-mail saying, in part “I’m answering 3–4 e-mails per day especially on this issue. ‘I want to make sure MY competition is in the match and that they don’t cheat.’ Well, this is another way of saying that if they cheat, then I will too! ... Have you ever seen this before? The distrust amongst program directors? I find it hard to believe that we are unique. Maybe this is [a] social science phenomenon?”

to participate in the Match. ... It therefore seeks to create rules that give both programs and applicants the confidence that applicants and positions will remain available to be filled through the Match and not withdrawn in advance of it.

This resolution addresses the issue that some applicants may be persuaded or coerced to make commitments prior to, or outside of, the Match ... Any applicant may participate in the matching process ... by ... resigning the accepted position if he/she wishes to submit a rank order list of programs ... The spirit of this resolution is to make it unprofitable for program directors to press applicants to accept early offers, and to give applicants an opportunity to consider all offers... (www.gastro.org/user-assets/Documents/04_Education_Training/Match/match_resolution.pdf)

The gastroenterology match for 2007 fellows was held June 21, 2006, and succeeded in attracting 121 of the 154 eligible fellowship programs (79 percent). Ninety-eight percent of the positions offered in the match were filled through the match, and so it appears that the gastroenterology community succeeded in changing the timing and thickness of the market.³⁶

Part of the success of the gastroenterology match was due to the fact that early movers couldn't impose a big negative externality on those who waited for the match, since pre-match exploding offers would not necessarily remove candidates from the market. This made it easier for everyone to wait for the match.

Not every attempt to facilitate a thick market by constructing an orderly marketplace takes this approach, and the markets for new orthopedic surgeons and law clerks are instructive in their similarities and differences from the market for gastroenterologists.

B. Orthopedic Surgeons and Law Clerks

The gastroenterology professional organizations turned out not to be able to prevent fellowship directors from making exploding offers, but they were able to empower applicants to decline pre-match offers that they had previously accepted. In contrast, orthopedic surgery has an unraveled market and is considering a match, and it is not clear whether young surgeons can effectively be empowered in this way. That is, it may be that young orthopedic surgeons are launching their careers in such a sufficiently small market with repeated interactions that there is nothing the professional organizations can do that will effectively allow young surgeons to renege on promises extracted from them by senior surgeons. However, at least some of the orthopedic surgery professional organizations feel that (unlike the gastroenterology organizations) they may be able to effectively require fellowship directors to stop making exploding offers (Christopher D. Harner et al. 2008). So it may be that in

³⁶The second match, in June 2007, attracted even greater participation (Niederle, Proctor, and Roth 2008).

orthopedic surgery, the option of requiring open offers will be effective, even if no policy allowing students to renege on offers would be.³⁷

The market for federal court clerks suffers the worst of both worlds, since the judicial organizations can neither restrain judges from making exploding offers nor effectively empower law students to renege on even a verbal contract with a judge. This market has periodically suffered serious market failures in which clerks are sometimes hired two years in advance, i.e., at the beginning of the second year of law school (Roth and Xing 1994; Avery et al. 2001). In the most recent effort to repair the market, beginning in March of 2002, a large majority of Federal appellate judges voted to hire only after Labor Day 2003, and to hire third year (and not second year) law students. That is, like the gastroenterology fellows market, the law clerk market sought to replace early and diffuse hiring with hiring at a specified time one year later than had become customary.

While the first four years following the new proposal succeeded in moving the hiring date, it appears that there has been substantial cheating on the precise opening time of the market, and an increase in offers that explode immediately (Avery et al. 2007). The modal respondent to the student survey, reported in Avery et al. (2007), accepted an offer on the first day that offers were allowed. So, although the market has so far largely succeeded in controlling the dates of appointment, for many participants the market remains very thin; applicants can often consider no more than a single offer.

Given that law students are almost never reported to renege on promises made to senior Federal judges (Avery et al. 2001), the results of the present experiment give us a clear prediction. The problem of early contracting will not be solved by the current attempt to change the date of the market while leaving the market rules and customs intact. Rather, more fundamental changes in the market culture of judges and law clerks will be needed.³⁸

C. Exploding Offers and Congestion

The main result of our experiment is that late, efficient hiring is fostered by open offers by firms, or the flexibility for applicants to change their minds about early acceptances, either of which prevents firms from permanently “capturing” applicants early through exploding offers. Note that this works because our market is uncongested. Firms that wait until the efficient time to make offers can make as many offers as they might need, so they don’t run an unacceptable risk of remaining unmatched.

In an environment in which congestion is a large problem (i.e., an environment in which many offers would have to be made to clear the market, relative to the time available), exploding offers have the potentially positive effect of combating congestion by increasing the number of offers a firm can make, and some early

³⁷ We know of another market (for clinical neuropsychologists) in which a policy of allowing applicants to renege on early exploding offers was rejected as being simply repugnant (Roth 2007).

³⁸ On this point, see also the discussion in Ernan Haruvy, Roth, and Ünver (2006), concerning other aspects of the market culture in the law clerk market. That experiment focuses on a different aspect of offers in the law clerk market, namely that applicants may not feel able to refuse the first offer they receive.

matching may be important to allow the market to clear.³⁹ However, exploding offers are also often used strategically to make the market thin for applicants, and, in many cases, this leads to markets that operate well before employment begins. That is, while exploding offers can sometimes be an efficiency-enhancing response to congestion, they lead to the problems discussed in this paper, which often leads to unraveling beyond what seems warranted by congestion problems.⁴⁰

This helps explain why the rules for offers and acceptances in the market for graduate students work well for that market. Graduate programs admit relatively large numbers of students, and are not too sensitive to the precise numbers in the entering class. So graduate programs can make many offers at once, and the successful control of the dates of appointment via policies that effectively prevent exploding offers goes a long way toward establishing the orderly and stable decentralized market that has persisted for many years.⁴¹

This is also why such a policy for discouraging early offers worked well in the market for gastroenterology fellows. By using a match, they solved the congestion problem. The remaining problem was whether gastroenterology programs and fellows could safely wait for the match. This is exactly the problem that our experimental results indicate can be solved with appropriate policies about offers and acceptances.

This also helps explain why, before the adoption of a match, policies encouraging open offers at a fixed time worked less well for medical residencies and fellowships. Those are markets in which it is hard to tolerate uncertainty about the number of residents and fellows. Accreditation requirements generally prevent fellowship and residency programs from offering even one more than their target number of positions. So residency programs had to make sequential offers. In the 1940s (for residency programs, and in the 1980s for fellowship programs), policies were tried that attempted to encourage program directors to make offers that would remain open for a longer period of time. However, these markets experienced large congestion problems, and were eventually replaced by a centralized match (Roth 1984, 2003). The policies that worked well in our experiments did not succeed in the absence of a solution to the congestion problem.

IV. Discussion

Market design has often been concerned with centralized marketplaces such as auctions or clearinghouses (Paul Milgrom 2004, and Roth 2002). But even centralized marketplaces need to make it safe for agents to participate, rather than transacting in a decentralized way outside of the marketplace. For the gastroenterology labor market, this involved adjusting the rules of the decentralized market.

³⁹ Previous experiments have looked at matching in congested decentralized markets in which early, exploding offers arise, in part, as firms seek to avoid being unmatched (see Kagel and Roth 2000).

⁴⁰ Note that in our experiment too, in the treatment with exploding offers and binding agreements, many markets lasted for more than five periods (the number of firms to be matched) and the first match often occurred before period four, that is, with more than six periods left in the market. Our experiment is designed to show that exploding offers with binding agreements can lead to unraveling even in the absence of congestion.

⁴¹ Despite the congestion around April 15, which may make the size of entering classes uncertain.

A striking feature of many markets is that market participants spend a good deal of effort addressing how and when offers are made, accepted, and rejected. The formal and informal rules, customs, and norms that result are a critical element of the widely different ways that the matching processes in these markets are organized. In some markets, exploding offers are the norm, and applicants for positions find themselves faced with offers that must be accepted or rejected before other offers may be considered or even received. In others, exploding offers are discouraged or made more difficult to use to advantage.

Observation of these markets suggests the hypothesis that exploding offers with binding acceptances are potent facilitators of inefficiently early matching. But, because these many markets are also quite different from one another in other respects than their norms concerning offers and acceptances, it is natural to look to the laboratory to isolate the effects of different rules and customs concerning exploding offers.

The results of our experiment confirm the hypothesis motivated by the natural markets, that exploding offers, together with binding acceptances, make it difficult to avoid early and dispersed transactions. In the same environment, late and thick markets are achieved by either allowing only open offers, or allowing applicants to renege on early acceptances.

The problem facing applicants who receive early exploding offers when acceptances are binding is that, to reject such an offer, an applicant must hope for a better offer later in the market, hence, an applicant must hope not only that he will be highly ranked in the later market, but also that high quality firms will not fill their positions early. Once some applicants are ready to accept early offers, they impose a negative externality on high quality firms, making the whole market move early. We have seen in other environments that to successfully halt unraveling, a major factor is that applicants must be willing to reject early offers (Kagel and Roth 2000; McKinney, Niederle, and Roth 2005; Ünver 2001; Haruvy, Roth, and Ünver 2006). When offers are open, or when applicants can renege on their acceptances, then the market does not have to depend on applicants' willingness to reject early offers to have most of its transactions happen efficiently late.

More generally, market design is about the details of how markets work, and so designs for different markets may sometimes be quite different. But it appears that there may be quite a bit of generality across markets about some kinds of details that are important. In particular, the details of the sometimes informal rules and practices governing how offers are made, compared, and accepted or rejected, can be critical elements of a market's design.

REFERENCES

- Avery, Christopher, Andrew Fairbanks, and Richard Zeckhauser. 2003. *The Early Admissions Game: Joining the Elite*. Cambridge, MA: Harvard University Press.
- Avery, Christopher, Christine Jolls, Richard A. Posner, and Alvin E. Roth. 2001. "The Market for Federal Judicial Law Clerks." *University of Chicago Law Review*, 68(3): 793–902.
- Avery, Christopher, Christine Jolls, Richard A. Posner, and Alvin E. Roth. 2007. "The New Market for Federal Judicial Law Clerks." *University of Chicago Law Review*, 74(2): 447–86.
- Fischbacher, Urs. 2007. "Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics*, 10(2): 171–78.

- Fréchette, Guillaume R., Alvin E. Roth, and M. Utku Ünver.** 2007. "Unraveling Yields Inefficient Matchings: Evidence from Post-Season College Football Bowls." *RAND Journal of Economics*, 38(4): 967–82.
- Harner, Christopher D., Anil S. Ranawat, Muriel Niederle, Alvin E. Roth, Peter J. Stern, Shepard R. Hurwitz, William N. Levine, G. Paul DeRosa, and Serena S. Hu.** 2008. "Current State of Fellowship Hiring: Is a Universal Match Necessary? Is it Possible?" *Journal of Bone and Joint Surgery*, 90(6): 1375–84.
- Haruvy, Ernan, Alvin E. Roth, and M. Utku Ünver.** 2006. "The Dynamics of Law Clerk Matching: An Experimental and Computational Investigation of Proposals for Reform of the Market." *Journal of Economic Dynamics and Control*, 30(3): 457–86.
- Kagel, John H., and Alvin E. Roth.** 2000. "The Dynamics of Reorganization in Matching Markets: A Laboratory Experiment Motivated by a Natural Experiment." *Quarterly Journal of Economics*, 115(1): 201–35.
- McKinney, C. Nicholas, Muriel Niederle, and Alvin E. Roth.** 2005. "The Collapse of a Medical Labor Clearinghouse (and Why Such Failures Are Rare)." *American Economic Review*, 95(3): 878–89.
- Milgrom, Paul.** 2004. *Putting Auction Theory to Work*. Cambridge, UK: Cambridge University Press.
- Mongell, Susan, and Alvin E. Roth.** 1991. "Sorority Rush as a Two-Sided Matching Mechanism." *American Economic Review*, 81(3): 441–64.
- Niederle, Muriel, Deborah D. Proctor, and Alvin E. Roth.** 2006. "What Will Be Needed for the New Gastroenterology Fellowship Match to Succeed?" *Gastroenterology*, 130(1): 218–24.
- Niederle, Muriel, Deborah D. Proctor, Alvin E. Roth.** 2008. "The Gastroenterology Fellowship Match—The First Two Years." *Gastroenterology*, 135(2): 344–46.
- Niederle, Muriel, and Alvin E. Roth.** Forthcoming. "The Effects of a Centralized Clearinghouse on Job Placement, Wages, and Hiring Practices." In *Studies of Labor Market Intermediation*, ed. David Autor. Chicago: University of Chicago Press.
- Niederle, Muriel, and Alvin E. Roth.** 2003a. "Relationship Between Wages and Presence of a Match in Medical Fellowships." *Journal of the American Medical Association*, 290(9): 1153–54.
- Niederle, Muriel, and Alvin E. Roth.** 2003b. "Unraveling Reduces Mobility in a Labor Market: Gastroenterology with and without a Centralized Match." *Journal of Political Economy*, 111(6): 1342–52.
- Niederle, Muriel, and Alvin E. Roth.** 2004. "The Gastroenterology Fellowship Match: How it Failed and Why it Could Succeed Once Again." *Gastroenterology*, 127(2): 658–66.
- Niederle, Muriel, and Alvin E. Roth.** 2005. "The Gastroenterology Fellowship Market: Should There Be a Match?" *American Economic Review*, 95(2): 372–75.
- Roth, Alvin E.** 1984. "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory." *Journal of Political Economy*, 92(6): 991–1016.
- Roth, Alvin E.** 1991. "A Natural Experiment in the Organization of Entry-Level Labor Markets: Regional Markets for New Physicians and Surgeons in the United Kingdom." *American Economic Review*, 81(3): 415–40.
- Roth, Alvin E.** 2002. "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics." *Econometrica*, 70(4): 1341–78.
- Roth, Alvin E.** 2003. "The Origins, History, and Design of the Resident Match." *Journal of the American Medical Association*, 289(7): 909–12.
- Roth, Alvin E.** 2007. "Repugnance as a Constraint on Markets." *Journal of Economic Perspectives*, 21(3): 37–58.
- Roth, Alvin E.** 2008. "What Have We Learned from Market Design?" *Economic Journal*, 118(527): 285–310.
- Roth, Alvin E., and Axel Ockenfels.** 2002. "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet." *American Economic Review*, 92(4): 1093–1103.
- Roth, Alvin E., and Xiaolin Xing.** 1994. "Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions." *American Economic Review*, 84(4): 992–1044.
- Roth, Alvin E., and Xiaolin Xing.** 1997. "Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists." *Journal of Political Economy*, 105(2): 284–329.
- Ünver, M. Utku.** 2001. "Backward Unraveling over Time: The Evolution of Strategic Behavior in the Entry Level British Medical Labor Markets." *Journal of Economic Dynamics and Control*, 25(6–7): 1039–80.