

## Chapter 10. Market Failure Due to “Unraveling”: Lemons and Matching Markets

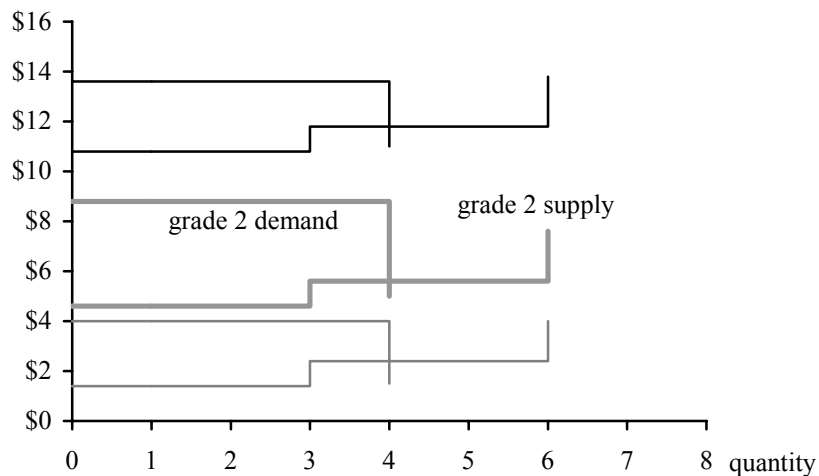
When buyers can observe both price and product quality prior to purchase, there is pressure on sellers to provide good qualities at reasonable prices. In contrast, a seller may be tempted to cut quality when it is not observed by buyers in advance, and the result will be disappointed buyers who are hesitant to pay the high prices needed to cover the costs to the high-quality seller. As a consequence, sellers may be forced to cut both price and quality in an unraveling process. The "lemons market" terminology is due to George Akerlof (1970), who explained why there is no market for a new car that has just been purchased from the dealer and offered for resale, since prospective buyers will assume it is a lemon. Akerlof showed how asymmetric information can cause quality to deteriorate to such low levels that the market may fail to exist. Unraveling has also been observed in labor markets, but for a different reason. At certain times in the past, job offers to students in professional schools have been made earlier and earlier by competing employers, and this process caused problems when employment matches were made years in advance of graduation. This chapter reviews experiments that have been used to study the unraveling process and the types of market mechanisms that improve performance by stopping that process. The lemons market game can be administered using the instructions provided or with the *Veconlab* program. This is a game in which buyers shop in sequence, so it should not be done with a large class, unless people are divided into groups.

### **I. Endogenous Product Quality**

The markets considered in this section have the feature that sellers can choose the qualities of their products. A high quality good is more costly to produce, but it is worth more to buyers. These costs and benefits raise the issue of whether or not there is an optimal quality. Most markets have the property that buyers are diverse in their willingness to pay for quality increases, so there will be a variety of different quality levels being sold at different prices. Even in this case, sellers who offer high-cost, high-quality items may face a temptation to cut quality slightly, especially if buyers cannot observe quality prior to purchase. Quality might be maintained, even with asymmetric information about quality, if sellers can acquire and maintain reputations reported or signaled by warranties and return policies. Before considering the effects of such policies, it is useful to examine how markets might fail if buyers cannot observe quality in advance.

For simplicity, consider a case where all buyers demand at most one unit each and have the same preferences for quality, which is represented by a numerical grade,  $g$ . The maximum willingness to pay for a unit of the commodity will be an increasing function of the grade,  $V(g)$ . Similarly, let the cost per unit be an increasing function,  $C(g)$ , of the grade. The net value for each grade  $g$  is the difference:  $V(g) - C(g)$ . The optimal grade maximizes this difference. The optimal grade may not be the highest feasible grade. For example, it is often prohibitively expensive to remove all impurities or reduce the risk of product failure to zero. The important behavioral issue in this type of market is the extent to which competition will force quality to near-optimal levels.

The main features of the market can be illustrated with a simple classroom experiment with 3 quality grades. Each buyer only demands one unit of the commodity, and buyers have identical valuations. The value of the commodity depends on the quality grade: \$4.00 for grade 1, \$8.80 for grade 2, and \$13.60 for grade 3. With four buyers and a grade of 2, for example, the market demand would be vertical at a quantity of 4 units for any price below \$8.80, as shown by the grade 2 demand line in the middle part of Figure 10.1. The other demand lines are similar, with cutoff prices of \$4.00 and \$13.60 for grades 1 and 3 respectively.



**Figure 10.1. Demand and Supply Arrays by Grade**

Market supply is determined by the costs given to sellers. Each seller has a capacity to produce two units, with the cost of the second unit being \$1 higher than the cost of the first unit. For a grade of 1, the costs for the first and second units produced are \$1.40 and \$2.40 respectively, so each individual seller's supply curve would have two steps, before becoming perfectly inelastic at two units for prices

above \$2.40. With three identical sellers, the market supply will also have two steps with three units on each step. The market supply for grade 1 is located in the lower part of Figure 10.1, and it crosses the demand for grade 1 at a price of \$2.40. The supply and demand curves for the other grades are shown above those for grade 1. The total surplus (for consumers and producers) corresponds to the area between the supply and demand curves for a given grade. It is apparent from Figure 10.1 that the sum of consumer and producer surplus is maximized at a grade of 2.

The results of a classroom experiment from Holt and Sherman (1999) are shown in Table 10.1. This was run with the instructions from the Appendix, and students were divided into groups and given roles corresponding to 4 buyers and 3 sellers. In the first three rounds, sellers chose price and grade and wrote them on record sheets. When all had finished, these prices and grades were written on the blackboard, and buyers were selected one by one in random order to make purchases. Notice that two of the three quality choices in the first round were at the maximum level. Each of these high-quality sellers sold a single unit, but the seller who offered a grade of 2 sold 2 units and earned more. The buyers preferred the grade 2 good since it provided more surplus relative to the price charged. All three sellers had settled on the optimal grade of 2 by the third round, and the common price of \$5.60 was approximately equal to the competitive level determined by the intersection of supply and demand at this grade.

Table 10.1 Results from a Classroom Experiment  
Source: Holt and Sherman(1999)

	<i>Seller 1</i>	<i>Seller 2</i>	<i>Seller 3</i>
Period 1 (full information)	\$11.50 grade 3 sold 1	\$6.00 grade 2 sold 2	\$12.00 grade 3 sold 1
Period 2 (full information)	\$5.75 grade 2 sold 2	\$5.50 grade 2 sold 1	\$1.90 grade 1 sold 1
Period 3 (full information)	\$5.65 grade 2 sold 1	\$5.60 grade 2 sold 2	\$5.60 grade 2 sold 1
Period 4 (only price information)	\$2.40 grade 1 sold 1	\$5.60 grade 2 sold 1	\$2.40 grade 1 sold 2
Period 5 (only price information)	\$2.40 grade 1 sold 1	\$1.65 grade 1 sold 1	\$5.50 grade 1 sold 2

The three “full-information” rounds were followed by two rounds in which sellers selected price and grade as before, but only the price was written on the board for buyers to see while shopping. Two of the sellers immediately cut grade to 1 in round 4, although they also cut the price, so buyers would be tipped off if they interpreted a low price as a signal of a low grade. In the final round, seller 3 offered a low grade of 1 at a price of \$5.50, which had been the going price for a grade of 2. The buyer who purchased from this seller must have anticipated a grade of 2, and this buyer lost money in the round. This pattern is similar to the quality unraveling observed by Holt and Sherman (1990) in a more complex market setting.

As noted by Akerlof (1970), market failure can occur even if sellers are not able to choose the quality grades of the products they sell, as long as there is asymmetric information in the sense that buyers do not observe quality prior to purchase. This market failure with exogenous quality grades may happen in a process of “adverse selection” in which sellers with high quality items drop out, since they do not expect to obtain high prices from buyers who cannot observe this quality. As the high-quality sellers drop out, the average quality in the market falls, and hence buyers’ willingness to pay declines, which drives additional sellers with higher-than-average quality goods out of the market. The end result of this “unraveling” process may be a situation in which no trade occurs, or a complete market failure. The key characteristic of this process is that the quality is not observed by traders on one side of the market.

There have been a number of experimental studies of the effects of practices and institutions that alleviate the lemons market problem. Lynch et al. (1986) show that performance in double auctions with information asymmetries can be improved with certain types of warranties, requirements for truthful advertising, etc. DeJong, Forsythe, and Lundholm (1985) allowed sellers to make price and quality representations, but the quality representation did not have to be accurate, and the buyer had imperfect information about quality even after using the product. In this experiment, the sellers could acquire reputations for providing higher quality based on buyer experience. This process of reputation building prevented the markets from collapsing to the lowest quality. Miller and Plott (1985) report experiments in which sellers can make costly decisions that “signal” high quality, which might prevent quality deterioration. Signaling models are discussed in chapter 33.

## **II. Clearinghouse Mechanisms and “Unraveling” in Labor Markets**

Just as firms might cut prices and quality grades, resulting in a low-grade outcome that is bad for all concerned, it is possible for a series of actions and reactions in labor markets to cause costly market failures. This kind of unraveling can occur when employers seek to beat the competition and congestion at graduation times for professional schools by making early offers, even though

early offers entail some costs on both sides of the market. In the market for federal appellate court clerks in the U.S., for example, judges who compete with each other for the best law students began to make earlier and earlier offers. This process continued until offers were being extended and accepted two years prior to graduation, on the basis of first-year grades in law school (Roth, 1984). An alternative, which is used in many medical labor markets, is to establish a centralized clearinghouse in which a computer-based job assignment is determined from ranked preferences submitted by both prospective employers and employees. Such a system was adopted in the U.S. in the 1950s for medical residents, and a modified version of this system is still in use (Roth, 1990). The success of the U.S. National Resident Matching Program was noticed in the U.K., where unraveling and early offers were prevalent in the 1950s and early 1960s. The computer matching algorithms subsequently adopted in the U.K. varied from place to place. Some of these systems have survived and prospered, whereas others have failed as unraveling did not diminish. This section describes the Kagel and Roth (2000) clearinghouse experiment, which compares the performance of the two prominent types of matching mechanisms in use.

Both matching algorithms begin by having each prospective employee (“worker”) submit a ranked list of employers (“firms”), with number 1 being the person’s top preference, etc. Similarly, each firm submits a ranked list of available workers. Under one system, which will be called the *priority product* system, the attractiveness of a specific worker-firm match is the product of the rankings from each side for the other. This product is 4, for example, if the worker is the firm’s second choice and the firm is the worker’s second choice. The process begins by enacting the match with the lowest priority product. The worker and firm for this match are then removed from the pool, with the lowest priority product for the remaining workers and firms determining the next match, etc. With equal numbers of workers and positions, this will result in no unmatched workers or firms.

The other matching method is known as the *deferred acceptance* system, since at each stage, workers keep their best offer on the table and reject others, which firms then send to workers lower on their lists. The computerized process is done mechanically on the basis of submitted rankings, but it is useful to think of it as beginning with firms sending an offer to the worker at the top of their lists. Some workers will get multiple offers, and in this case they will keep the offer that they have ranked the highest, rejecting the others. If a firm has a rejected offer, it can submit it in the next stage to the worker highest on its list who has not previously rejected its offer. A worker who keeps an offer on the table in one stage may reject it later if another offer arrives that is higher on the worker’s preference list. Thus each worker’s satisfaction with the match on the table will

not go down. The computer stops this process after a stage in which there are no rejections to be implemented on the basis of the submitted rankings.

The deferred acceptance process produces an outcome that is *stable* in the sense that it would be impossible to find a worker and firm who would prefer to be matched with each other instead of keeping their assigned partners. To see this, note that if a firm prefers a particular worker to the person with whom it is matched, then it must have previously proposed a match to that preferred worker. The worker who rejected that offer previously would reject it again since the value of the worker's match does not go down from one stage to the next. Thus it is a Nash equilibrium for workers and firms to submit their true rankings under this deferred acceptance system. In contrast, the priority product system does not always produce a stable outcome, as will be shown next by looking at an example.

The setup in the experiment involved 6 workers and 6 firms. Three of the workers were of high productivity, and a match with them was worth about \$15 to each firm, regardless of its own productivity. In contrast, a firm who was matched with a low productivity worker earned \$5 on average. Similarly, three of the firms were high productivity firms in the sense that a worker matched with one of them would earn \$15 on average; matches with one of the other three firms only provided the worker with \$5 average earnings. The actual earnings for each match were equal to the average, plus or minus a small, privately observed, random deviation that was less than \$1 in absolute value. These random deviations are meant to model idiosyncratic match values, e.g. that a medical student may have a small preference for a hospital that is located near friends and relatives. A simplified version of this setup, with only 4 workers and 4 firms, is shown in Table 10.2. For example, the high productivity firm F1 would earn \$14.90 if it is matched with worker W2, who is preferred to W1, W3, and W4 in that order.

If the true rankings in the table are submitted under the deferred acceptance procedure, then firms 1 and 2 send initial offers to W2, who keeps F2 on the table and rejects F1. Firms 3 and 4 send offers to W1, who keeps F3 and rejects F4. The low productivity workers get no offers in the first stage. The F1 offer that was rejected initially is resent to W1, who keeps it and rejects the offer it had (F3). At this stage, the high productivity workers, W1 and W2, are provisionally matched with firms F1 and F2 respectively, and these pairings cannot be overturned by a low-productivity match. It can be shown that the process ends when W3 is matched with F3 and W4 is matched with F4 (see question 5). In the end, W4 would rather be matched with F3, but F3 would not want to switch, since W3 is preferred to W4 by this firm. A similar observation can be made for the high productivity firms and workers, so the outcome is stable.

Table 10.2. Payoffs for Each Matching

High productivity firms		Low Productivity Firms	
F1	F2	F3	F4
W2 = 14.90	W2 = 15.50	W1 = 15.30	W1 = 14.90
W1 = 14.20	W1 = 14.60	W2 = 15.00	W2 = 14.50
W3 = 5.90	W3 = 5.30	W3 = 5.50	W4 = 5.50
W4 = 5.60	W4 = 4.60	W4 = 4.70	W3 = 5.10

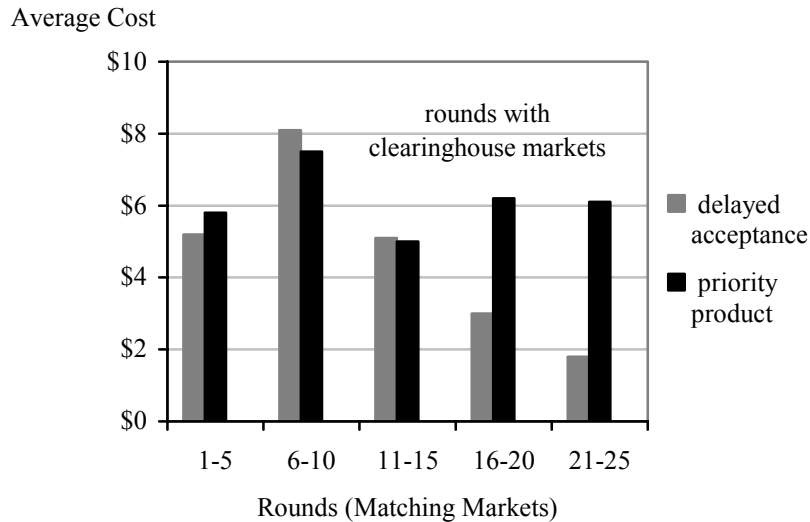
High Productivity Workers		Low Productivity Workers	
W1	W2	W3	W4
F2 = 14.90	F2 = 15.50	F1 = 15.50	F1 = 15.50
F1 = 14.80	F1 = 14.60	F2 = 14.60	F2 = 15.10
F3 = 5.60	F3 = 5.00	F3 = 5.30	F3 = 4.90
F4 = 5.30	F4 = 4.60	F4 = 4.60	F4 = 4.50

Next consider the priority product method under the assumption that participants submit their true preference rankings. Since W2 and F2 are each other's first choice, the priority product is 1 for this match, which can be shown to be the lowest product. This match is then implemented, taking W2 and F2 out of the picture. Next consider whether F1 will be matched with W1. They are each the other's second choice, so the priority product is 4. This is higher than the product for F1 and W3, which is 3 since F1 is W3's first choice and W3 is F1's third choice. The product for F3 and W1 is also 3, and these are the lowest priority products remaining after those in the initial match (W2, F2) were removed (questions 6 and 7). Thus the algorithm will match W1 and F3 in one pair, and W3 and F1 in another. This is unstable since both of the high productivity agents (F1 and W1) would prefer to be matched with each other.

An unstable system will create regret and dissatisfaction after matchings are announced, so participation may decline as workers and firms seek to match with each other bilaterally. The way the process worked in the experiment is that each firm could send an offer to a single worker in each period. There were three periods in the process, -2, -1, and 0 (graduation day). Such bilateral matches incur a cost of \$2 if accomplished at time -2, and they incur a cost of \$1 if accomplished at time -1. There is no cost for making a match in time 0. Anyone who arranges a match must not make or accept subsequent offers. If there is a clearinghouse, it operates for unmatched people and firms in time 0. Anyone who remains unmatched after time 0 earns \$0. Each session began with 10 rounds or repetitions of this 3-period market process, with no clearinghouse. Match value profiles were shuffled randomly from round to round, but these profiles stayed the same for the periods -2, -1, and 0 of a round. After 10 rounds with no

clearinghouse, one of the clearinghouse mechanisms was implemented at period 0 of round 11 and in all subsequent rounds.

One way to evaluate the outcomes is to look at the average costs associated with early matches, which are graphed in Figure 10.2. The first two sets of bars, on the left side, are for the 10 rounds prior to the imposition of a clearinghouse in round 11. Notice that costs of early matching rise from the first five rounds to the second five rounds, which shows the unraveling as earlier matches generate higher costs. Then costs for both treatments fall in rounds 11-15, as each type of clearinghouse reduces costs for the sessions in which it is operative. But in the final rounds 16-25, costs increase for the sessions with the priority product algorithm, and costs fall for the deferred acceptance sessions. Although not shown in the graph, the most costly matches in time period –2 vanish completely for the delayed acceptance sessions in markets 21-25.



**Figure 10.2. Average Costs of Early Matchings**  
**Source: Kagel and Roth (2000)**

This experiment is quite innovative from a methodological perspective. First, note that the sequential (within-subjects) design is used because the clearinghouse mechanisms were implemented in a natural experiment in settings where unraveling had already occurred, so any sequence effects were a focus of attention, not a nuisance. The experiments enabled the authors to observe the patterns of behavior *during the process of adjustment* at a level of detail that was not available from field data. In the experiments, the sudden availability of a clearinghouse option did not deter firms from making early offers just as they had before. What happened, however, was that workers who were not matched in

periods  $-2$  and  $-1$  ended up trying the clearinghouse and becoming satisfied with its results, especially for the deferred acceptance mechanism. The priority product mechanism did improve matching initially before unraveling re-emerged. In addition, the laboratory experiment, with only 6 agents on each side of the market, was able to reproduce both the unraveling patterns observed in the field and the superior performance of the mechanism that generated stable matches. Without a laboratory experiment, one might have wondered whether the survival of the deferred acceptance method in two locations in the U.K. and its demise in two other locations had been due to chance differences in demographics, social norms, or some other unobserved aspect of the field context. The laboratory used subjects selected from the same pool for both treatments, so the performance differences could be attributed to the incentive properties of the two mechanisms. Taken together, the laboratory experiment and the natural experiment provide a clearer picture of adjustment and behavior in these types of markets.

### **III. Baseball, Dorm Rooms, School Choice, Deep Space, Sorority Rush, ... but Marriage?**

The matching mechanisms discussed in section II involved both sides of the market. The deferred acceptance procedure was developed by Gale and Shapley (1962), who discussed applications to college admissions and marriage. Two-sided mechanisms are widely used in sorority rush, where “unraveling” might mean making a choice soon after arriving on campus, before having time to find a range of friends. Nalbanthian and Schotter (1995) report market-like experiments motivated by matching of free agents in professional baseball. In addition, there are many interesting experimental studies of one-sided matching problems, e.g. assigning students to dorm rooms when money-based allocations like auctions are precluded (Chen and Sonmez, 2004). Olsen and Porter (1994) study both auction-like and non-price rankings-based mechanisms that are motivated by the Jet Propulsion Lab’s allocation of time slots on its Deep Space Network of antennas. Chen and Sonmez (2005) use experiments to evaluate the “Boston System” of assigning students to schools on the basis of submitted rankings. They conclude that alternatives like the Gale-Shapley mechanism would improve performance. The Boston system looks good on paper because students tend to get assigned to the school that they rank the highest. Chen and Sonmez note that students act on incentives to misstate their preference rankings under this system, which makes the assignments look better than they are in terms of actual preferences, which are induced and known in their experiment.

#### **Questions**

1. Use the unit costs and redemption values for the market design in Figure 10.1 to calculate the total surplus for each grade level.

2. In a competitive equilibrium for a given grade, the price is determined by the intersection of supply and demand for that grade. Compare the equilibrium profits for each grade for the market structure in Figure 10.1.
3. For the market structure shown in Figure 10.1, consider a small increase in the cost for each seller's first unit *for grade 2 only*. An increase of \$0.25 in this cost would mean that the competitive equilibrium profits for each seller are lower for grade 2 than for the other grades. What effect do you think this cost increase would have on the tendency for grade choices by sellers to converge to the optimal grade (2)?
4. Groucho Marx once remarked: "Please accept my resignation, I don't want to belong to any club that will accept me as a member." How is this situation different from (or similar to) the exogenous quality model discussed at the end of Section I?
5. Assuming that the preference orderings implied by Table 10.2 are submitted under the deferred acceptance program, show that W3 is matched with F3 and W4 with F4.
6. Please fill out the table of priority products for the example from section II of the chapter, assuming that the submitted preference orderings are those in Table 10.2. The second column and row of the priority product table have been completed for you, and this column and row can be crossed out since the top priority product (1) causes W2 and F2 to be matched first. Use the resulting table to determine all four matches.

	W1	W2	W3	W4
F1		2		
F2	2	1	6	8
F3		6		
F4		8		

7. Your answer to the previous question should show F1 being matched with the low productivity worker W3. If F1 had foreseen this unfortunate outcome, how could F1 have changed the submitted ranking of workers to achieve a better match with a high-productivity worker? Explain.
8. Consider how you might modify one of the match value numbers in Table 10.2 to show that the deferred acceptance algorithm does not necessarily generate an efficient outcome in the sense of maximizing total earnings of all firms and workers. To do this, consider a very large change in a match value that does not alter the outcome of the deferred acceptance procedure, but that does cause a firm to miss out on an extremely productive match with a worker who does not rank that firm very high.