Draft: Comments are welcome.
Please do not quote without
author's permission.

Using Laboratory Experiments to Teach Introductory Economics

Jeffrey Parker

Department of Economics
Reed College
Portland, OR 97202
(503) 517-7308
Fax: (503) 777-7776
Internet: parker@reed.edu

August 1995

Acknowledgements

This book is a significantly revised version of a document entitled *Economics 201 Instructor's Laboratory Manual*. Many individuals contributed to the preparation of that document. Most of all, my colleagues in teaching Economics 201 at Reed, Noelwah Netusil, Denise Hare, and Zenon Zygmont, have all contributed significantly to the design of the course, the lab, and the experiments. The potential benefits of lab experiments as a teaching tool in economics was first made clear to me through a seminar conducted by Donald Wells and Arlington Williams sponsored by the National Science Foundation at the University of Arizona in May 1989. The preparation of the manual was supported by the Sloan Foundation through a New Liberal Arts Grant to Reed College.

1 Experiments and Introductory Economics

The use of controlled experiments to test economic hypotheses is not new; Roth (1995, 5) traces the origins of experimental economics back at least to the 1930s. However, it is only in the last three decades that a recognizable research field of "experimental economics" has developed. The field has now reach two important milestones of maturity: the appearance of a comprehensive reference volume (Kagel and Roth 1995) and the publication of a textbook in the field (Davis and Holt 1993).

Research experimentation has thrived in applications where it is impossible to observe data arising from naturally occurring experiments with sufficient clarity to test important economic hypotheses. For example, early experiments focused on direct tests of risk aversion and the expected-utility framework. Experiments testing individual choice theory have been common in both economics and psychology since the 1960s. Other areas that have spawned a large experimental literature within economics are testing of the efficiency of various auctions and other forms of market organization, prisoners' dilemma situations and other simple game-theoretic applications, public-goods provision and free ridership, and various kinds of bargaining frameworks (Roth 1995).

The systematic use of experiments as pedagogical tools has only begun to become wide-spread in the 1990s. The growing popularity of classroom experiments has been largely due to a series of seminars held at the University of Arizona by Donald Wells and Arlington Williams under the sponsorship of the National Science Foundation. The purpose of this volume is to introduce the use of experiments to teachers of introductory economics and to describe some common experiments that have been adapted for classroom use.

Many economists now use experiments actively in their teaching. The many creative applications presented at recent conferences demonstrate the broad scope of problems to which experiments can be applied as teaching tools.[1] This book makes no attempt at a comprehensive review of all of the classroom experiments in use; indeed, the use of experiments has progressed too far to make compilation of such a review possible. Instead, it focuses on a handful of common experiments that have proved to be successful tools for demonstrating to students key ideas involved in the typical introductory economics course.

All of the experiments discussed in this book have been used at Reed College in the laboratory that accompanies the introductory Economics 201 course. While the results

[1]The Economic Science Association meetings in November 1994 and the Western Economic Association conference in July 1995 both had multiple sessions devoted to using experiments in the classroom. New experimental applications are published regularly in the semi-annual newsletter *Classroom Expernomics*.

certainly vary from year to year and section to section, there is enough consistency to the outcomes that we now consider these experiments to be "bullet-proof," in the sense that we can run them with confidence that the results will prove pedagogically useful.

Why Use Experiments?

For most teachers, running experiments during class time represents a dramatic departure from the normal format of the introductory economics class. This provides strong arguments both for and against using experiments. Because they are distinctive and more participatory than ordinary class session, students are likely to enjoy and remember experiments and the lessons associated with them. However, for these same reasons preparation for the experiment and the follow-up lessons are likely to take more instructor's effort (at least the first time) than delivering the same old lecture once again.

Educators have long recognized [need some citations here] that students remember lessons better if they are actively involved in them than if they experience the lessons merely as a reader and listener. Experiments, when used as this book suggests, involve each student first as a participant then as an analyst, trying to draw on his or her growing knowledge of economic theory to interpret the results of the experiment.

The experience component of the experiment can be very important. Many introductory economics students have had little experience making economic decisions, especially those on the supply side of the market. Most have never thought seriously about the market processes at work in the markets in which they have participated. Precisely because the experimental setting is so simple, it demonstrates very directly to students the basic decisions involved in being a buyer, seller, producer, taxpayer, policy-maker, or whatever other role they may assume. They often remember subtle aspects of the experiment that would be impossible to get from media or textbook accounts of actual events. For example, the process by which market participants acquire and use information is remarkably transparent to those who have actually done it, but often difficult to describe in the abstract. Moreover, the experience they gain is common to all students who participate. This common experience can form the basis for useful references in subsequent class sessions—experiments are ready-made case studies. For example, in explaining how the invisible hand affects buyers and sellers, an instructor might say "Remember how you felt in the double-oral auction when the price had converged and no one would buy or sell at a price higher or lower than equilibrium?" Each class member who participated in the experiment is likely to recall the experience to which the instructor is referring.

The opportunity to try out the tools of economic theory on data arising out of an actual observed event helps students see the relevance of the theories that are presented in their textbooks. It encourages the students to think about theories critically, assessing under what conditions the textbook conclusions are likely to be correct, rather than blindly accepting (or rejecting) the textbook's assertions. The use of experiments allows the instructor to cast himself or herself as a "scientific observer" testing the validity of theories rather than as a "preacher" asserting their truth and relevance. There are aspects of nearly every experiment that follow the textbook's predictions accurately and other aspects in which surprises occur. Follow-up assignments and discussions can encourage students to think about how the simple structure of the experiment compares to the simple assumptions of the theory they are using. Where there are differences, students can often see how these deviations led to the unexpected results.

Finally, students (and instructors) usually think experiments are fun! By enhancing students' enjoyment, experiments can improve the "mood" of the introductory economics class. Happy and excited students learn more and are more likely to recommend economics to their friends and to continue their own study of economics at higher levels.

What Makes a Good Classroom Experiment?

Before beginning the discussion of some of the ways in which experiments have contributed to teaching at Reed and elsewhere, it is useful to think about what kinds of topics lend themselves to classroom experiments. Although most of the experiments currently used in the classroom have their origins in the experimental economics research literature, the criteria for suitability as a research experiment and as a classroom demonstration experiment are quite distinct. As noted above, researchers use experiments when data generated under controlled conditions in the laboratory allow a clearer test of a hypothesis than is possible using other observations. This condition is neither necessary nor sufficient for classroom experiments to be useful.

The most important criterion for a good classroom experiment is that it must relate to a central issue of the course. For example, simple auction experiments relate to the conept of market equilibrium, one of the central issues in introductory microeconomics. Similarly, public-goods experiments help students to understand the crucial concept of market failure. Other research experiments, such as bargaining experiments or those dealing with coordination failure, may not "fit in" as well to the standard introductory economics course.[2]

A second criterion for a good classroom experiment is that the concepts it demonstrates to students are not already obvious or easily understood without the experiment. An "experiment" in which students eat ice cream until they get sick is one way of demonstrating the concavity of indifference curves. However, reflection on their own consumption experiences together with a little imagination has, in my experience, usually been sufficient to cement the concept firmly in their minds. By contrast, the

[2]This is not to say that these topics are unimportant or that they do not justify time in an introductory course, only that most introductory texts, and thus presumably most introductory courses, do not spend much time on them.

nature of market equilibrium and how markets converge to equilibrium are often opaque when described in lectures and on the blackboard. Experiments can allow students to see and experience how information and competition drive markets to clear.

Other criteria are important for classroom experiments as well. It must be possible for students to learn the roles of participants quickly. If an hour is available for the experiment, students must be able to assimilate the instructions (and perhaps have a "practice round") in no more than about fifteen minutes. Some complex research experiments require considerably longer. Since fun is part of the allure of experiments, it is wise to avoid experiments in which participation becomes quickly tedious. Participants in the classroom are more than "experimental subjects," they are your students who will be back in class week after week throughout the term. A poorly chosen or poorly designed experiment can alienate rather than stimulate them. Finally, the results of the experiment should be fairly predictable and lend themselves to straight-forward interpretation using the theories being studied. It would be risky, for example, to begin the term by illustrating equilibrium with a "posted-offer" market experiment, where convergence is slower and less reliable than in the double-oral auction.

The Four Phases of an Experiment

The student's experience with an experiment can be broken down into four principal phases:

- preparation for the experiment,
- participation in the experiment,
- individual attempts at interpretation of the results, and
- in-class analysis of the results.

Preparation for the experiment is the exposure that the student receives to the theoretical models and tools that will be required to participate in the experiment and to interpret the experiment's results. Normally this consists of textbook reading and lectures examining basic theories. Experiments vary widely in the amount of preparation required. Some (for example, the voluntary-contribution experiment) require little or no preparation; others require that the student be able to apply basic theories (such as the profit-maximization rule MC = MR in a posted-offer experiment) in order to be effective participants. In the discussions of individual experiments below, considerable attention is given to the issue of appropriate preparation of students for each experiment.

The most memorable aspect of the experiment for the student is likely to be participation in the experiment itself. From the instructor's point of view, the actual experiment is little more than a process by which data for analysis are generated. However, careful design and execution of the experiment is very important. A tedious, excessively stressful, or otherwise unpleasant experiment is likely to alienate students rather than making them curious about relationship between the experience and economic theory.

After the experiment has been performed, students may be given a few days to attempt to interpret the results individually with the help of some leading questions asked by the instructor and a description of the experimental setup and results. These interpretations can be submitted as homework or "lab reports," which are useful not only as an evaluation tool but especially as a diagnostic device and a focus around which students may organize their thoughts about the experiment.

After the students have submitted individual interpretations of the results, it is highly desirable to devote a class period (or part of one) to group discussion of the experiment. In this setting, the instructor can guide discussion to make sure that major connections between important theoretical concepts and experiment results are covered in detail.

The Role of the Instructor

The instructor's responsibility begins well before the actual experiment with the design and preparation of experiment materials. As discussed below in chapters dealing with individual experiments, each student typically receives a piece of paper at the beginning of the session with a description of his or her role in the experiment and cost or value information. These materials are prepared by the instructor along with appropriate forms for recording experiment results.

At the beginning of the actual experiment, the instructor must read instructions to the participants. This step is of great importance, but is often underestimated. If one or two students do not understand the instructions and therefore engage in loss-making trades or similar activities, the behavior of the market may be dramatically altered. Not only does this make it more difficult to demonstrate the applicability of economic theory in the experimental setting, but it makes it far more difficult for students to interpret the experiment results. Instructions should be read slowly, repeating important points using different words and providing simple examples (e.g., "If you are a buyer and your sheet says you value the good at $5.00, then do not pay more than $5.00 for it or you will make a loss."). In market experiments, it is highly desirable to track participants' behavior after the first trading period by somehow examining the transaction prices to check that each trade was profitable for both parties. One way of doing this is for the instructor or an assistant to visit briefly with each participant, asking her to show her transaction price (if any) and her information sheet. This allows the instructor immediately to identify individuals who may have misunderstood the instructions and to explain why they made a loss rather than a profit. This is particularly true of more complicated experiments in which participants may have opportunities to buy or sell multiple units or alternative goods.

A common problem in disseminating instructions is late-arriving students. At best, they necessitate starting over when halfway through the instructions; at worst they show up when the experiment is already in progress. Unless the number of participants is too small initially, it is better not to allow new participants to join an experiment in progress. It is much more difficult to evaluate the results of experiments if the number of participants changes from period to period as the expected results may vary across periods and if the equilibrium changes, the process of convergence is often undermined. The students should be admonished to arrive on time on days when experiments are scheduled. Late-comers should be used as observers or assistants (collecting forms, writing data on board, etc.) rather than added in the middle of an experiment.

Once the experiment is running and all participants seem to have the procedure under control, the instructor's role should usually be limited to the essential tasks of administering the experiment: opening and closing trading periods and collecting results. Although an occasional prod ("Buyers, how come you are paying so much?") may be helpful in nudging the market more quickly toward equilibrium, students should feel like the market mechanism and not the instructor's direction are steering the overall behavior of the experiment. Indeed, one of the main lessons of these experiments is the degree to which decentralized market systems achieve predictable, equilibrium results. If students leave the experiment feeling like the instructor's intervention was largely responsible for convergence to equilibrium, a main point of the experiment is lost.

There should usually be a follow-up assignment or report to be completed by the students based on a summary of the results prepared by the instructor. This assignment should be designed to focus the students' attention on the connection between the textbook theory and the experimental data. Students find it helpful to be told which theoretical tools might help them analyze the experiment (e.g., telling them that it might be helpful to know the demand curve), but you may often leave the actual application of the tools to them (e.g., figuring out what the actual demand curve in the experiment was).

At least part a subsequent class should be devoted to a discussion of an experiment. If at least some of the students have understood the experiment well, they should be able to carry the discussion. The agenda is set by the instructor both through questions asked on the report assignment and through leading the direction of the discussion. The instructor should make certain the all main lessons from the experiment are well developed in the discussion and provide answers to problems that the students are collectively unable to untangle.

Motivating Participants

One area of considerable controversy among those who do experiments in the classroom is how to achieve "saliency." In research experiments, a modest monetary reward for each participant based on his or her success is used to provide motivation for participants to pursue the individual behavioral objectives that the experimenter wishes to induce. Using conventional levels of rewards, a typical experiment costs $100 or more. Few teachers have a budget that allows for such payments, so the issue of how to induce the desired objectives in participants is a potential problem.

The simplest option, which has worked well in most cases at Reed College, is simply to ask the students to behave as if they were earning money based on the outcome of the experiment. In most cases, the motivation of "doing well" in the experiment has proved sufficient to generate results that are satisfactory for classroom analysis and comparable to those of research studies.

In cases where the induced incentives of the experiment may present real costs to the student, the absence of a tangible reward becomes a significant factor. For example, in the voluntary contribution mechanism, students who are discovered (or believed, whether correctly or not) to be free riding are subject to real disdain from their peers. In a small college, where most students know each other and have ongoing academic and nonacademic relationships, avoiding the resentment of one's peers can outweigh the incentive of "doing well" on the experiment. In this experiment, the classroom results often differ substantially from those in the research literature. When asked to explain the difference, students often point to the absence of any cash reward in the classroom experiment. Their recognition of the importance, in some circumstances, of the magnitude of rewards and of the importance of the institutional setting and the relationships among agents is itself a useful concept taught by the experiment.

Some professors use grade-related rewards in classroom experiments. While the use of grade points or extra-credit points should give students some motivation toward the experimenter's induced objectives, it is sometimes difficult to give enough points to motivate students without overly compromising the overall grading system.[3] If many experiments are being used, then it may be feasible to accumulate points over the term and assign an overall experiment grade based on the cumulative outcome. This has the advantage, through the law of large numbers, of lessening the impact of good or bad luck on an individual experiment, such as being a low-cost seller in a market experiment.[4]

[3]Our experience at Reed College (on which much of this book is based) may be unique in this respect. Motivation by grades is explicitly avoided as a matter of college policy. Reed students are not informed of their grades unless they inquire and individual assignments normally have no grade attached. Nonetheless, I generate some motivation by telling them that part of their evaluation on the lab part of the course will be based on their performance in the experiments, though the bulk will be on their written and oral analysis.

[4]In his comments at a session on classroom experiments at the July 1995 Western Economic Association meetings, Professor Terrence Langan of St. Thomas University described a scheme of "Econ Dollars" which accumulated

Another alternative, worthy of consideration if self-motivation, monetary payments, and grade incentives are unsuccessful or infeasible, is a merchandise reward of some kind. We use food as a reward in some Reed experiments, either intrinsically by using food as the actual commodity being studied or more abstractly by bringing a heterogeneous product such as assorted doughnuts to the next class session and offering the first choice of variety to those with the highest experimental "earnings."

A variation on this theme would be to tell students at the beginning of the term that there will be an auction of merchandise using accumulated experimental earnings over several experiments. The auction could then involve relatively inexpensive items such as college mugs, school supplies, or something similar. It may be possible in some cases to solicit donations to supply items for the auction (souvenirs from the college bookstore, services from campus facilities such as bowling alleys, pool halls, etc., or even a home-cooked dinner at the instructor's house). If the students are not told the details of the items to be auctioned until after the experiments are over, there will be no loss of incentive for people who do not happen to need or want the items available. Moreover, the auction itself might provide an interesting setting for a late-semester class social gathering.

The Plan of This Book

The remaining chapters of this book are devoted to a series of experiments that are suitable for use in an introductory economics course. Each experiment serves to highlight one or more important concepts that are commonly taught in economics. Few instructors will be able to find time for all of them, but most of the experiments can be run independently of the others, allowing individual instructors to pick and choose among the experiments quite freely.

Chapter 2 describes the work-horse of classroom experiments, the double-oral auction. This experiment seems to work almost flawlessly under a wide range of conditions: with as many as 100 or as few as 12 students, with unstructured trading or trading guided by a "pit-boss," with experienced participants or with total novices. The double-oral auction is widely used because it is simple, realiable, and demonstrates very effectively some concepts of market equilibrium that can otherwise be very difficult for students to grasp.

Chapters 3 through 5 describe three experiments that use a posted-offer trading institution similar to that found in American retail stores. Sellers post prices (and usually set an upper limit on sales volume) and buyers shop among sellers looking for the lowest

---

throughout a one-month course in experimental economics and which were used to help determine grades.

available price. The first of these experiments is designed to be roughly symmetric to the double-oral auction experiment, to provide a bridge between the behavior of the double-oral market and markets that are more commonly recognized by students. The second variant of the posted-offer market bring in the role of risk by forcing sellers to declare production levels for each period before the market opens and penalizing them for unsold units. The final version demonstrates the effects of market concentration as a posted-offer market moves from a relatively large number of sellers to an oligopoly and finally a monopoly.

Chapter 6 examines an experiment testing the ability of markets to process buyers' and sellers' signals about the optimum quality of a good to produce. In this experiment, widgets are available in two grades, gold and silver. As the parameters of the experiment vary from beginning to end, the optimal level of production changes from gold to silver to a mixture of the two. The market institution used here is the double-oral auction. A variation of this experiment can be used to demonstrate the "lemons" problem of adverse selection. If sellers are not able to reveal the quality of the widget they are selling, all will end up selling silver widgets.

Chapters 7, 8, and 9 describe experiments in the economics of information. The first is a hands-on attempt by students to use various indicators to evaluate the quality of an actual product such as orange juice, pumpkin pie, or chocolate-chip cookies. The second is an experiment that demonstrates the basics of search theory. Chapter 9 looks at several brief experiments dealing with auctions for a good of uncertain value.

Chapter 10 presents another work-horse, the voluntary contribution mechanism. This experiment tests the willingness of class members to sacrifice personal gain for a greater collective gain. Chapter 11 considers other public-goods and free-rider experiments that can supplement or provide alternativcs to the voluntary contribution experiment. Chapter 12 considers experiments involving social decisions about the optimal income distribution, including testing the applicability of Rawlsian justice.

2 Double-Oral Auction Experiment

The double-oral auction market is the most commonly-used experiment in introductory classes. One introductory textbook now has a "kit" for running such experiments that is provided to faculty adopting the text. This experiment is ideally suited for the introductory class because both buyers' and sellers' roles are easily understood, it has a lively and interesting pace, and it nearly always achieves market equilibrium in a few trading periods.[5]

The double-oral auction market also has historical significance in experimental research. Chamberlin's (1948) use of this experiment is often regarded as the birthof modern experimental economics.

Overview of the Experiment

The word "double" in the title of the auction refers to the fact that both buyers and sellers can make price bids/offers in hopes of striking a deal. The "oral" implies that these prices are called out by mouth, though some recent implementations have used computer networks to replace oral communication.

There are two general approaches to managing the double-oral auction. Under one, trading is organized by a "pit-boss" who accepts bids from buyers, offers from sellers, and acceptances of outstanding bids or offers from sellers and buyers. The other approach allows traders to find each other with no central coordination via general "floor trading." Both approaches seem to achieve rapid and reliable convergence toward equilibrium. The pit-boss method is less noisy and chaotic and is more easily used in a classroom or auditorium with unmovable furniture since traders need not move around. The floor-trading method works well where enough space for a trading floor can be made available and where noise is not a problem. The obviously unstructured interaction of the floor-trading scheme can be advantageous, since one of the pedagogical aims of the double-oral auction is usually to demonstrate the credibility of the invisible hand. Even the limited role of a pit-boss keeps the instructor at the center of the activity, and students may leave with the mistaken notion that something the instructor did led the market to equilibrium.

This experiment can be run with as few as eight or ten students or with fifty or more. Half of the students are designated as buyers and half as sellers of a hypothetical commodity (widgets). Each buyer receives a piece of paper indicating the *value* of one widget to him. This value can either be explained to the student either as the consump-

---

[5]For more detailed information about research designs and results using double-auction markets, see chapter 3 Davis and Holt (1993) or Holt (1995).

tion value of the widget or as a price at which you as the experimenter are willing to buy widgets from him at the end of the experiment. Each seller receives a sheet indicating her *cost* of producing one widget.

Values vary among buyers and sellers. Some buyers receive a value well above the market equilibrium price (though they will not realize this initially), others are less fortunate. Some buyers get values below the equilibrium price, so they will be excluded from the market unless they can persuade a low-cost seller to accept less than the going price. Cost values of sellers, similarly range from well below the equilibrium price to somewhat above it.

The experiment proceeds by a series of "trading periods," each of which is a replication of the basic experiment. Each buyer and seller may make only one transaction in each trading period.[6] Under the floor-trading method, as soon as the trading period begins, buyers and sellers move around on the trading floor offering to undertake transactions with one another. When a buyer/seller pair has agreed on a price, they come to the instructor's table and report the transaction. Information about prices may be disseminated to the entire group of participants as transactions happen, at the end of the period, or not at all. The trading period can be brought to a close "softly" when trading activity naturally ceases or at the end of a specified period of time.

If the pit-boss method is adopted, the instructor is positioned at the front of the class with a blackboard for keeping track of the current outstanding bid and offer. He or she recognizes individual traders who raise their hands. When a seller is recognized, she may either accept the currently active bid, or make an offer to sell at a price lower than the currently active offer. In the former case, a transaction occurs and is recorded, the accepted bid is erased from the blackboard, and trading resumes. If the seller makes a lower-price offer, then the instructor (or an assistant) erases the previous offer and replaces it with the new offer price and the identification of the seller making the offer.

When a buyer is called on by the pit-boss, he may either accept the currently active offer or make a bid at a higher price than the active bid. Acceptance of an offer completes and transaction and erases the accepted offer. A higher bid replaces the active bid with the new one. Trading proceeds until action stops or until a predetermined time limit is reached.

An experiment usually consists of five or more trading periods. Experience suggests that the market will have converged to a narrow range of transaction prices after four or five periods in the floor-trading system, or somewhat quicker in the pit-boss system.[7]

Once general convergence to equilibrium has occurred, variations may be introduced in later periods, such as encouraging collusion by either buyers or sellers or imposing a price ceiling or floor on the market. Although one must be careful not to try

---

[6]This aspect of the setup is strictly for convenience and simplicity. Many research experiments have used double-oral auctions with multiple units for each buyer and seller. We use a single unit per period to avoid the need to explain why successive units are valued less than the first and so that students do not have to try to remember how many transactions they have already undertaken in order to know their marginal values.

[7]This does not mean, however, that the pit-boss experiment necessarily takes less clock time. Because of the sequential nature of trading, each period usually lasts somewhat longer under the pit-boss system, so convergence in three pit-boss periods may require as much time as convergence in six floor-trading periods.

to push a single experimental session to cover too many theoretical concepts, these are obvious and quick extensions that probably would not warrant a separate experimental session.

The process of handing out information sheets, reading instructions, and performing five to eight periods typically requires about 30 minutes. Since class periods are usually at least 50 minutes in length, this means that two experiments can usually be run. There are two good reasons for running two separate replications of the double-oral auction in a single class. First, a second experiment can be designed to demonstrate the effects of a shift in supply or demand, allowing students to observe (in their ex-post analysis) the comparative static effects on price and quantity exchanged. Second, during the first experiment, several students (buyers with very low values and sellers with high costs) quickly find out that they are unable to find profitable trades.[8] These students become bored observers by the third or fourth period. To involve them more effectively into the process, the second experiment reassigns the values and costs of each buyer and seller so that those who were at a disadvantage in the first experiment have profitable trading opportunities in the second, and vice versa. The need for reassignment will be obvious to all students and provides a logical justification for handing out new values. If they are not told, students may not recognize that market conditions have changed, so convergence to the new equilibrium price usually begins from the first-experiment equilibrium.

Details of the Experiment Design

Participants Excluded from Market. Buyer values and seller costs will vary across participants. It is possible to design an experiment in which everyone can make a profitable transaction. Such a design may be necessary for experiments with only ten or twelve participants. However, it is pedigogically useful to have the demand and supply curves extend beyond the equilibrium quantity, thus excluding some participants from making transactions at the equilibrium price. Many students enter with the presumption that it is optimal for every participant to buy or sell. One potentially important lesson of the experiment is why it is socially desirable for some potential buyers and sellers *not* to trade in widgets (i.e., because their desire for widgets does not exceed marginal production cost). Only by having some "outsiders" who are excluded from the market can this point be made.

As noted above, exclusion also provides a motivation for running two distinct experiments in order to bring those who were outsiders in the first into the action in the second. This inclusion/exclusion feature also provides useful "cover" by masking another purpose of the second experiment (shifting supply or demand). Although the presence of

[8]This is the principal disadvantage of allowing only a single unit for each trader: either there are no unprofitable units or some traders are shut out completely.

a few outsiders are desirable, the marginal benefit of having more than two or three excluded buyers and sellers is negligible and just leads to more people standing on the sidelines.

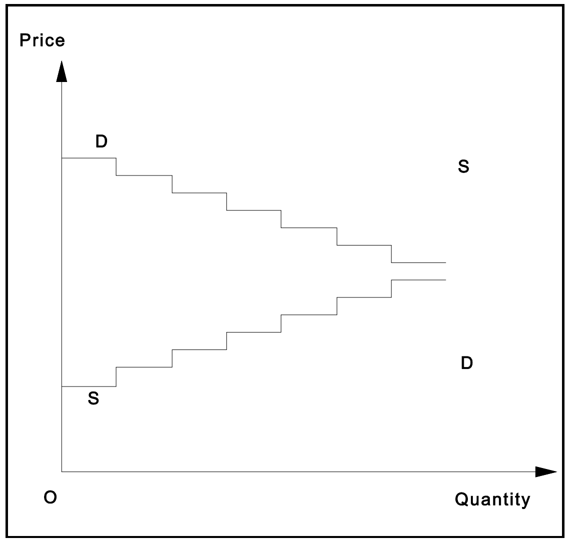

The Nature of Equilibrium. Since fractional widgets are not possible, the demand and supply curves will have the form of step functions, as shown in Figure 2-1. Step-function curves can intersect in four ways: (a) on vertical segments of both curves, (b) on horizontal segments of both curves, (c) with a vertical segment of the demand curve cutting a horizontal segment of the supply curve, or (d) with a vertical segment of the supply curve cutting a horizontal segment of the demand curve. These four possibilities are shown in Figure 2-

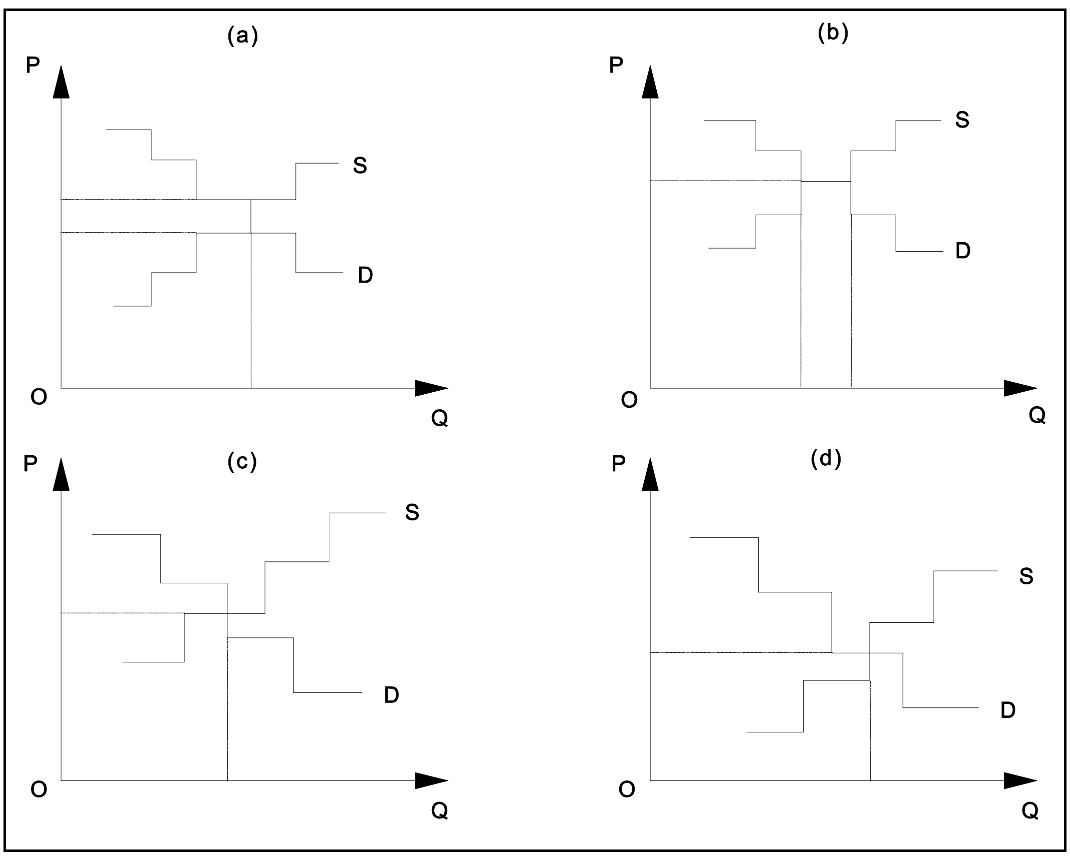

2. Although the differences among these kinds of equilibria are trivial for most economic applications, they are potentially important for experiments with small numbers of participants.

In model (a), there is a unique equilibrium quantity exchanged but a range of possible equilibrium prices. In (b), the equilibrium price is unique, but at least two quantities exchanged are possible. In both of the "swastika" models, both price and quantity are determined uniquely. However, in model (c), there are (at least) two sellers who are just willing to sell at the competitive equilibrium price and only one buyer to purchase a widget from the two sellers. Since at least one potential seller will be disappointed in this case, the notion of equilibrium as a state in which everyone is happy with the outcome of the widget market (given his or her preferences) is not satisfied. Similarly, model (d) leaves at least one buyer who is willing to buy at the equilibrium price but cannot find a seller.

This suggests that model (a) or (b) may be preferable in spite of the nonuniqueness of equilibrium. Between the two, model (b) has a distinct disadvantage in that the individuals considering making the "final" transaction have zero profit in doing so. Thus, the result may vary depending on whether individual participants decide that making zero-profit transactions are worthwhile. In contrast, the "price tunnel" model of (a) yields well-defined sets of buyers and sellers for whom transactions are desirable or undesirable at any price in the equilibrium range. Interpretation of the results as "equilibrium" is also somewhat easier, since the price should converge only to the interior of the range, not necessarily to a single value. For these reasons, the double-oral auction experiment at Reed is performed with a price-tunnel equilibrium such as that of panel (a) of Figure 2-2.

The Number of Buyers and Sellers. In class situations, the instructor often does not know the exact number of students that will participate. This makes it important to design an experiment that will work well with a range of possible numbers. The double-oral auction can be designed quite easily for a wide range of numbers.

If there are an even number of participants, the number of buyers and sellers are normally set to be equal. Conceptually pairing the values of participants in the design (Buyer 1 with Seller A, etc.) allows the experiment to be arranged so that the equilibrium price (range) is independent of the number of participant pairs. The instructor can then calculate the expected equilibrium price before the experiment and know as it progresses whether equilibrium is being achieved.

To make the equilibrium price independent of the number of participants, begin by deciding on a value for the midpoint of the equilibrium price range, say \$3.00. Then set values for each buyer-seller pair so that if the buyer's value is greater than \$3.00, then the paired seller's cost is below \$3.00. And if the buyer's value is below (so that he is an outsider who will not trade at equilibrium), then the seller's cost is above (so she is an outsider, too). By pairing *m* potential participants in this way and assigning the first $n/2$ pairs (where n # m is the number of students who show up for the experiment) to participants, for each arriving pair that is added to the experiment either both can trade at $3.00 or neither can trade at $3.00, so the market is still in equilibrium at $3.00.

With an odd number of participants, the equilibrium price will change depending on whether or not the "extra" person is not an outsider. An extra buyer or seller that is not an outsider will change the nature of the equilibrium from a price tunnel to a swastika intersection. Because of these complications, it is often easier to incorporate the last-arriving student as an assistant who writes results on the blackboard if an odd number of participants would result.

Steepness of Demand and Supply Curves. The steepness of the demand (supply) curve determines the amount of consumer (producer) surplus or profit in the market. For experiments in which cash payments are used, the total surplus is constrained by the amount of desired payments per period. Without cash payments there is no such constraint.

There is no necessary reason why the demand curve and supply curve should or should not be symmetric relative to the equilibrium price. Making either curve more steeply sloped than the other skews the distribution of surplus or profits toward the participants whose behavior is represented by the steep curve.

Interesting variations on the double-oral auction can be obtained by making the demand and supply curves "lumpier," i.e., creating curves with fewer and steeper steps rather than small, smoothly declining ones. The limiting case of this is the "box" model where all buyers have the same value and all sellers have the same cost. As the curves become lumpier, there is a greater likelihood of getting stuck on one of the lumps and not converging to equilibrium, especially in experiments where each buyer and seller can make more than one transaction.[9] Since a major purpose of conducting this experiment in the classroom is to demonstrate that markets *can* converge to equilibrium, it seems better to avoid such potential complications and use smoothly-stepped demand and supply specifications.

Another design issue is the width of the equilibrium price range. If the range is made very large, there is a wide region of uncertainty about where the price will settle. Making it very small eliminates most of the gains from exchange on the last transaction, so the number of transactions is more likely to fall short of the equilibrium quantity. If prices are denominated in cents, then an equilibrium price range of 10 cents means that the least profitable buyer and least profitable seller actually trading can share 10 cents of profit.

---

[9]See Holt (1995, 393) for a discussion of double auctions in which lumpy supply curves with multiple units convey market power on sellers, leading to noncompetitive outcomes.

Information. One of the key parameters that must be set in the design of the double-oral auction experiment is the amount of information each participant will obtain as the experiment progresses. Initially, each buyer and seller knows his or her own value, but not those of other traders. As they make bids and offers in the market, participants begin to discover what prices are feasible.

If the experiment is run with a pit boss, then information about bids and offers is public and information travels instantaneously through the market. In the floor-trading setup, there are several choices of information transmission. The most common is to post the price of each exchange on the blackboard as it occurs. This gives buyers and sellers up-to-date information on all transactions in the market, not just on their own attempts to make an exchange. The other extreme of the information spectrum is to provide no public information about transactions. An intermediate procedure is to post prices (or perhaps just an average price) at the close of the trading period.

The speed of convergence to equilibrium is quite sensitive to the spread of price information in the market. If participants are informed quickly about the prices of other transactions, they will adjust their own bargaining rapidly in response. This leads to quick and reliable convergence. Since that is an important objective of the classroom experiment, it seems desirable to facilitate convergence with generous information provision in most cases.[10] One potential problem with continuous posting is that traders may delay making a transaction (especially in the first period) in order to obtain information about other prices before committing themselves. The posting of prices after the end of the period rather than continuously would eliminate this problem should it arise.

What the Students Should Know

Participation in the experiment requires no knowledge of economics at all. If desired, it could be run on the first day of class. In practice, it works well to place it about the second week of the semester when the general subject of market equilibrium is being taught.

Should students be familiar with the what the textbook says about these subjects before the experiment? The answer is not clear; the experiment can be usefully employed either way. The fundamental lessons that are illustrated by this experiment are the meaning of demand and supply curves, the concept of market-clearing equilibrium and convergence to it, and the measurement of gains from exchange. If variations are incorporated into the experiment, other issues such as the effects of price controls and

[10]If time allows, it is most instructive to run a second double-oral auction with no information posted. Students will quickly discover who much more difficult their job is and convergence is usually slower. They may also engage in more careful search, sampling more potential partners during their trading to replace the information they are provided about other traders' prices in the standard experiment.

the instability of collusion may also appear. If students have studied these ideas before the experiment, they should be able to write more complete interpretations of the experiment in their lab reports. The experiment can serve as an example to reinforce ideas they have already seen. Alternatively, if students study the material for the first time immediately after the experiment, the interpretation of the experiment can be a focal point around which to center their study of the subjects.

In the lab at Reed, the students usually have a minimal introduction to supply and demand curves and the idea of equilibrium before the experiment. The main theories are developed in the class periods between the lab in which the experiment is performed and the lab in which it is discussed. The experimental results are used in the discussion as an example of the theoretical concepts. Students write their lab reports using the theories at about the same time that they are being presented in class.

Instructor Preparation and Materials

To prepare for the double-oral auction experiment, the instructor must first decide on the values for buyers and sellers that define the demand and supply curves, then prepare the information sheets to be handed out to the participants. Careful instructions need to be prepared to be read to the students. The instructor should also plan what, if any, market distortions are to be added to the experiment. Possibilities include trying to induce collusion among sellers or buyers (this sometimes happens spontaneously), imposing a price ceiling or floor, or imposing a quantity rationing scheme. Collusion usually breaks down very quickly, demonstrating to students the tendency of cartels to be unstable. That price controls and rationing lead to inefficient outcomes is obvious to students, particularly to those who are left out of the market.

If the buyer and seller values are designed according to the principles suggested above, then the equilibrium price range can be determined in advance without knowing how many students will actually participate. It is wise to know this information before the experiment because sometimes nonconvergence can be easily rectified if it is recognized, as when one participant does not understand the experiment and mistakenly makes unprofitable trades.

The materials that should be ready before the experiment are:

- information sheets for all participants for each part of the experiment,
- a page of written instructions to be read to the participants,
- a writing instrument and some simple forms for recording the transactions that occur during each period of the experiment (a clipboard is also handy),
- a calculator, if average prices in each trading period are to be posted.

Samples of each of the written materials can be found in Appendix A. It is desirable to have a worksheet as part of each participant's information sheet on which he or she can record transactions and calculate profit. If a student misunderstood the instructions and is making unprofitable transactions, it should be quickly apparent if he or she is recording the results and calculating the profit correctly. It is also a written record that the instructor can easily check (see below).

Running the Experiment

Setup. It is always desirable for the instructor to arrive in the classroom at least five minutes before the experiment is to begin. The floor-trading version of this experiment works best when there is a large trading floor. If possible, furniture should be moved to the perimeter of the room so that students can sit around the outside of an empty circle or square to receive instructions. The instructor should have a table or desk at the front of the trading area with a blackboard nearby.

As the students arrive, they should find seats around the room and wait until the experiment is to begin. They should be warned in previous class meetings that it is extremely important not to be late for the experiment. As noted above, latecomers are disruptive because they may miss part or all of the instructions. Once trading in the experiment has begun, it is difficult to add new participants (the supply and demand curves shift, changing the equilibrium quantity exchanged).

When it is time for the experiment to begin, the information sheets are handed out, with the buyers on one side of the room and the sellers on the other side. As they are handed out, the instructor should suggest that participants not share the information on their sheets with others. Then the instructor reads the instructions for the experiment (see example in Appendix A), slowly and carefully with examples as necessary. The importance of effectively communicating the "rules of the game" cannot be overemphasized—nothing is more difficult for students to interpret in their analytical reports than an experiment that did not converge to the equilibrium price because of one student behaving unpredictably. Students should have an opportunity to ask questions if they do not understand the instructions. Examples are very helpful, such as, "Buyers, if your information sheet says that the value of a widget to you is $25.00, then you should try to buy a widget as long as you don't have to pay more than $25.00."[11] When the students do not have any more questions and seem to have understood the instructions, the experiment is ready to begin.

[11]In order not to provide accurate or unaccurate clues about the equilibrium price, the numbers chosen for the numerical examples in the instructions should be far away from the numbers in the experiment. For example, if the range of price in the experiment is $4.00 to $5.00, use a number such as $0.25 or $300 in the example.

Initial Trading Period. In the first period, prices will be erratic. Participants have no information about what prices are likely to be, so prices of the first few transactions are likely to be very high or low. Once transactions begin to occur, buyer/seller pair will begin to report transactions. For each transaction, the instructor should note in his or her records the identities of the buyer and seller and the price. If prices are to be posted as they occur, then they should be written on the blackboard as they are reported.

If the trading period has a time limit, the instructor should remind traders of the amount of time remaining as the deadline approaches. Since most transactions take place quickly, a fixed limit is not usually necessary (and can be imposed in subsequent periods if a particular group of participants causes a significant delay).

There will often be a few stragglers still bargaining after most of the participants have made transactions. Some of these are "outsiders," who are destined to be disappointed in the widget market. Others may be potential transactors who are still bargaining down to the last penny. To keep the experiment moving in the absence of a fixed time limit, the instructor should warn them that they need to finish up. Usually, those who are going to be successful then agree on a price and report their transactions, while the others retreat in disappointment, hoping for better luck in the next period.

Once trading has concluded, the instructor announces that the first period is over and tells the students to mark their transaction prices on their information sheets and calculate their profits. It is a good idea at this point for the instructor (and assistants if any) to circulate around the room and look at each student's sheet. If a student is found to be making unprofitable transactions, it can be pointed out and explained, so that such problems do not persist into later trading periods.

Subsequent Trading Periods. The first period generally takes more time than subsequent periods. Once students become familiar with the procedure and begin to know what price to expect, transactions occur more quickly. If all participants are maximizing profit, the range of prices usually converges (more or less) to the equilibrium range in three to four periods.

The experiment becomes monotonous very quickly once prices have converged, but there is some pedagogical benefit to this monotony. If students are asked at this stage why they don't spend more time trying to find a better price, they usually respond, "I had no choice. Nobody would sell for less (buy for more)." At this point they have become price takers. If there is time (and it does not take very long), a couple of periods of monotonous equilibrium may provide a reference point for future discussions of price taking in perfectly competitive markets. It is up to the individual instructor to weigh the benefits of additional periods of equilibrium against the opportunity cost, which is more time to implement variations such as market distortions.

Collusion. Once the market seems to have converged, the instructor has the opportunity to introduce distortions such as collusion or price controls. Collusion can be induced by simply calling the sellers together quickly and suggesting that they attempt to maintain price at a given level higher than the equilibrium. If they do this, trading will usually halt as buyers are initially unwilling to pay the higher price. One or two buyers (with high demand values) sometimes cave in and buy at the collusion price, but the seller cartel usually breaks down before the end of the first period. Competitive equilibrium is usually restored after one or two more periods.

Enterprising students will often perform this experiment spontaneously. If this happens, be sure to write down for future reference which side was attempting to collude and what their initial price was. This information can be included in the write-up of experimental results that is given to students.

Changing Experiments. As discussed above, running two separate experiments in a single session allows full participation by those left out of the market in the initial experiment and enables a shift in the demand or supply curve to be unobtrusively performed. If two experiments are planned, the change from experiment one to experiment two (with new information sheets) can occur after about five to eight periods (roughly three to converge to equilibrium, two at equilibrium, and three of attempted collusion if that is included). Allowing experiment one to continue any longer than this is very difficult on the outsiders, who have not been participating actively since the initial couple of periods.

As discussed above, the students should believe that the purpose of the second experiment is to reshuffle values to include the outsiders more actively. Any questions about other changes can be answered ambiguously, so that they do not know whether to expect the same or different prices.

Trade in the first period of the second experiment often starts from the first experiment's equilibrium price, then eventually moves toward the new equilibrium. There is often considerable inertia despite excess demand or supply at the old price. Market theory predicts that a frustrated participant who would like to buy (sell) at the market price but cannot find a partner will begin to bid a higher (offer a lower) price, which will lead the market to its new equilibrium. In practice, the student in this position may take several periods to realize that this is his or her best strategy. Although it would be strictly forbidden in a formal research experiment, it may lead to pedagogically beneficial results if the instructor quietly suggests to the relevant student(s) in this situation that they may want to raise or lower their price slightly to try to attract a partner. Such a "nudge" away from the inertial point compromises the experiment as a test of the ability of the market institution to find the new equilibrium spontaneously (which is well established in the research literature on double-oral auction markets). However, with only about fifty minutes to demonstrate several concepts to students, the loss of formal rigor must be balanced against the benefit of rapid convergence.

Price Controls. A price ceiling or floor is easily implemented in this experiment. The instructor simply announces to the participants that transactions above (below) the ceiling (floor) price will not be accepted. Since all transactions must go through the instructor, there is no possibility of evading the controls. (This point can be made in subsequent discussions.)

Periods are completed very quickly under price controls as (in the case of ceilings) buyers race to be the first to transact with the limited number of sellers who can sell at the ceiling price. Two or three periods are usually sufficient to demonstrate the effects of controls.

The Report Assignment

The raw materials for students' analyses of the experiment include complete descriptions of the setup of the experiment and its results. The setup information should include the values for all buyers and sellers, which can be arranged in a table as shown in the example in Appendix A.

In the first years of this experiment at Reed, students were encouraged to write down all the transaction information from the blackboard either as it happened or at the end of the experiment. There are large economies of scale in this process, so it is now performed by the instructor or an assistant. A sample summary of results is shown in Appendix A. It shows the price and the identities of the buyer and seller for every transaction that occurred. The context of each period is also shown, indicating when observable changes in market behavior occurred (price controls, organized or spontaneous collusion, entry or exit of participants, etc.).

Most students who have attended only a fewr meetings of their first economics class find it difficult to structure their thinking about the experiment around the central concepts of market theory. Experience suggests that simply turning them loose with instructions to "explain the results of the experiment using economic theory" will generate many reports that read like the play-by-play of a football game: "Then, in period 5, the buyers seemed to gain the upper hand as the average price of a widget fell from $4.67 to $4.63." A set of guiding questions such as those in the example in Appendix A will help students to focus the analysis in the proper direction. However, the instructor needs to make clear whether the assignment is a report, to be written as a unified document addressing the issues raised by the questions, or a problem set in which the student simply answers each question in turn.

The usual timetable for this experiment at Reed has the instructor or assistant preparing this information between the end of Thursday afternoon's lab and the beginning of Friday morning's classes. Students' lab reports are then due sometime the following week and are discussed in lab the following Thursday. Having reports due Monday or Wednesday allows the instructor to look at them before the Thursday lab discussion to determine what concepts students understood well and what ones were missed by most students. Postponing the due date until Thursday gives the students more time for analysis and writing. Practice at Reed varies among instructors.

Reading and Discussing Student Reports

Few, if any, students master the fundamental concepts well enough to prepare a fully correct analysis of the experiment. It is very important for the instructor to give strong encouragement to them; many will feel very inadequate when they see in the lab discussion that their analyses were incorrect in one way or another. Many will approximate the demand curve, supply curve, and equilibrium correctly, but there are several details that many students miss: particularly the step-function nature of demand and supply curves and the measurement of potential gains from exchange. Each student should realize that these are subtle concepts that are not usually picked up spontaneously on the first pass.

For the class discussion, the instructor can prepare appropriate handouts or transparencies to show

- exact demand and supply curves and the market equilibrium for both experiments,
- range of experiment prices (and perhaps quantity exchanged) in relation to equilibrium values,
- potential gains from exchange (consumer and producer surplus), and
- efficiency of the experimental market as measured by the percentage of the potential gains from exchange that were achieved.

A sample handout for this experiment is shown in Appendix A. However, simply handing out this sheet at the beginning of the lab period may turn a potentially active discussion into a lecture. Students should be allowed to present their own analyses (and to correct one another's mistakes) as far as they are able to go. If the students are collectively able to reach the correct analysis, then the instructor's prepared information simply confirms the students' conclusions. If the students reach a dead end or fail to grasp one of the central points, the instructor can try to lead the discussion in an appropriate direction and use the prepared information as an expositional tool. There are several main points that should be covered in the discussion, which are discussed in the sections below.

The Concepts of Supply and Demand. Analyzing this experiment forces students to operationalize supply and demand curves. To do this, they can simply use the data on buyer and seller values to address the question, "At a price of *X*, how many widgets would buyers want to buy (or sellers want to sell)?" Using this approach, both the general concept of what the curves represent and the step-function shape of the experimental curves are easily illustrated.

Convergence to Equilibrium. Most double-oral auction experiments illustrate effectively the process of convergence to equilibrium and the character of market clearing. Since they have actually performed the experiment, the students have "felt" the forces leading the market toward the equality of supply and demand.

One useful question for them to consider in the discussion is whether the market in the double-oral auction experiment is perfectly competitive.[12] It fails to meet most of the textbook conditions for competition: there are (depending on the size of the class) relatively few buyers and sellers, information is (at the beginning) very imperfect, some transactions occur at nonclearing prices. The only assumption that is clearly fulfilled is homogeneity of the good being traded. Nonetheless, by the third or fourth trading period, the results of the market usually seem to approximate those of perfect competition quite closely. This can be used as a lever to motivate discussion of "How many buyers and sellers is enough for competition?" and "What information is necessary for buyers and sellers and how do they get it?" Through consideration of these questions, students can be shown that although the pure model of competition is rarely, if ever, observed in actual markets, the basic ideas of supply, demand, competition, and equilibrium hold more broadly.

Another useful issue for discussion is the role of information in price convergence. The students can be asked to consider how the market would have behaved differently had prices not been written on the board. As participants, they grasp very easily the importance of price information and its role in shaping their behavior as buyers and sellers. Some discussion of the ways in which buyers and sellers get information in real-world markets is easily motivated here.

A final point on equilibrium deals with the nature of equilibrium itself. In practice, there are usually a few transactions lying just outside the equilibrium price range, even in later periods. The market may also be one transaction short of the equilibrium number in some periods. Some instructors may wish to use this to develop the idea of equilibrium as a central tendency toward which natural forces tend to move markets, rather than the normal state of markets. This leads students to see the theory of markets as a general set of principles with wide but often imprecise application, rather than a detailed prediction that is contradicted whenever diequilibrium or price disparity is observed.

Gains from Exchange and Efficiency of Equilibrium. The easiest way to measure efficiency of a market experiment is by the share of potential gains from exchange in any trading period that are realized by participants. If cash rewards are used, this is simply the share of the money the participants could have earned that they

---

[12]Textbooks and courses vary in when the definition of perfect competition is introduced. This suggestion clearly applies only to courses in which students have encountered the definition before analyzing the experiment.

actually did earn. (Even if no cash rewards are used, this is an effective way to illustrate gains from exchange.)

Even in experiments in which full equilibrium is not achieved (i.e., some transactions are outside the range of competitive equilibrium prices or the number of transactions is one or two short of equilibrium), the percentage of gains realized is likely to be very high. Those participants with the largest potential surplus (sellers with low costs and buyers with high values) have the most to gain from transacting and the easiest task finding a partner. If the market "misses" one potentially profitable transaction, it is usually between two participants with relatively small potential gains. Thus, it is not unusual for experimental periods in which prices vary considerably and quantity exchanged is one short of equilibrium to achieve 90 percent or higher efficiency. This is important for students to observe, since it probably applies as well to actual markets.

Effects and Instability of Collusion. Collusion arrangements can occur spontaneously in these experiments, or they can be induced by the instructor. In either case, they provide fruitful material for discussion. Collusion among sellers (buyers) usually raises (lowers) the price for a few transactions as high-value buyers (low-cost sellers) give in and make profitable transactions at the cartel price set by sellers (buyers). However, it usually breaks down within the first trading period as sellers (buyers) are left over after demand (supply) is exhausted at the cartel price. (The indivisibility of output makes collusion even more difficult to sustain in this experiment than in most natural markets.) Asking the sellers (buyers) who broke ranks to explain why they did so leads naturally into a general discussion of the incentives to cheat on cartel arrangements.

Sometimes the cartel succeeds in maintaining a high (or low) price for one or more periods. In this situation, a similar discussion can be motivated by asking the sellers (or buyers) who were not able to sell (buy) at the cartel price why they did *not* break ranks and offer to sell (buy) at a lower (higher) price.

Effects of Shift in Supply or Demand. By the time the experiment is discussed in lab, the students will usually have seen basic comparative-static results from the supply-demand model in class. The shift in one curve that occurs if two experiments are run can be demonstrated by repeating the derivation of the curves for the second experiment. The effect on the equilibrium price range can then be tracked and compared to the actual experimental result. Most students find this application straight forward.

Effects of Price Controls. Price ceilings (floors) lead to excess demand (supply) at the controlled price. A convenient starting point for the discussion is the question, "What determined which buyers (sellers) got to buy (sell) and which did not?" The students will be able to answer this better than anyone who did not participate in the market, including the instructor. Common responses about who the successful traders were include those who had traded with low-cost sellers (high-value buyers) in previous periods and who found each other quickly, people with loud voices who could be heard easily, and people with fast feet who got into the trading area quickly. This situation can be contrasted with the market-clearing situation in which everyone wanting to buy (sell) at the prevailing price was able to find a seller (buyer).

The experiment also makes it easy to demonstrate the lost gains from exchange due to controls. Using the concepts of consumer and producer surplus, the instructor (or students) can show how much surplus is lost if the highest-value buyers (lowest-cost sellers) are the ones who are able to make transactions, and why this is a lower bound on the actual lost surplus. By evaluating the actual surplus achieved in these periods, students can see that the actual loss can be quite a lot larger.

Using this Experiment in Later Classes

Several of the central concepts illustrated by this experiment are useful reference points for later classes. The notion that markets converge toward (even if they do not fully persist in) market-clearing equilibrium can be made more credibly after students have observed the spontaneous convergence of the decentralized experimental market. In particular, explaining the forces that drive markets toward equilibrium should be easier when students have seen these forces in action. The constraints faced by a price taker should be familiar to those who have participated in several periods of sustained market clearing. Instability of collusive arrangements is usually covered much later in the course (when studying oligopoly), but an example often occurs in this experiment.

Variations

There are many ways in which the basic double-oral auction framework described here can be varied. The basic framework has worked well at Reed for the purposes outlined above. Variations may be useful for instructors with additional time to devote to experiments or with particular points that they wish to emphasize to their classes.

Variations in Information Provision. One key element leading the experimental market to equilibrium is information obtained by traders. Generally, the more quickly they get information about the prices at which exchanges are taking place, the more quickly the range of prices narrows. The experiment can be done with no price information written on the blackboard, so students need to "shop" in order to obtain price information from traders on the opposite side of the market. Alternatively, information can be provided only at the end of each period, so traders are uninformed about the current period's activities but know the range of prices from the last period. These experiments are interesting for purposes of comparison, but since they usually converge more slowly they do not allow as wide a range of activities to be undertaken in a lab session of given length. Performing the first experiment with reduced information may mean that it is not feasible to incorporate collusion, a shift in one curve, and price controls in the same session.

Alternative Setups. As noted above, there are four different kinds of equilibria that can arise with step-function demand and supply curves. If time permits, experiments exploring these alternatives can be performed. Another interesting variation is the "box model," in which every buyer has the same value and every seller has the same (lower) cost. The supply and demand curves in this model are shaped like a backward L and a backward Γ, which makes the equilibrium diagram resemble a rectangular box. The range of equilibrium prices is the entire range between the buyers' value and the sellers' cost (assuming there is an equal number of each). By adding an extra seller or an extra buyer, the intersection becomes a swastika and the equilibrium price moves either to the bottom or the top of the original range.

Especially when there are a small number of students in the class, it is tempting to introduce experiments in which each buyer and seller can trade more than one widget. In experiments at Reed, this has enjoyed only limited success. If each buyer and seller can trade, say, three widgets at the same value or cost, then nothing prevents them from simply trading three units when they find a partner rather than one. Most likely, the equilibrium of the market would be the same. (An interesting but untried variant would be to prevent multiple trades in the same period from between any pair.) When buyers and sellers are given multiple units with varying values at this stage of their economics training, they often become confused and engage in transactions that are unprofitable at the margin. Such experiments may be better left to a later stage when students have a more sophisticated understanding of "marginal" concepts.

3 Posted-Offer Market Experiment

The double-oral auction market usually pursuades students that decentralized market mechanisms with multiple buyers and sellers are capable of achieving market-clearing equilibrium and generating large gains from exchange. As such, it serves a crucial role early in the introductory economics course. However, the trading institution of the double auction is one that students associate more closely with financial markets and with bazaars in developing countries than with the consumer markets they see around them.

Most American consumer markets set prices through a posted-offer instutition; sellers post a price at which they offer each item for sale and sometimes set an upper limit on the amount they are willing to sell at that price. Buyers then decide how much of each seller's product to buy at the posted price. Depending on the details of the market and the good or service being produced, unsold units offered may or may not incur production costs and it may or may not be possible to carry them over into future periods (possibly at some cost) as inventories. In service markets and some produce-to-order goods markets, variable costs are typically not incurred unless units are sold but no inventory carry-over is possible. In markets for perishable goods, full costs are usually incurred whether or not the unit is sold, but inventories are still impossible or very costly. For less perishable goods, costs are incurred for unsold goods but inventory carry-over is feasible at moderate cost.

Given the prevalence of posted-offer markets in the experience of American students, one or more experiments with posted-offer institutions can help bridge the gap between theory and their outside lives. Participation in the posted-offer market can be somewhat more complicated than in the double-oral auction. Moreover, research experiments have found that posted-offer markets are consistently slower to converge to equilibrium and are somewhat less efficient in capturing gains from exchange than double auctions (Holt 1995, 368-72). Because the posted-offer experiments are more difficult to perform and more likely to generate ambiguous or anomalous results, one is wise to run a double-oral auction experiment first, then to present the posted-offer market as a second institutional setting for comparison rather than skipping the double auction and jumping directly to the posted-offer market.

This chapter and the following two chapters discuss experiments based on the posted-offer instutition. In this chapter, a posted-offer experiment for a service or produce-to-order market is developed that is directly comparable to the double-oral auction of the previous chapter. The purpose of running this experiment in the classroom is to examine the impact of trading and price-setting institutions on convergence to equilibrium and market efficiency. In the next chapter, several alternative structures are explored in which sellers incur costs on unsold units. At the instructor' choice, these units can be useless, can be sold in a "clearance sale," or can be carried over as inventory to the next period. Chapter 5 describes a posted-offer experiment in which the number of sellers is reduced as the experiment progresses to examine the effects of market concentration on price and output.

Overview of the Experiment

Unlike the symmetric roles of the double-oral auction, sellers in the posted-offer market have a much more interesting and challenging role than buyers. Sellers set price and decide on a production level; buyers merely buy as many units as are profitable from the lowest-price sellers currently having units for sale. For this reason, the posted-offer experiment is often run with the instructor or an assistant acting as "simulated buyers," and all student participants (either individually or in small groups) being sellers.

Besides the simplistic nature of the buyer's role, there can be logistical difficulties in running large-scale posted-offer experiments with human buyers. A simultaneous trading mechanism such as the double-oral auction's trading floor can lead to collisions among buyers as they all rush to find the lowest-price seller and buy from her. This problem is usually solved in research experiments by having buyers make their purchases one at a time in random sequence. With a large number of buyers, each period can take a relatively long time to complete as each buyer must be chosen randomly and then shop.

Where there are significant differences in their treatments, separate subsections will discuss human-buyer and simulated-buyer experiments in each section of this chapter.

With Human Buyers

Each buyer receives an information sheet showing the values of widgets to him. Each buyer may be constrained to buy only one widget or may be allowed to buy more than one with marginal values that may be consstant or may decline as subsequent units are purchased.[1] Each seller receives an information sheet telling her the cost of producing widgets. As with buyers, the sellers can be constrained to one unit as in the double-oral auction or they may be allowed to produce more than one unit (up to some capacity limit) with constant or increasing marginal cost.

Sellers begin each trading period by deciding how many units to offer and what price to set. (All units produced by a single seller must sell for the same price.) Once all sellers have made their decisions, they are collected by the instructor and a table like Table 3-1 is written on the board. It is desirable to have seller decisions submitted in written form to avoid the possibility that late-reporting sellers may change their decisions based on the revealed decisions of the first sellers to report.

Next the buyers must be sequenced randomly and do their shopping. Numbered balls, beans, or beads can be drawn without replacement (Bingo tokens work well) or

[1]The idea of declining marginal valuation can be difficult for some students to grasp, making it more difficult to be sure that the instructions have been effectively understood. However, most posted-offer experiments use multiple units.

several sequences of random numbers can be prepared ahead of time. As each number is drawn, the instructor calls out the identity of the next buyer, who then decides how many units to buy and from what seller(s) to buy them. As noted above, the buyer's decision is quite easy—look for the lowest price on the board and compare to his valuations to decide how much to buy. As units are bought from a particular seller, the instructor reduces the number of units that seller has for sale in the table on the blackboard. The buyers proceed sequentially until all have had an opportunity to buy, then the results of the trading period are totalled up (average price, total quantity exchanged, etc.) and the next period begins with seller decisions.

With Simulated Buyers

If the instructor is to simulate the role of buyers, then sellers' decisions need not be posted on the blackboard.[2] Once the decisions are submitted, the instructor orders them from lowest price to highest and moves down the aggregated demand schedule until no additional profitable purchases can be made. He or she then reports back to each seller (either on the board, orally, or by paper) how many units of that seller's units were sold.

Seller	Price	Units for Sale
A	$4.00	5
B	$3.70	4
C	$3.80	4
D	$4.60	6
E	$3.75	5
F	$3.90	4

Buyers and sellers should total their profits each period and mark them on a table provided on their information sheets. The experiment proceeds through as many rounds as the instructor desires. Since convergence of the posted-offer market usually requires more periods than the double auction and since each period may be more time-consuming because of the sequential buying process, there is less time available for pursuing variations such as collusion or price controls than in the double auction.

Since the role of buyer is less interesting, experiments with human buyers may find it useful to run two posted-offer expeiriments and to reverse buyers' and sellers' roles from one to the next. For example, one could run an initial experiment with each trader limited to a single unit. This setup maximizes the symmetry with the double-oral auction, allowing students to compare the results directly. A second experiment, in which buyers and sellers are interchanged, could allow multiple units to be sold or purchased by each agent. Waiting until the second experiment to introduce multiple units also spreads out the process of instructing participants: they learn the institutional

[2]Note, however, that if all decisions are not posted, this withholds information from sellers that would be available in the human-buyer setup. This change in information regime may affect the behavior of the market.

setup of posted-offer markets in the first experiment, then learn to deal with multiple units and (possibly) declining marginal values and increasing marginal costs in the second experiment.

Details of the Design

Number of Buyers and Sellers. In the single-unit experiment with human buyers, the number of buyers, sellers, and units is determined by the number of participants, as in the double-oral auction. In multiple-unit experiments and in experiments with simulated buyers, the number of buyers and sellers must be decided.

When buyers and sellers can each exchange multiple units, there is no need for the number of buyers to equal the number of sellers in order for the total number of potential widgets on each side of the market to be equal. For example, there could be six sellers each of whom can produce up to twelve widgets and twenty-four buyers who can buy up to three widgets each. There is the potential to produce 6 H 12 = 72 widgets and the potential to buy 24 H 3 = 72 widgets, which retains a degree of symmetry between demand and supply. (Since the experiment will most likely have an equilibrium quantity exchanged well below 72 units, the attractiveness of this symmetry is more aesthetic than practical.)

Since one purpose of this experiment is to examine whether posted-offer markets do as well as double auctions in achieving equilibrium and exploiting gains from exchange, this is not the time to reduce the number of sellers low enough to approach oligopoly levels. In practice, six sellers seems to be sufficient to approximate a competitive outcome, though there is no reason why the number cannot be larger.

The optimal distribution of buyers and sellers in multiple-unit experiments depends to some degree on the size of the class. With a class of 30, a 15/15 split with each buyer and seller transacting a maximum of three or four units would work well. Alternatively, the 6-seller, 24-buyer setup discussed above would be feasible. In a class of 60, it may be preferable to reduce the number of effective participants by pairing students up, running two separate experiments with 30 participants, or by having only half of the class participate with the remainder just observing. Doubling the size of the 30-participant experiment would lead to extremely long and tedious trading periods as 30 or 48 buyers have to be quizzed about their purchasing decisions.

In an experiment with simulated buyers, each person in the class is a seller. There is no conceptual limit on the number of sellers in a single market, but a large number of sellers will make the process of determining how much was bought from whom much more time-consuming. With six sellers, this can be accomplished in a minute or so for each period by an experienced instructor. With 30 sellers, it probably takes at least three or four minutes, long enough that the students may be distracted from the experiment and begin socializing. Thus, again, it may be that a large class should be split for a simulated-buyer experiment with several parallel experiments being run simultaneously with 6-12 sellers in each. Alternatively, students could be paired or some students could observe rather than participating. Although the role of buyer is not very exciting, it can absorb a relatively large number of students, so classes of 24 or more students may find the human-buyer market more convenient.

Cost Schedules. In the single-unit experiment, the cost schedule is a single point: each seller can produce one unit at a given cost. In multiple-unit experiments, an upward-sloping marginal-cost schedule is usually used.[3] Linear marginal cost schedules work well, though other specifications would probably be equally effective. Some care in designing the cost schedules can assure that all sellers produce at partial capacity when the price is at equilibrium. The cost schedules shown in Table 3-1 have linear (increasing) marginal cost and zero fixed cost. They have been designed so that each seller's marginal-cost schedule reaches a common level of 40-50 at an intermediate range of output. The market demand schedule for the experiment is then set in a way that puts the equilibrium price level in the 40-50 neighborhood.

Some sellers' MC schedules are quite steep; others are much flatter. This leads to wide differences in the amount of profit that they earn in equilibrium; the sellers with steeper marginal-cost schedules earn much higher profits on inframarginal units. For example, using the marginal costs in Table 3-1, if the equilibrium market price is 50, then both seller B and seller C produce 7 units, but C makes 154 units of profit while B makes only 76. In experiments in which actual payments are made, fixed costs that vary across sellers could be used to even out (and reduce) the equilibrium earnings of sellers.

Market Demand Schedule. In asymmetric (more buyers than sellers) experiments with human buyers, the instructor often does not know in advance the exact number of buyers who will participate. With the market demand schedule varying, the exact equilibrium price will depend on the number of students who participate. One can reduce this dependence by constructing the market demand schedule to be quite elastic near the presumed equilibrium value, so that extra buyers or missing buyers would not have much effect on price. For example, in an experiment with up to 30 buyers each of whom can buy up to 3 units, the demand calues can be constructed so that each buyer's second unit falls in the 40-50 range. With the six sellers shown in Table 3-1, 31 units are produced at a price of 40 and 44 at a price of 50. A market demand schedule with buyers' second-unit values distributed over [40,50] would assure that the equilibrium price would be between 40 and 50 if as few as 16 buyers participate or for as many as 30.

In symmetric experiments—where the number of buyers and sellers are equal—the discussion in Chapter 2 is relevant. Buyers and sellers can be "paired" to add either inframarginal or extramarginal units on both sides of the market as additional pairs of

[3]While flat marginal-cost schedules and demand-value schedules are possible in multiple-unit experiments, there is no substantial difference between using *both* flat curves and using single units. The number of units exchanged is just increased by the appropriate multiple. Flat marginal-cost schedules and downward-sloping demand-value schedules can be used together.

students arrive. This make the equilibrium price independent of the number of participants. When simulated buyers are used, it is possible to design separate market demand schedules to be used with 5 sellers, 6 sellers, or other number. This allows an appropriate adjustment to be made to the demand side of the market to achieve any desired equilibrium price with varying numbers of sellers.

Output	A	B	C	D	E	F
1	15	30	7	24	12	15
2	20	33	14	28	18	20
3	25	36	21	32	24	25
4	30	39	28	36	30	30
5	35	42	35	40	36	35
6	40	45	42	44	42	40
7	45	48	49	48	48	45
8	50	51	56	52	54	50
9	55	54	63	56	60	55
10	60	57	70	60	66	60
11	65	60	77	64	72	65
12	70	63	84	68	78	70
13	75	66	91	72	84	75
14	80	69	98	76	90	80
15	85	72	105	80	96	85

What the Students Should Know

The single-unit, symmetric implementation of the posted-offer market requires no prior knowledge on the part of the student. The roles of buyers and sellers are very similar to those in the double-oral auction. Buyers buy if a unit is available at a price lower than their demand values; sellers offer units at a price higher than their costs. The major difference is in the method of market interaction.

However, when multiple units are introduced, as in the 15-unit marginal-cost schedules in Table 3-1, participation in this experiment requires some background in the relevant economic theory. Sellers must be comfortable with the principle that they should produce only units for which price equals or exceeds marginal cost. Failure to apply this principle consistently may lead to erratic behavior by one or more sellers, which will impact the equilibrium of the market in this experiment.

Similarly, multiple units on the demand side introduces the complication that students must understand that they should buy additional units only if the *marginal* value exceeds price. Neither of these concepts is especially difficult and both are central to understanding classical microeconomic theory. Both must be mastered by all students in the course of learning microeconomics. However, a successful experiment that yields the results predicted by economic theory requires that *all* students achieve a functional mastery of these concepts, as well as the posted-offer trading mechanism, based on a five-minute set of oral instructions. If this is likely to be a problem for the students at hand, multiple-unit experiments may be postponed until after the relevant concepts have been covered in class. As a compromise strategy, it may be effective to run a single-unit experiment first to get them used to the posted-offer institution, then proceed to a multiple-unit experiment (in the same class period) after explaining to them very carefully how to understand the concept of marginal valuation (cost) with the multiple units.

Instructor Preparation and Materials

After deciding on the structure of the experiment(s) to be run—the number of simultaneous or sequential experiments to be run, whether to use simulated or human buyers, single or multiple units for each agent, and the numbers of buyers and sellers—the main task of the instructor is to prepare the relevant information sheets for all agents and an appropriate set of instructions.

If a single-unit experiment is being run, then the information sheets can be relatively simple like those described for the double-oral auction in Chapter 2. A multiple-unit experiment requires more information to be given to participants and entails more complex record-keeping. Table 3-2 shows an example of a seller information sheet for a hypothetical seller of up to 10 units. The table at the top shows marginal and total costs; the worksheet at the bottom aids the student in computing profit each period. The terms "marginal" and "total" are used despite the possibility that students may not yet know these terms. The amount of explanation in the instructions may be varied by the instructor depending on the familiarity of the concepts and terms to the students in the class. The sheet shown in Table 3-2 has zero fixed cost, but positive fixed cost can be incorporated easily by inserting a line above or below the top table telling the student what her fixed cost is and adding the appropriate amount to each total cost number. Alternatively, an additional line in the table could be used for fixed cost.

If human buyers are to be used, information sheets such as the one shown in Table 3-3 must be prepared for each. Again, the terms "marginal" and "total" are used in their ordinary economic contexts. The worksheet portion of this sheet is made more complicated by the possibility that buyers may pay different prices for each unit they buy.

Seller A Information

Unit Number	1	2	3	4	5	6	7	8	9	10
Marginal Cost of That Unit	25	30	35	40	45	50	55	60	65	70
Total Cost of That Many Units	25	55	90	13 0	17 5	22 5	28 0	34 0	40 5	47 5

Record of Profits by Period

Period	**Units Sold**	**Price**	**Total Revenue** (units H price)	**Total Cost** (from table)	**Profit Earned** **(rev. - cost)**
1					
2					
3					
4					
5					
6					
7					
8					
9					
10					

Buyer 1 Information Sheet

Unit Number	1	2	3
Marginal Value of Unit	68	53	36

Record Sheet for Computing Profits

Period	First Unit			Second Unit			Third Unit			Period Profit
	Value	Paid	Profit	Value	Paid	Profit	Value	Paid	Profit	
1	68			53			36			
2	68			53			36			
3	68			53			36			
4	68			53			36			
5	68			53			36			
6	68			53			36			
7	68			53			36			
8	68			53			36			
9	68			53			36			
10	68			53			36			

If the experiment is being run with simulated buyers, then the instructor or assistant who is determining the number of units purchased from each seller needs a table representing the market demand schedule. This can just be a table of prices with corresponding quantities demanded. As noted above, it may be desirable to adjust the demand schedule based on the number of sellers participating. If you choose to do this, you will need to prepare several demand schedules corresponding to ranges of possible values for the number of sellers, e.g., one for 10-13 sellers, one for 14-16 sellers, and one for 17-20 sellers.

The other paperwork that may be prepared before the experiment is a form on which sellers transmit their decisions in each period to the instructor. It was argued above that written transmission is superior to oral communication since the former allows all sellers' decisions to be revealed at the same time, after all decisions are irrevocable. Sellers may just write their decisions on blank slips of paper saying something like "A: 5 @ 53" to indicate that seller A is willing to sell up to 5 units at a price of 53. Alternatively, the instructor may prepare a large number of printed slips with blanks for period number, seller identification, quantity offered for sale, and offer price. The latter avoids any confusion with sellers forgetting to put their identification on the slip or with confusion about whether a particular slip applies to the current period or is left over from a previous period.

In experiments with human buyers, some method of random selection of buyers must be prepared ahead. Numbered balls, beans, or tokens can be drawn, or many sequences of random orderings can be prepared and printed by a computer program ahead of time.

Running the Experiment

Physical Setup. The ideal physical setup for this experiment will depend on the design chosen. If more than one simultaneous experiment is being run, it is desirable either to use different prices for the two or to run them in physically separated locations, so that price information from one session is not used by participants in the other.

Sellers should be positioned in a way that facilitates the collection of their decision slips. If human buyers are used, they can be further from the front of the room since their communication in the experiment is oral.

Instructions. Once students are seated appropriately, the instructor should slowly and carefully read the instructions, using an example to clarify the importance of not producing units for which marginal cost exceeds price or buying units for a price higher than the marginal value of the unit. The instructions should clearly describe both the behavioral objectives of sellers and buyers and the mechanism of market interaction under the posted-offer institution.

Initial Trading Period. In the initial period, sellers have no idea what the demand for their product is. Emphasize to them that they are likely to make losses in the first few periods (as new firms usually do in their first few years), but that they will gradually acquire more information about market demand and the behavior of other sellers. Each student makes a decision about how many widgets to produce and the price at which they are to be offered. These are written on a slip of paper which is collected by the instructor. Once all sellers have submitted their decisions, they are displayed in a large table such as Table 3-1 on the blackboard and the demand side of the market swings into operation. The instructor should also maintain a record of sellers' decisions, either on a record sheet maintained for that purpose or by keeping the decisions slips.

Human Buyers. If human buyers are used, the random sequencing scheme is then invoked to select the first buyer. The first buyer decides how many units to buy and from what seller(s) to buy them. The instructor and/or assistant then records the decision on a permanent, written experiment record and revises the table on the blackboard to reduce the number of units available from the seller(s) that the buyer chose by the number of units bought. The second buyer is then selected and the process is repeated, continuing until all buyers have shopped.

Simulated Buyers. If simulated buyers are used, then the instructor should sequence the sellers' offer slips from lowest to highest offer price and determine the number sold by each seller by the following method. For the lowest-price seller, determine (from the demand schedule) the demand price of the number of units that seller has produced. If it is higher than the price offered, then all units are sold. If not, then the number sold is the total market demand at that price. Fill in the appropriate number on the seller's slip and the entries required on the experiment record (and on the blackboard if each seller is to be informed of the sales of all sellers).

If the lowest-price seller failed to sell all units produced, then sellers with higher prices will sell zero. This information can be entered and the slips returned to all sellers. The next period then begins. If the lowest-price seller sold all units (the usual case), then proceed to the next-lowest-price seller. Compare this seller's price to the demand price of the cumulative number of units to be sold—the number produced by this seller plus the number sold by the lowest-price seller. If the demand price is higher, then all units sell and we proceed to the next seller. If the offer price is higher, then determine the maximum number of units that can be sold at the offer price and subtract the number sold by the lowest-price seller. The result is the number sold by the next-lowest-price seller. If the next-lowest-price seller fails to "sell out," then all sellers with higher prices will sell zero.

Continue this process as long as additional units can be sold. At each stage, compare the current seller's offer price to the demand price of the cumulative number of units produced (the number produced by the current seller plus the number sold by sellers with lower offer prices). If the offer price is lower, then all units sell and the next seller can be processed. If the offer price is higher, then the number of units sold by the current seller must be determined by finding the maximum cumulative number of units that is demanded at the current seller's offer price and subtracting the number sold by sellers with lower prices. In this case, no seller with a higher price than the current seller will sell any units.

Particularly in the later periods of an experiment, several sellers may offer widgets at the same price. In this case, the group of sellers must be processed together. If not all units produced by the group can be sold, they should be divided as equally as possible among the sellers.

This calculation process should normally be completed in one minute or less in order to keep the experiment moving. It may be useful for the instructor and assistants to practice the process a few times before the experiment to avoid awkward delays.

Once the calculations are complete, each seller should receive whatever information is specified by the design. Minimally, each seller must know how many of her units were sold. Additional aggregate information that may be given to sellers includes the range of prices at which widgets were offered, the range at which widgets were sold, and/or the total number of units sold in the market. It is also quite reasonable to give full information about the prices and sales of all sellers, since this information would naturally be available in the human-buyer experiment.

Since sellers who do not correctly understand the profit-maximizing relationship between marginal cost and price may make errors by producing too much or too little, it may be desirable to remind each student at the end of the first period to check to make sure that the marginal cost of the last unit produced is less than (or equal to) the price for which it was offered. If it seems necessary, the instructor can meet briefly with each seller to make sure that she understands and to suggest a lower (or higher) production level when that seems appropriate. Human buyers may likewise buy units for more then their value, and can be reminded of how to use marginal values to maximize their gains from the experiment. Remember that the purpose of this experiment is to teach students by experience, not to provide a rigorous test of market institutions.

Subsequent Trading Periods. Each subsequent trading period is similar to the first. The pace of the experiment should pick up as students become more confident of the decision process and learn more about the market. Convergence to competitive equilibrium is usually much slower and less reliable than in the double-oral auction experiment.[4] The transmission of information among buyers and sellers is less fluid in the posted-offer market than in the oral auction. If there are relatively few sellers, then changes made by one seller affect the demand curves of other sellers, so it can be difficult for them to learn about their demand. For example, a seller may offer the same number of units at the same price in two successive periods, but sell all units in the first period and none in the second due to price reductions by competing sellers. Essentially, each seller can get reliable demand information only by changing his or her offer *conditional* on the offers of other sellers. Since other sellers are usually changing their offers, reliable information is difficult to obtain in this experiment.

The Report Assignment

As a basis for writing a report, students should have access to the full results of the experiments. If several different groups, classes, or sections run experiements, then it is desirable to provide results of all, so that students can compare the results of their own classes to others. Results should include, for each period, the price and number of units offered and sold by each seller and the price paid and number bought by each buyer. Students should also be provided with a complete set of cost schedules and value schedules for all buyers and sellers in the experiment.

There are several directions in which a student report on the posted-offer experiment might be focused. The concepts of demand, supply, and market equilibrium are present here as in the double-oral auction. If students seem to understand them well

[4]This is consistent with the results of the research literature, see Davis and Holt (1993, 177-84).

based on the double-oral experiment and discussion, then it is probably not necessary to review them at length here. However, it is probably useful for students to compute the equilibrium price for each posted-offer experiment, since this will allow them to evaluate the performance of the experimental market.

If a single-unit experiment is run, then students may compare the results of the posted-offer market with the single-unit double-oral auction. If, as is commonly observed, the posted-offer market converged to equilibrium more slowly, they may consider how the flow of information between buyers and sellers differs between the two market institutions and what effect this might have. If several different market experiments are observed, did the posted-offer markets tend to converge more quickly if the price began below equilibrium than if it started out above of vice versa?

In a multiple-unit experiment, the individual behavior of sellers and buyers is another possible focus. Whether this is their first exposure to the concept of marginal cost and valuation or whether it has been covered in class, a set of leading questions to get them to explain why decisions of buyers and sellers are made on the basis of marginal value or cost rather than total or average will reinforce these important concepts.

Once again, it is useful to ask the students to compare the conditions of the market and the experiment's results to those of short-run perfect competition. A competitive results usually emerges eventually despite a relatively small number of buyers and sellers and a highly-imperfect mechanism for transmitting information. Moreover, it is very important for them to recognize that a seller who "sets" a price in a posted-offer market may still be a price taker in the sense that he is pushed to set the price that has been collectively determined by the market. A question that asks students to consider what would happen (in later periods) if a seller's price deviated substantially from that of other sellers may elicit thought about the distinction between how markets determine prices as opposed to the institution by which they are posted.

Reading and Discussing Student Reports

If students are asked to write about the demand and supply curves, market equilibrium, and gains from exchange, then this assignment can serve as a check on how much they learned from the similar exercise in the double-oral experiment.

If the experiment is performed near the time that perfect competition is covered in class, then the concept of the market supply curve as the aggregation of sellers' marginal cost curves can be easily motivated from the experiment. Students who were successful in constructing the supply curve will have little difficulty recognizing that what they did amounted to marginal-cost aggregation.

As in other experiments, students who participated usually have considerable insider experience with the decision-making process and the transmission of information in the market. They can probably carry on a successful discussion about the process of price dynamics with little guidance from the instructor.

The issue of price-setting vs. price-taking is quite important. Many students view the price-taking assumption of perfect competition as being irrelevant to markets in which they participate simply because sellers usually affix a price tag to each item. This experiment, if it converges to equilibrium, provides a counterexample. Even though each seller sets the price at which she is willing to sell, her decision is tightly constrained by the market. If she sets the price a little to high or too low, substantial profits may be lost. This separation of the process of determination of price from the process of physically posting the price is one of the most important lessons of the posted-offer experiment.

Using the Experiment in Later Classes

This experiment may complement the double-oral auction in establishing in students' minds the relevance of market equilibrium an the notion of price-taking behavior. It also provides a convenient illustration of the theory of market supply under competition. Even if the experiment is done a month or more before the class segment on perfect competition, students are likely to remember enough of the experiment to allow the instructor to use it as an example of how supply curves are constructed from producers' marginal cost curves.

As students begin to study imperfect competition, this experiment provides a good example of the distinction between competition and oligopoly. Asking them whether or not they explicitly took account of the effects their decisions would have on their rival when setting prices and production limits motivates them to think about an important difference between the two market forms. The market-concentration experiment described in Chapter 5 develops this connection more directly by successively concentrating a posted-offer market to oligopoly then to monopoly.

Variations

There are interesting ways in which the structure of the market can be varied if time allows. The amount of time that the basic experiment takes will depend on how quickly the students adapt to the the posted-offer setup and the concept of marginal decision-making, the number of participants (especially in a human-buyer experiment), and the length of time it takes for the market to converge.

Taxes. At some point in the experiment, a per-unit tax on producers can be imposed. Sellers will see this tax as an increase in marginal cost, and will correspondingly produce less and charge a higher price. As a result, fewer units will be exchanged and the market price will rise. This is a good opportunity to demonstrate to students the shifting of the incidence of a tax levied on producers—a sticky point of great importance for public policy. The posted-offer market is well suited to handling the tax distortion since there is an asymmetry between the roles of buyers and sellers that may make it less obvious to students a priori that the costs of the tax will be shared.

The posted-offer setting can also be used to examine the short-run effects of a lump-sum tax on sellers or buyers. Since such a tax does not affect marginal costs or values, it should simply take money away from the taxed individuals without affecting their decisions. However, one must be very careful in the interpretation of such an experiment. A lump-sum tax on the sellers (or buyers) *in an industry* will induce some to leave the industry, which will affect price and quantity exchanged. Since there is no mechanism for entry or exit in the posted-offer experiment proposed here, students may tend to make unwarranted generalizations from the results of the experiment.

Price Ceiling. Introducing a price ceiling into the market forces sellers to cut back on output relative to the competitive case. If price ceilings were performed in the double-oral auction, there may be little gained from repeating the exercise here. However, one interesting feature that arises in the human-buyer version of the experiment is that the selection of which buyers are able to satisfy their demands is explicitly random. This should result in gains from exchange that are far below even the maximum attainable if the units available at the ceiling price are allocated to the highest-value buyers.

Price Floor. A price floor generates an interesting sort of disequilibrium in the posted-offer market. At the higher-than-equilibrium price, each seller would like to sell a lot, but knows that he or she can sell only a little unless he or she can undercut the other sellers. With a homogeneous good and perfect enforcement of the price control (as occurs in the experiment), there is little that the individual seller can do except to set price at the floor and hope for the best. Since sellers are not penalized by incurring cost on unsold units, they can simply set the floor price and limit quantity at the relatively high amount they can supply at that price. (Chapter 4 presents an experiment in which sellers decide on production in advance and incur costs on all units produced, whether or not they are sold.) With all sellers setting the same price, market demand will be divided equally among sellers. The frustration of being unable to use changes in price or output to improve their situation should give students a vivid understanding of the efficiency costs of price controls.

4 A Posted-Offer Market with Advance Production

As noted at the beginning of Chapter 3, the nature of a market depends, among other factors, on what happens to unsold units of the good or service. The simple double-auction and posted-offer experiments of the previous two chapters have operated under the assumption that the good being produced is nonstorable so that units cannot be carried over from one period to the next, and that it is "produced-to-order" so that production costs are incurred only if the unit is sold.

Although this set of assumptions corresponds most directly with the concept of static competition presented in most textbooks, it does not match the features of some industries in the real economy. This chapter explores some classroom experiments designed to demonstrate the differences that are introduced by advance production ("produce-to-inventory") and storability.[5]

The basic setup of the experiment is identical to the posted-offer experiment of the previous chapter, except that sellers set their production levels in advance of the market interaction (corresponding to the limit quantities they set in the earlier experiment) and that costs are calculated based on this ex-ante production quantity rather than the ex-post sales quantity. The behavior of the market may vary considerably depending on the potential for disposing of produced but unsold units.

In the basic experiment described here, unsold units simply vanish, leaving large potential losses for producers. In this setting, sellers quickly learn to underproduce to avoid the possibility of these large losses. This provides a convenient pedagogical entrance for the importance of uncertainty. In particular, the loss function in this model is asymmetric: underproducing costs only the (small) difference between marginal revenue and marginal cost on the missing units; overproducing costs the entire (larger) marginal cost of the extra units. Unsurprisingly, student sellers tend to underproduce, settling the economy at a higher price and a lower quantity than competitive equilibrium.

Two significant variations on this experiment are also developed in some detail in this chapter. One, following Mestelman and Welland (1992), allows producers to have post-market "clearance sales," in which unsold units can be offered at a lower price and buyers then have a second opportunity to buy. Mestelman and Welland find that the introduction of clearance sales brings market efficiency much closer to competitive equilibrium. However, the clearance sales may introduce strategic behavior by human

---

[5]Research experiments relating to this issue have been published by Mestelman and Welland (1987, 1988, 1991) and Mestelman, Welland, and Welland (1987).

buyers who withhold demand in the regular phase of the experiment hoping to get a better price on the clearance sale. If buyers carry this strategy to an extreme, buying no units in the original market, then the clearance sales cease to have any meaningful function. The sales that would occur in the basic experiment occur here, except they occur in the clearance-sale round rather than immediately. With simulated buyers, the instructor must set a buying rule: presumably that buyers do not anticipate the clearance sale and buy all units they want at the prices posted in the first round, then any additional units they want at the lower prices in the clearance round.

The second variation allows unsold units to be carried forward to the next period, either with or without a storage cost. Under this rule, the penalty for overproduction is reduced relative to the basic experiment. Indeed, if marginal cost is linear in output, there are no storage costs, and there is zero probability that this is the last trading period in the experiment, the allocation of output over time is indeterminate. Because of this, it is probably desirable to make sure that at least one of these conditions is violated.

While this experiment can add significantly to the richness of students' understanding of the behavior of producers and markets, it tends to bring in ideas such as uncertainty and risk that are often omitted from introductory courses. Each instructor must decide whether his or her students will find the complications introduced by this experiment stimulating or confusing. None of the subsequent experiments described in this book builds on this experiment, so it can be omitted by those instructors who are pessimistic about its impact. Because it uses the posted-offer market institution, it is highly advisable to run the basic posted-offer market experiment of Chapter 3 prior to this experiment. Not only is it helpful to have students who are familiar with the posted-offer mechanism, but the previous experiment is a good basis against which to compare the results of this experiment.

Overview of the Experiment

As in the previous chapter, the instructor must decide in advance whether human buyers or simulated buyers will be used. Although either is feasible, there are definite advantages and disadvantages of each for the variants of the experiment. In the basic experiment and in the inventory experiment, the role of buyers is identical to Chapter 3, so buyers are likely to become bored quickly. However, using human buyers can lead to very different results in the clearance-sale experiment, since by holding back demand waiting for the sale they may dramatically alter the behavior of the market. The ensuing text will assume simulated buyers except in the clearance-sale experiment. Instructors wishing to use human buyers may easily alter the experiments following the discussion of the previous chapter.

While this experiment could be conducted with single units for each buyer and seller, the environment is much richer with multiple units. With but a single unit, a seller can avoid the risk of overproduction only be withdrawing from the market entirely. Producers of multiple units can cut back production in response to risk, but still participate actively in the experiment. The discussion below assumes that sellers can produce multiple units and that human buyers, if present, can purchase multiple units.

The process of running the experiment is very similar to that of Chapter 3. It may be possible to perform two or even three experiments within a session, depending on the size of the class and the length of the session. One may begin with the basic advance production experiment, then introduce either clearance sales or inventories as a second experiment (or, if time allows, do both variations after the basic experiment). If more than one experiment is run, the parameters of the model should be shifted so that the equilibrium price is different for each experiment. It may also be desirable to exchange roles between experiments: if human buyers are used, to rotate buyers into seller roles and vice versa. If several small groups are running parallel experiments, changing the composition of the groups may avoid situations where individuals begin to become very familiar with the behavior of other individuals and begin to tailor their strategies accordingly.

The competitive equilibrium in the basic experiment is a Nash equilibrium. However, the result to be expected, at least in early periods, is that sellers will produce a smaller quantity and sell at a lower price than competitive equilibrium. This behavior arises because of the asymmetry of risks in the model. As noted in the introduction to this chapter, the penalty for producing more than can be sold at the offer price is much larger than the penalty for leaving excess demand at the offer price. This creates incentives for sellers to set both output and offer price slightly lower than otherwise. The combination of lower output and lower prices makes it likely that a situation of excess demand could be sustained for a considerable number of periods in this experiment as the risks to any seller of raising output and price are great.[6]

To see the nature of the asymmetric risk clearly, consider the sample experiment described in Table 4-1. The top part of the table gives the inverse market demand schedule (showing the highest price that simulated buyers are willing to pay for any given number of units) while the bottom gives the cost curves of four sellers who make up the entire market. The demand curve is highly elastic to avoid giving market power to the four sellers. The market clears at a price of 48 with 12 units exchanged and each seller producing 3 units.

Consider the situation of seller A if all other sellers produce three units and offer them at a price of 48. Table 4-2 shows the profit and the number of unsold units that result from a range of price/output choices for seller A. Note first that profit is maximized at 3 units and a price of 48, the competitive equilibrium. This is consistent with the competitive equilibrium being a Nash equilibrium for this experiment. Under conditions of certainty, seller A would choose (3, 48) and the competitive equilibrium would result. However, notice the shape of the profit function in the neighborhood of the optimum. To the northwest, the profit function is quite flat because lowering price

[6]Mestelman and Welland (1988) find that output is slightly below competitive equilibrium and price is within the competitive equilibrium range in a series of advance-production posted-offer experiments.

or output does not result in unsold units. To the southeast, profit drops precipitously even for small movements. When the seller is subject to uncertainty about the behavior of other sellers or about the exact value of her optimal output and price, she pays a huge profit penalty for setting price or output too high, but very little for setting them too low. Consequently, underproduction and a relatively low price may emerge in this experiment.

Inverse Market Demand Schedule

Units Demanded	Price	Units Demanded	Price	Units Demanded	Price	Units Demanded	Price
1	59	7	53	13	47	19	41
2	58	8	52	14	46	20	40
3	57	9	51	15	45	21	39
4	56	10	50	16	44	22	38
5	55	11	49	17	43	23	37
6	54	12	48	18	42	24	36

Sellers' Marginal Cost Schedules

Unit Number	Seller A	Seller B	Seller C	Seller D
1	35	33	39	34
2	40	39	43	39
3	45	45	47	44
4	50	51	51	49
5	55	57	55	54
6	60	63	59	59

Introducing a clearance-sale market in which buyers are active in both subperiods (the regular trading round and the clearance round) can mitigate this risk, since sellers can ask avoid some of the losses otherwise incurred on unsold units. Likewise, the ability to carry over inventories reduces the cost of unsold units. Both of these variations may allow the market to approach competitive efficiency. Moreover, both are realistic responses to the problem of advance production that students can readily identify in the real world. If students are dismayed by deviations from efficiency that appear in the posted-offer market with advance production, they may have some faith in markets restored by these variations.

Profit Outcomes of Price/Output Choices

	45	46	47	48	49	50	51
0	0	0	0	0	0	0	0
1	10	11	12	13	14	15	-35*
2	15	17	19	21	23	-25*	-75**
3	15	18	21	24	-22*	-70**	-120***
4	10	14	18	-26*	-72**	-120***	-170****
5	0	5	-37*	-81**	-127***	-175****	-225*****

*Each asterisk represents one unsold unit at that price/output choice.

Details of the Design

Since the strategic interaction among sellers is quite complex, it is probably wise to keep the individual seller's task relatively easy. Having a relatively small number of potential units (perhaps in the range of 4 to 8) makes the cost schedule less daunting.

The possibility of large losses due to unsold units creates a problem in experiments in which subjects are paid. For example, in the sample experiment of Tables 4-1 and 4-2, in which the values seem reasonable upon initial examination, a bad decision that is not too far from the optimal one can bring losses in the neighborhood of 100 value units. Since equilibrium profits are only 24 per period, a couple of early losses can cause a seller to get so far in the hole that she has virtually no hope of earning money for the entire experiment, which may nullify the profit motive for her since experimenters are typically reluctant actually to take money from subjects. Since classroom experiments do not usually use money as a reward, this difficulty may not arise; a grade incentive can work just as well to induce participants to avoid losses as to earn profits, for example. However, if an incentive mechanism is used, one must take care in designing it so that one or two periods of large losses do not eliminate the incentive to earn profits in subsequent periods. Designing the experiment to keep expected profits high and expected costs low may help avoid putting sellers into large-loss situations.

The clearance-sale experiment presents unique issues of design. Human buyers may adopt a strategy of underbuying in the first round of trading, waiting for prices to fall in the clearance sale. If they carry this to extreme and purchase nothing in the first round, then the trading institution collapses into a single-round posted-offer market with all exchanges occurring in the clearance round, where prices and quantity should approximate competitive equilibrium (or be slightly lower) as in the basic experiment. Thus, if all purchases are postponed, the clearance round no longer offers bargains to buyers.

The other polar strategy for buyers is to buy their full demand in the initial round as though there were no clearance round, then purchase extra units if the price falls sufficiently in the clearance round. Where between these poles behavior will fall depends on the individual buyers in the experiment.

Since the behavior of human buyers in the clearance-sale experiment is potentially complex, it is more difficult to capture buyer behavior in a simple rule to simulate buyers. If simulated buyers are to be used, buying all units that are profitable in each round is probably the best simple rule to adopt. In subsequent discussion, students (all of whom were sellers) can be asked about how they would have behaved as buyers and how any changes from the simulated behavior should affect the market.

Final-period effects may be important in the inventory experiment. The instructions must specify clearly what happens to units that are unsold in the final period. One choice is that they simply rot. Another is to establish a fixed, positive salvage payment at which the instructor buys leftover units (presumably at a price well below the market price). If leftover units have low value and students know that a particular period is going to be the last (because the number of periods has been specified in advance or perhaps because they realize that the class period is ending), they may lower price and/or quantity as in the basic experiment to avoid the heightened cost of having unsold units.

What Students Should Know

To participate in a multiple-unit experiment, students must be able to make decisions based on marginal cost and value. Whether they need to have covered such concepts in class prior to participating or whether they can learn it "on the fly" will depend on the individual students. Research experimenters typically use untrained subjects and rely on the instructions they give to provide an understanding of the decisions that must be made. However, research experiments have different goals than classroom activities. In the research setting, the experiment is often designed to test whether people can or will make decisions based on marginal concepts. The research experiment can typically run longer than one hour if that is necessary to make sure that all participants understand the situation. In the classroom setting, one is looking for a relatively quick demonstration of how firms that behave in certain ways (i.e., make decisions based on marginal concepts) will interact in a market. This leaves less room for the kinds of errors in decision-making that a totally inexperienced participant might make in the first few periods.

The topics introduced by the presence of asymmetric risk in this experiment are beyond the scope of introductory textbooks. Senior economics majors, graduate students, and even many economics professors have never considered how markets would behave if overproduction led to much greater losses than underproduction. Thus, this experiment does not really "fit" into the introductory curriculum anywhere. The most logical place to insert it is in the discussion of the theory of the firm or of perfect competition.

Instructor Preparation and Materials

The preparation and materials required for this experiment are nearly identical to those discussed in Chapter 3 for the regular posted-offer market. The seller's profit worksheet must be changed slightly to account for the fact that production and sales need not be the same. In experiments with clearance sales or inventory carry-over, these must be incorporated into the worksheet as well. Table 4-3 shows sample worksheets for the three experiments. It is clear from the worksheets that the student's task in computing profits becomes somewhat more complex when clearance sales or inventories are introduced. With clearance sales, there are revenues from both regular sales and clearance sales to add in. With inventories, there are both production costs and inventory carrying costs to subtract from revenues. Sellers might find a calculator handy to help make these calculations.

For students who have already participated in the posted-offer experiment of Chapter 3, there should be little difficulty understanding the market mechanism. Nonetheless, it is important to review how the market works carefully in the instructions, in case someone was absent for the previous experiment or has forgotten how it worked. The instructions must clearly stress the differences between the present experiment and the previous one: the nature of advance production and the disposition of unsold goods. If a different buyer mechanism (human vs. simulated) is being used than was used in the original posted-offer experiment, then detailed instructions about how the demand side works are important.

Running the Experiment

The period-to-period operation of the basic experiment is similar to that of Chapter 3. Participants make large "mistakes" in the first period or two of many experiments because they do not yet know even approximately the equilibrium price level. In this experiment, some will make large losses because they will have unsold units. It is important to make sure that all sellers are calculating profit correctly, basing costs on total production rather than on sales. Stepping through the calculation with them slowly after the first trading period may be useful.[7]

[7]Many research experiments perform a "practice period" in which participants perform their trading and record their results, but which does not "count" toward the final experiment results. Unless formal statistical tests are being applied to the results of the classroom experiment, it seems unnecessary to throw out the first period. Students may choose to ignore it in their analyses, or may contrast the behavior of the market in the first period when mistakes may occur to behavior in later periods.

In the clearance-sale experiment, the regular trading session occurs as in the standard posted-offer market. After all buyers have finished their purchases (in random order, or by the instructor if simulated buyers are used), sellers should mark their sales in the "Regular Sales" column of their worksheets and multiply by price to get "Regular Revenue." Then each seller must submit a price at which she is willing to sell all left-over units in the clearance sale. The clearance sale then proceeds as a second posted-offer market within the period, with either the same or a new random sequence of buyers. After the clearance sale, sellers enter the number of units sold and the clearance price on their worksheets and multiply to get "Clearance Revenue." Adding the two sources of revenue and subtracting total cost gives the final profit or loss for the period.

In the inventory experiment, the usual posted-offer trading period occurs. At the end, some sellers may have unsold inventories. These should be entered on their worksheets as "Ending Inventory" for the current period and "Beginning Inventory" for the next period. If there is a storage charge for holding inventories, this number can be multiplied by the number of units of inventory and the product entered in the "Inventory Cost" column for the next period. Sellers then make and submit their production decisions for the next period and posted-offer trading occurs once again, with each seller offering a number of units equal to the sum of carried-over inventory and current-period production. Profit for each period is revenues less the sum of production costs and inventory costs.

Basic Experiment

Period	Units Produced	Total Cost	Units Sold	Price	Total Revenue	Profit
1						
2						
3						

Experiment with Clearance Sales

Period	Units Prod.	Total Cost	Regular Sales	Price	Regular Revenu e	Clrnce Sales	Price	Clrnce Revenu e	Profit
1									
2									
3									

Experiment with Inventory Carry-Over

Period	**Begin Inv**	**Units Prod.**	**Units to Sell**	**Prod. Cost**	**Inv. Cost**	**Units Sold**	**Price**	**Total Rev.**	**Profit**	**Ending Inv.**
1	**0**									
2										
3										

The Report Assignment

There is much fertile ground for discussion and analysis in the student reports on this series of experiments, but a relative shortage of hard theory in their textbooks to apply to the problems. Thus, a report on the posted-offer experiment may be a wonderful opportunity for the bright student to engage in creative speculation without the constraints of conforming to the textbook. It may also be unusually frustrating for the weaker or less creative student who cannot find the answers to the problems in the textbook. The challenge of designing a report assignment for this experiment is to allow the best students to explore beyond classroom theory while also keeping at least some parts of the assignment within the capabilities of students who are not able to develop their own tentative theories.

An evaluation of the efficiency characteristics of the basic advance-production experiment and a comparison of the outcome to competitive equilibrium should be within the scope of most students, and constitutes a good starting point for the experiment. These results can then be compared to those of the produce-to-order experiment of Chapter 3. Finally, students can be asked to characterize their own behavior in each experiment and how it differed between the two. These comparisons should provide the raw materials from which the best students can identify asymmetric risk as a source of differences in behavior.

Table 4-2, showing the result of various decisions for one seller given equilibrium behavior of other sellers can be useful in showing students the nature of the risk in this experiment. Students could either be asked to prepare such a table (with hints in the instructions as necessary to get them started) or the instructor could prepare the table for the follow-up discussion.

It is also possible to prepare alternative versions of Table 4-2 corresponding to non-equilibrium behavior by sellers B, C, and D. For example, one might redo the table assuming that the other sellers all offer 4 units instead of 3 at a price of 47 instead of 48. This results in large losses (zero sales) if A follows the competitive equilibrium plan of 3 units at 48. This sensitivity of A's profit to the behavior of other sellers underlines the meaning of uncertainty, taking it beyond merely making mistakes about one's optimal decision and considering the effects of one's rivals mistakes as well. One important lesson that can be learned from constructing these alternative tables is that A's optimal production and price decision depends on those of its rivals. This is a useful introduction to the idea of reaction functions and the basis of game theory. It also helps explain why equilibrium may be hard to achieve: it is not in A's interest to make the equilibrium decision unless other sellers do. In terms of two-dimensional game theory, the reaction functions are not vertical and horizontal.

Reading and Discussing Student Reports

Obviously, a sophisticated analysis of the game-theoretic interactions of players in this experiment is not to be expected from introductory economics students. However, having participated as sellers, students should be able to pinpoint why this experiment is different than the earlier posted-offer market. Once they recognize the importance of caution in setting price and quantity, it is not hard for them to see why a situation of excess demand might be sustained in this market.

Although the underlying analytics of this experiment are very complex, the intuitive connection to real-world applications comes more easily. The experiment leads smoothly into a discussion about the benefits to producers of production to order rather than advance production (Why are bananas are produced to inventory but fighter aircraft are produced to order?). If the clearance-sale or inventory variants are performed, these also correspond to ways that real-world producers deal with the risk of being caught with unsold goods.

If one wanted to get into a quite sophisticated discussion of market dynamics, one could think about the idea of randomly occurring shocks to demand (see variations below). Although fluctuating demand is not a feature of the basic experiments that are suggested above, it should not be too hard for students to imagine how conservatively they might behave if the demand curve was subject to periodic, random declines of 20%. One can also introduce the idea that despite demand fluctuations, firms must not lose money on average in the long run or they will leave the industry. Thus, in the long run firms will leave the industry until prices are sufficiently high to earn enough profit in the good periods to cover losses in the depressed periods. This is a rather advanced idea for an introductory student to grasp, but it is not too far removed from the experiment. Moreover, it has very important applications in everyday economics. Just because a personal-computer manufacturer happens to report large profits for one quarter or year does not necessarily indicate that the industry is not competitive. The extension of the ideas in this experiment suggest the possibility that this is just a "high-demand" period in which firms recoup past losses or build up reserves against future bad times.

Using the Experiment in Later Classes

As noted above, this experiment deals with issues that are beyond the analytical level of most introductory texts. However, it can play a very important role in lectures and class discussions on perfect competition and efficiency. This experiment can serve as a "reality check" on the theoretical model of competition and provides examples of variations on the model to account for important features of actual markets.

Perfect competition is presented in the introductory texts as a static model under certainty. It is important for students to recognize that these are simplifications made for analytical purposes and that changing them may cause important changes in the model's conclusions. This experiment serves as a shred of evidence about the effects of relaxing these assumptions in a couple of simple and empirically relevant ways.

There are several important questions that can be asked about markets using this experiment. One is whether the introduction of advance production under uncertainty, which does characterize many real-world markets, dramatically reduces the efficiency of resource allocation. Evidence from research experiments suggests that the loss in efficiency may be relatively modest even compared to situations of certainty. This may help students to believe that the classroom model of competition can be (carefully) extended beyond the bounds of its strict assumption.

However, efficiency comparison between uncertain and certain worlds are not empirically useful unless one is choosing between them. To the extent that the uncertainty in this experiment results from unpredictable behavior by other sellers, one could imagine a central planner reducing this uncertainty by coordinating seller behavior. If the experiment shows relatively small efficiency losses, it may be pointed out that the market is capable of providing the necessary coordination even when important elements of the environment are not known with certainty. If not, then an interesting discussion can ensue about the degree to which a planner could reduce the relevant kinds of uncertainty and what other potential efficiency problems might be introduced by explicit coordination.

The latter parts of the experiment, if they are used, provide examples of how market institutions adapt to reduce the impact of uncertainty under advance-production conditions. Students often enjoy discussing such features of markets as inventory managements and the strategic use of clearance sales. When they have made the relevant decisions in the context of the experiment, they also usually find it fairly easy to understand the fundamental economic principles involved—ideas such as balancing the costs and benefits of carrying inventories and the balancing of sales between regular prices and clearance sales, which can also lead into a discussion of price discrimination.

While these topics take one considerably outside the usual lessons of introductory micro theory, they are useful and relevant. Through discussions such as these, students may begin to gain an appreciation for the richness of the economic environment that we try to caricature in our textbook models. This is likely to add to their interest in economics as well as helping them to see the subject as more relevant to modern decision-making.

Variations

There are several built-in variations in the experimental setup discussed above: human and simulated buyers, ranges of values for inventory costs, and the use of clearance sales have all been considered. As in the earlier posted-offer experiment, one can also vary the amount of information provided to sellers in the simulated-buyer case, either giving full information about prices and amounts produced and sold or only informing them of their own sales.

Sequencing Decision-Making. Another source of variation is in the sequencing of decision-making. In the basic advance-production experiment, production and pricing decision are made and revealed simultaneously. For goods with long production periods, producers may have pretty good information about their rivals' quantities in production before making price decisions. Thus, an interesting variation on the experiment is to separate the decision-making process into two rounds. In the first round, each seller chooses her output level. These are revealed to the other sellers, then each seller submits her price and transactions begin. This sequence of information would allow sellers to lower prices in order to avoid unsold goods if another seller increased production.

Random Fluctuations in Demand. Although this experiment probably introduces enough complexity, it is straight-forward to add explicit randomness to a simulated demand curve in the model. The easiest way is to have two levels of demand: high and low. Sellers know the probabilities associated with the two states, but do not find out which state will prevail until after their decisions have been finalized.

In the basic experiment, sellers must choose between potentially large losses due to overproduction if demand is low and missing sales if demand is high. Depending on the relative magnitude of the losses and the probabilities associated with the two states, they are likely to vary their production levels and prices between aggressive and conservative. This is a good exercise in choice under uncertainty.

The random-demand experiment can also be used with inventories to illustrate the nature of macroeconomic "inventory cycles." If inventories are relatively cheap to carry, then sellers will probably produce for high demand and carry unsold units into future periods. The left-over units will depress production in the period(s) after the decline in demand. Subsequent discussion can elucidate the effects that this would have on suppliers and how the inventory cycle could be transmitted through the economy.

If the sellers' marginal cost schedules are extremely steep and inventory costs are low or nil, then they might be induced to use inventories for production-smoothing. As demand fluctuates between high and low, sellers will minimize costs by producing at an intermediate level and allowing inventories to act as "buffer stocks" to avoid having to match production to sales.

5 *Market-Concentration Experiment*

The multiple-unit, posted-offer experiment in Chapter 3 can be easily extended to examine the effects of market concentration on price, quantity, and efficiency. By successive mergers, the relatively large number of sellers (six or more) of the original posted-offer market can be merged into an oligopoly or duopoly, then into a monopoly. Students find it interesting to compare the efficiency results in the intermediate, oligopoly situation to those at the extremes.

In the monopoly setting, students must attempt to maximize profit by exploring their demand curve. Since the successive mergers will have produced a group of individual sellers who must now act together, it provides an interesting example of group decision-making.

Overview of the Experiment

The initial setup for this experiment is like the basic posted-offer experiment discussed in Chapter 3. Producer/sellers interact in a posted-offer market (without advance production) with a moderate number of sellers (four to eight). After several periods, producers are paired up and merged, with their production capabilities reflecting the combined capacity of both "factories." Finally, all producers merged into a single giant monopoly firm with capacity reflecting all the factories in the market.

As with other experimental implementations of the posted-offer market, either human or simulated buyers can be used in the market-concentration experiment. Since all of the interesting decision-making is on the side of the sellers, it is probably desirable to use simulated buyers. However, it is quite unwieldy at the monopoly stage to have more than about eight merged sellers trying to make a single decision, so even moderate-sized classes will probably need to be divided into two or more possibly identical but completely separate experiments in the two halves of the classroom.

This experiment can be performed effectively with four, six, or eight participants in each widget market, but the number must generally be determined at the time the materials for the experiment are prepared. This makes it somewhat difficult to deal with unpredictable attendance. To have full participation (i.e., no nonparticipating observers) in a class of 30, the instructor might prepare three parallel experiments with eight sellers in each. If more than 24 students attend, the extras can be doubled up with another seller or asked to help with instructor's chores.

Each producer receives an information sheet showing marginal and total cost at various production levels. They have no initial information about the cost functions of other producers or about demand, except that they know the number of producers and are told that the instructor will simulate the total demand of a certain number of buyers who may collectively purchase up to a given maximum number of units if the price is sufficiently attractive. The given maximum is chosen to equal the total production capability of the collection of sellers. In other words, if there are eight sellers who can each produce up to 10 units, then the total potential demand is 80.

In their instructions, sellers are reminded to set a price higher than the *marginal cost of the last unit produced* in order to avoid losing money on producing it. If students have participated in other multiple-unit posted-offer experiments or if they have studied the theory of the firm in class, then this concept should be familiar.

Initially, each producer/seller acts independently. The experiment progresses through several trading periods (at least four or five). Because the the number of units each producer sells depends on the decisions of other producers, each change by one producer affects the market for the others. As one producer responds to last period's price signal, she changes this period's signal to everyone else. Thus, full convergence to competitive equilibrium is a complex interactive process and is unlikely to occur in a few periods. However, the range of prices often settles near the equilibrium range fairly quickly.

After a few periods, the producers are merged by pairs. New information sheets give marginal and total costs for twice as many potential units as before to be produced by the combined factories. There is no change in technology, merely a consolidation of control over the same factories that produced in the first part of the experiment. The experiment runs for several periods with half as many independent sellers, then is concentrated further by merging all producers into a single monopoly seller, with cost curves that consolidate all of the original factories.[8] After several periods of monopoly (which can be quite time consuming if the sellers have strong disagreements about their collective decisions), the experiment ends.

Details of the Design

Number of Sellers. The number of sellers should be even to facilitate mergers in the second part of the experiment—four, six, and eight are plausible numbers, twelve is not totally impossible. Increasing the number beyond six, and especially beyond eight, can be problematic when all are merged to a monopoly. The more individuals that are involved in making the monopoly's decisions, the longer it usually takes them. If time is likely to be a problem, a more streamlined group will work better. Reducing the number of sellers to four makes the first stage of the experiment very concentrated. With six or more, it is more likely that there are enough sellers that each can ignore the effects of her decisions on the decisions of others.

[8]If eight or more sellers were used in the initial periods, it is possible to add an additional round with two sellers, giving four total rounds with 8, 4, 2, and 1 sellers.

Seller X Cost Schedules

Units	Marginal Cost	Total Cost	Units	Marginal Cost	Total Cost	Units	Marginal Cost	Total Cost
1	15^{A}	15	11	42^{B}	345	21	60^{B}	870
2	20^{A}	35	12	45^{A}	390	22	63^{B}	933
3	25^{A}	60	13	45^{B}	435	23	65^{A}	998
4	30^{A}	90	14	48^{B}	483	24	66^{B}	1064
5	30^{B}	120	15	50^{A}	533	25	69^{B}	1133
6	33^{B}	153	16	51^{B}	584	26	70^{A}	1203
7	35^{A}	188	17	54^{B}	638	27	72^{B}	1275
8	36^{B}	224	18	55^{A}	693	28	75^{A}	1350
9	39^{B}	263	19	57^{B}	750	29	80^{A}	1430
10	40^{A}	303	20	60^{A}	810	30	85^{A}	1515

Output	A	B	C	D	E	F
1	15	30	7	24	12	15
2	20	33	14	28	18	20
3	25	36	21	32	24	25
4	30	39	28	36	30	30
5	35	42	35	40	36	35
6	40	45	42	44	42	40
7	45	48	49	48	48	45
8	50	51	56	52	54	50
9	55	54	63	56	60	55
10	60	57	70	60	66	60
11	65	60	77	64	72	65
12	70	63	84	68	78	70
13	75	66	91	72	84	75
14	80	69	98	76	90	80
15	85	72	105	80	96	85

Cost Schedules. A simple quadratic total cost function (linear marginal cost) works well in this experiment, though other specifications would probably be equally effective. The cost schedules for six sellers shown in Table 5-1 have linear (increasing) marginal cost and zero fixed cost. They have been designed so that each seller's marginal cost curve reaches a common level at an intermediate range of output. Some sellers' MC curves are quite steep; others are much flatter. This leads to wide differences in the amount of profit that they earn in equilibrium. In experiments in which actual payments are made, fixed costs should be used to even out (and reduce) the equilibrium earnings of sellers.

For later stages of the experiment, the cost functions of merging sellers are merged together. Table 5-2 shows the information sheet for seller X, which results from the merger of sellers A and B in Table 5-1. To clarify how seller X's cost schedule is constructed, superscripts have been added to the cost figures showing whether the unit in question is produced by the factory of seller A or B.

Demand Function. The same demand function is generally used at all stages of the experiment. Since the focus of the experiment is on the effects of market concentration, it is probably unwise to confound these effects with a shift in demand. The step-function equivalent of a linear function is strongly recommended, such as the one shown in Table 5-3. Not only is the linear function simple, but smoothness of the demand function is of critical importance in the monopoly stage of the experiment. At each level of output, the marginal revenue of a monopoly depends both on the demand price and the change in the demand price from the previous unit produced:

$$MR = Q \text{ H } P(Q) \; ! \; (Q \, ! \, 1) \text{ H } [P(Q) \; ! \; P(Q \, ! \, 1)],$$

where $P(Q)$ is the inverse demand function. If price declines linearly with increases in output, then the bracketed term is a constant, so marginal revenue declines smoothly. If $P(Q) \; ! \; P(Q\text{-}1)$ increases with Q (as along a concave demand curve), then revenue declines even faster than in the linear case, which should not confuse the experiment but will result in the monopoly's profit-maximizing level of output being smaller. However, problems result if $P(Q) \; ! \; P(Q\text{-}1)$ declines for some values of Q. If this case, marginal revenue may not be uniformly declining and multiple local maxima are possible. This is *very* confusing to student-monopolists who are groping along their (unknown) demand function and should be avoided. Since it is almost impossible to design a quasi-linear demand schedule for an uncertain number of buyers, human buyers should be used in this experiment only when the number of buyers can be predicted confidently prior to the experiment.

Price	Demand Quantity	Price	Demand Quantity	Price	Demand Quantity
90	2	75	32	60	62
89	4	74	34	59	64
88	6	73	36	58	66
87	8	72	38	57	68
86	10	71	40	56	70
85	12	70	42	55	72
84	14	69	44	54	74
83	16	68	46	53	76
82	18	67	48	52	78
81	20	66	50	51	80
80	22	65	52	50	82
79	24	64	54	49	84
78	26	63	56	48	86
77	28	62	58	47	88
76	30	61	60	# 46	90

What the Students Should Know

Participation in this experiment is enhanced by significant background in the relevant economic theory. It helps for students to be comfortable with the principle that profit is maximized when price (or marginal revenue) equals marginal cost. Failure to apply this principle consistently may lead to erratic behavior by one or more sellers, which will impact the demand for the output of other sellers in this experiment.

If students have participated in the posted-offer market of Chapter 3, they should be able to handle this experiment even without substantial classroom exposure to the relevant concepts. However, they may overlook the possible gains from monopolization if they continue to apply the price = marginal cost rule blindly. However, such naivete can be useful pedagogically—in the follow-up analysis students can be led to the observation that higher profits were attainable by raising price and restricting output.

Students should compute their profit each period so that they can track how they are doing. They may find it desirable to have calculators to facilitate these computations.

The instructor or assistant should check the first period's calculations (or step through them with the group) to assure that they instructions have been correctly understood.

Instructor Preparation and Materials

The main preparation for this experiment consists of preparing the information sheets containing the cost information for sellers for each phase of the experiment. For each experiment (remember that it may be necessary to divide large classes into two or more separate experiments), one set of (four, six, eight, or twelve) disaggregated seller information sheets for the first stage, one set of (two, three, four, or six) information sheets for the second stage, and several copies of the single monopoly information sheet must be prepared.[9] Sellers must also have slips of paper on which to communicate their decisions about price and production to the instructor. In addition, the instructor or assistant who is running each experiment must have a sheet of instructions to read, a table giving the demand schedule, and a sheet on which to record the results of each period of the experiment.

Running the Experiment

Setup. The physical setup of the experiment can vary depending on the nature of the room and its furniture. If more than one experiment is being run in the room, they should be separated as much as possible to avoid sharing of information and possible confusion. Within each experiment group, the sellers should initially be separated sufficiently to allow each to make private decisions without being observed by other sellers. In the second stage, the two paired sellers simply move together physically as their factories merge. In the monopoly stage, all sellers move close together to facilitate group decision-making. The instructor or assistant running each experiment should be located near the sellers, but far enough away to assure that the sellers cannot observe his or her actions.

Once students are seated appropriately, the instructor should slowly and carefully read the instructions, using an example to clarify the importance of not producing widgets for which marginal cost exceeds price. The instructions should also include a suggestion to make decisions relatively quickly, so that the experiment progresses

[9]If an additional stage of concentration is inserted, then a separate set of sheets must, of course, be prepared for it as well.

rapidly. It is not necessary to tell students at the beginning that the number of sellers will be reduced at the experiment progresses. Since the number of sellers in each experiment must be set ahead of time and the number of students will usually not be an integer multiple of that number, it is usually necessary to "double up" students in a few cases. This should also be done with any late-comers, so that any students who may have missed the instructions can learn the process from a partner who was present.

Initial Trading Period. In the initial period, sellers have no idea what the demand for their product is. Emphasize to them that they may make losses in the first few periods (as new firms usually do in their first few years), but that they will gradually acquire more information about market demand and the behavior of other sellers. Since this experiment is often difficult to complete in fifty minutes, it may be necessary to hurry them a little in the first period or two; there is little to be gained by agonizing over a decision about which you have little relevant information.

Each student makes a decision about how many widgets to offer and the price at which they are to be offered. These are written on a slip of paper which is collected by the instructor. Once all sellers have submitted their decisions, the instructor should sequence them from lowest to highest offer price and determine the number sold by each seller by the following method.

For the lowest-price seller, determine (from the demand schedule) the demand price of the number of units that seller has produced. If it is higher than the price offered, then all units are sold. If not, then the number sold is the total market demand at that price. Fill in the appropriate number on the seller's slip and the entries required on the experiment record.

If the lowest-price seller failed to sell all units produced, then sellers with higher prices will sell zero. This information can be entered and the slips returned to all sellers. The next period then begins. If the lowest-price seller sold all units (the usual case), then proceed to the next-lowest-price seller. Compare this seller's price to the demand price of the cumulative number of units to be sold—the number produced by this seller plus the number sold by the lowest-price seller. If the demand price is higher, then all units sell and we proceed to the next seller. If the offer price is higher, then determine the maximum number of units that can be sold at the offer price and subtract the number sold by the lowest-price seller. The result is the number sold by the next-lowest-price seller. If the next-lowest-price seller fails to "sell out," then all sellers with higher prices will sell zero.

Continue this process as long as additional units can be sold. At each stage, compare the current seller's offer price to the demand price of the cumulative number of units produced (the number produced by the current seller plus the number sold by sellers with lower offer prices). If the offer price is lower, then all units sell and the next seller can be processed. If the offer price is higher, then the number of units sold by the current seller must be determined by finding the maximum cumulative number of units that is demanded at the current seller's offer price and subtracting the number sold by sellers with lower prices. In this case, no seller with a higher price than the current seller will sell any units.

Particularly in the later periods of an experiment, several sellers may offer widgets at the same price. In this case, the group of sellers must be processed together. If not all units produced by the group can be sold, they should be divided as equally as possible among the sellers.

This calculation process should normally be completed in one minute or less in order to keep the experiment moving. It may be useful for the instructor and assistants to practice the process a few times before the experiment to avoid awkward delays.

Once the calculations are complete, each seller is given information about the number of units sold. Typically, full information about all sellers is given publically by writing quantities sold and prices for each seller on the blackboard, however the experiment can also be run with less information. If two or more simultaneous experiments are being run in the same room, then the experiments should either have different equilibrium prices and quantities (so that looking at the other experiment's results would not be relevant) or the method of reporting the results back to sellers should not allow those in the other experiment to see.

Since students who do not correctly understand the profit-maximizing relationship between marginal cost and price may make errors by producing too much or too little, it may be desirable to remind each student at the end of the first period to check to make sure that the marginal cost of the last unit produced is less than (or equal to) the price for which it was offered. If it seems necessary, the instructor can meet briefly with each seller to make sure that she understands and to suggest a lower (or higher) production level when that seems appropriate. Remember that the purpose of this experiment is to teach students by experience, not to provide a rigorous test of market institutions.

Subsequent Trading Periods. Each subsequent trading period is similar to the first. The pace of the experiment should pick up as students become more confident of the decision process and learn more about the market. Changes made by one seller affect the demand curves of other sellers, so it can be difficult for them to learn about their demand. For example, a seller may offer the same number of units at the same price in two successive periods, but sell all units in the first period and none in the second due to price reductions by competing sellers. Essentially, each seller can get reliable demand information only by changing his or her offer *conditional* on the offers of other sellers. Since other sellers are usually changing their offers, reliable information is difficult to obtain in this experiment.

Merging to Oligopoly. After several periods of "competition" with the full number of sellers acting independently, the instructor announces that mergers are occurring: Sellers A and B merge to form Seller X, Sellers C and D to form Seller Y, etc. Ideally, the initial mergers should occur after competitive equilibrium has been established, but since convergence to competitive equilibrium in posted-offer markets is sometimes slow and unreliable, this may not be possible.

New cost information sheet are passed out reflecting the combined production capacity of the paired sellers. The students will typically not see immediately how the joint cost schedule is obtained from their individual schedules. If they are curious, it can be easily explained that the two original factories are still operating with the same costs as before, but they are now controlled jointly. Showing them how the units produced by the original factories have just been resequenced in order of increasing marginal cost should demonstrate the nature of the merger.

In the first period after merger, some pairs may be somewhat tentative as they begin trying to make joint decisions for the first time. Their actual decisions are often the aggregation of the decisions that they made individually in the last period under competition. (Is this true of real-world mergers as well?)

This stage of the experiment usually runs for three to five periods. Behavior can vary from a continuation of the approximately competitive market of the first stage to various kinds of explicit or implicit collusion. The instructor should exercise judgment about whether the additional periods will have pedagogical value for the students. If the market is doing something interesting that will provide useful material for later analysis, this stage should be prolonged. If it is just replicating the earlier stage, then it is best to move on to monopoly.

Merging to Monopoly. In the final phase of the experiment, all sellers merge into a single monopoly enterprise with the production capacity of all the original factories. The instructor hands out several copies of the single monopoly cost information sheet. The principal difficulty with the monopoly stage of the experiment is encouraging efficient decision-making by the "committee" controlling the monopoly. One individual should be appointed as the "chief executive officer" of the monopoly and charged with getting a consensus of the group and reporting the decision to the instructor. They should be reminded that only by trying various possible levels of output and price can they find out about the market demand curve and the levels of profit associated with various choices.

At this point in the experiment, the interaction between the instructor and the group can be quite intimate. As soon as the group makes a decision, it can be orally reported to the instructor who can immediately refer to the price on the demand schedule and respond orally with the number sold. In this way, many periods can be completed in a few minutes, but only if the students can make decisions quickly.

In the "heat of battle," students will usually not be able to derive the marginal revenue curve for the monopoly as they accumulate information about the demand curve. However, they sometimes find the profit-maximization price/quantity combination by trial and error if enough periods are performed. A crucial step in obtaining and processing demand information is recognizing the only when they offer *more* units than are sold do they actually find out a point on their demand curve. If all offered units are sold, they cannot be sure that even more units would not have been bought at that price had they been produced.

If the experiment is constrained to one hour or less, the monopoly stage may be truncated by the end of the time period. If it should happen that the monopoly converges to the profit-maximizing equilibrium (and recognizes it) before the expiration of the class period, the experiment can be ended early or one of the variations discussed below can be implemented.

The Report Assignment

If students are to write a report on the results of the experiment, they should be provided with all the relevant information, which includes the cost schedules, the market demand schedule, and a record of the number of units produced and sold and the prices charged by each seller in each period of the experiment. Cost information can be provided in detail (as in Table 5-1, for each seller at each stage of the experiment) or in summary (just the monopoly's cost schedule, which incorporates all of the original factories). With the more limited information, students can calculate the competitive market supply curve and the monopoly's profit-maximizing levels of output and price from the monopoly's cost schedule, though they cannot evaluate the performance of individual sellers in the pre-monopoly stages. If the students have access to a computerized spreadsheet program, it may be helpful to provide the cost and demand schedules to them in that form.

A set of leading questions helps to focus the students' attention on important lessons that can be learned from this experiment. They should be led to find the levels of output and price that would prevail under competition or monopoly and to compare the actual results of the various stages of the experiment to these benchmarks. They should also be encouraged to think about the implications of competition and monopoly for consumer surplus and aggregate profit. It should be noted that even in the "competitive" case profit will not be zero unless fixed costs have been rigged to produce this result. All of these exercises can be easily conducted on a computer spreadsheet with the number of units produced and sold in one column and the associated cost and demand price information in other columns. A large table can easily be generated with the levels of profit and consumer surplus for each hypothetical level of output.

Regardless of the actual experiment results, this report provides a numerical exercise in evaluating the behavior of a monopoly producer versus a competitive market. To the extent that they can fit the results of the experiment into the framework of the theory, they can evaluate the performance of the posted-offer institution and of the class members who participated.

Reading and Discussing Student Reports

Student performance on this report may vary considerably. Strong students are usually able to create a table of profit and consumer surplus at each level of output. This allows them to calculate the profit-maximization solution readily. Many of the students who successfully find the point of profit maximization will not use marginal revenue to obtain it. To get the competitive equilibrium, they must recognize that the supply curve is the marginal cost curve—some will miss this and many will not note that the sum of profit and surplus is at its maximum there.

The instructor should enter the discussion of the experiment prepared to demonstrate the properties of competitive and monopoly equilibrium for the experiment. For the analysis of competition, students should be reminded that a competitive firm's supply curve is its marginal cost curve (average variable cost is less than marginal cost for all levels of output in this experiment). The marginal cost curves of the firms can then be aggregated to get the market supply curve, which will coincide with the marginal cost curve of the monopoly. This sets the stage for conducting further analysis using only the monopoly cost information and the demand schedule.

The large spreadsheet table discussed above may be a useful handout. It shows output, demand price, total revenue, marginal revenue, total cost, marginal cost, profit, consumer surplus, and total gains from exchange for each level of output. The students can see that the highest level of profit is achieved where marginal revenue equals marginal cost, but that the highest level of gains from exchange (profit plus consumer surplus) occurs where price equals marginal cost. This example can then illustrate the nature of the deadweight loss from monopoly—the reduction in consumer surplus is larger than the increase in profit.

Using the Experiment in Later Classes

Depending on its results, this experiment can illustrate several principles. If the results of the first and second stages of the experiment are similar, it can be pointed out to students that (at least under certain circumstances) two or three sellers can be enough to lead to a roughly competitive outcome. This is an important issue since it is difficult to sell the assumptions of perfect competition as realistic, yet we spend many class hours developing the desirable implications of competitive markets.

As noted above, the experiment is an extended numerical example of monopoly theory, allowing the evaluation of both competitive and monopoly solutions. Even if no report is assigned or discussed in lab, the experimental market can be a useful and familiar classroom example.

Variations

Unless more than an hour is available for this experiment, it is unlikely that there will be time for any interesting variations to be performed within a single lab. However, there are interesting ways in which the structure of the market can be varied if time allows.

Taxes. At some point in the experiment, a per-unit tax on producers can be imposed. Sellers will see this tax as an increase in marginal cost, and will correspondingly produce less and charge a higher price. This is a good opportunity to demonstrate to students the shifting of the incidence of a tax levied on producers—a sticky point of great importance for public policy.

Price Ceiling. Introducing a price ceiling into the market forces sellers to cut back on output relative to the competitive case. However, introducing a price ceiling in the monopoly setting may enhance gains from exchange. A monopoly price ceiling experiment may be useful in a discussion of optimal regulation of monopoly.

Price Floor. If there are no human buyers in the experiment, the unsatisfied demand in the price ceiling case may go unnoticed. However, a price floor generates an interesting sort of disequilibrium in the competitive stage of the market. At the higher-than-equilibrium price, each seller would like to sell a lot, but knows that he or she can sell only a little unless he or she can undercut the other sellers. With a homogeneous good and perfect enforcement of the price control (as occurs in the experiment), there is little that the individual seller can do except to set price at the floor and hope for the best. With all sellers setting the same price, market demand will be divided equally among sellers. The frustration of being unable to use changes in price or output to improve their situation should give students a vivid understanding of the efficiency costs of price controls.

Divestiture. In the experiment outlined above, the seller side of the market becomes increasingly concentrated. The results might be quite different if the experiment were performed in reverse sequence, beginning with monopoly and then divesting the individuals into independent firms. Time permitting, this could also be implemented by appending an additional oligopoly stage after the monopoly stage of the original experiment. With knowledge of what level of price maximizes industry profit (assuming that the monopoly converged to that level), the oligopolists may be more likely to attempt to collude to retain the previous high prices. Each oligopoly should see clearly that it is costly as an individual to keep price that high, given the behavior of the other sellers. This should illustrate clearly the incentives toward instability inherent in cartels.

6 Product-Quality Experiment

This product-quality experiment is very similar to the double-oral auction experiment of Chapter 2, but widgets are now of two quality levels: gold and silver. Each buyer and seller can make one transaction in each trading period, buying or sellering either a gold or a silver widget. Separate markets in gold and silver widgets occur simultaneously, with each buyer and seller bidding or offering widgets of one or the other quality level at will in order to maximize profit.

The pedagogical lessons of this experiment focus on allocative efficiency. Students should observe from the results of the experiment that well-informed auction markets can be very successful at choosing the level of product quality that maximizes social gain and at matching up buyers and sellers according to their comparative advantage at producing and preferences for consuming various goods.[1]

Overview of the Experiment

In the product-quality experiment, the group is divided into equal numbers of buyers and sellers. Each buyer is given a demand value representing willingness to pay for a silver widget and a higher demand value for a gold widget. Each seller has a cost of production number for a silver widget and a higher cost number for gold. Each participant may buy or sell one widget of either kind in each trading period. Traders interact in a double-oral auction.

The experimental session consists of three separate experiments of common form. Each participant receives a different set of values for each experiment. In the first experiment, the values are such that the greatest gains from exchange are achieved if all buyers and sellers exchange gold widgets. In the second experiment, the values are adjusted so that all buyers and sellers gain more from trading in silver at equilibrium prices. The third experiment challenges the market to produce an efficient outcome when the comparative advantage of some buyers and sellers is in gold and that of others

[1]Although this experiment was developed independently at Reed (at the suggestion of one of the introductory students who participated in an early set of experiments), it is similar in design to the research experiment of Lynch, Miller, Plott, and Porter (1986). The emphasis of the research experiment is on the effects of accurate information about quality, while the basic experiment suggested here imposed perfect quality information. An extension to imperfect information along the lines of the Lynch experiment is proposed in the final section of this chapter.

is in silver. The record of this experiment in class labs at Reed College is that the double-oral auction reaches the efficient choice of product quality quickly and reliably in each of the three experiments. Student participants rarely buy or sell the less profitable grade of widget beyond the first trading period.

The results of the experiment can be used to illustrate principles of efficient resource allocation. For this purpose, the experiment can be interpreted in either of two ways. Under one interpretation, the gold and silver widgets represent any pair of commodities. Consumers enter the market with established demands for the two; sellers enter with known production costs for them; the market must match buyers and sellers in an optimal way to assure that maximum gains from exchange are realized in the two markets. The other interpretation treats the gold and silver widgets as different quality levels of the same good. Consumers are going to buy a widget, but must choose whether to buy the "standard" (silver) one or the "premium" (gold). Suppliers can produce either kind. The market must match up those buyers (if any), who are willing to pay a sufficient amount for the added quality of the gold widget with those sellers able to produce the gold widgets at the lowest cost (relative to silver).

Details of the Design

The original double-oral auction experiment (described in Chapter 2) demonstrates to students that such markets usually converge to a market-clearing equilibrium. This point need not be reiterated here. Thus, there is no gain to having individuals be excluded from participation in the market by high seller cost or low buyer value. If all participants can profitably participate in either market, then the demand and supply curves for each market will resemble those in Figure 6-1, with the range of equilibrium prices given by (P_1, P_0).

As in the double-oral experiment, the number of participants can vary over a wide range without affecting the viability of the experiment. However, note that the equilibrium in Figure 6-1 requires that the number of buyers and the number of sellers be equal. If there is an "extra" buyer or seller, then the equilibrium becomes a swastika (if the horizontal steps include more than one participant at each price) and one participant who could have made a profitable transaction will be disappointed in each period. It also implies that there is a unique equilibrium price at (or just inside) one boundary of (P_1, P_0). This means that one or more participants will have a negligible profit at equilibrium and may decide not to trade. If the horizontal steps in Figure 6-1 represent single individuals, then the extra buyer or seller will be excluded from the market by his or her poor value. Thus, it is desirable to have equal numbers of buyers and sellers. If an odd number of students are present, one student should be designated as an assistant, helping to record transactions or to post them on the blackboard.

The range of equilibrium prices should be sufficiently wide to assure that everyone trades. Again, convergence to equilibrium was established in the double-oral auction experiment; the emphasis here is on choosing *which* market, not *whether* to trade.

Table 6-1 shows the buyer and seller values for a 20-participant experiment for each of the three separate parts of the experiment. Values vary little among individual buyers and sellers; the vertical steps in the demand and supply schedules are only 1 cent. The equilibrium range of prices at which all participants can make a profitable transaction is relatively wide: at least 10 cents except for gold in Experiment #2 and silver in Experiment #3. If fewer than twenty students participate, the range will be wider since the higher-numbered buyers and sellers are further to the right on the demand and supply curves.

Notice that in Experiment #1, each buyer values gold widgets by 30 cents more than silver, while it costs each seller only 20 cents more to make gold. An efficient allocation of resources thus has all sellers and buyers transacting in gold widgets. In Experiment #2, this is reversed: although gold is now even more valuable to buyers (35 cents more than silver), if costs 45 cents more for each seller to produce. Thus, all participants should shift to silver in Experiment #2. In the last experiment, buyers 1, 3, 5, 7, and 9 have 40 cent preferences for gold, but buyers 2, 4, 6, 8, and 10 are only willing to pay 20 cents extra for gold. Sellers A, C, E, G, and J can produce gold for 20 cents more than

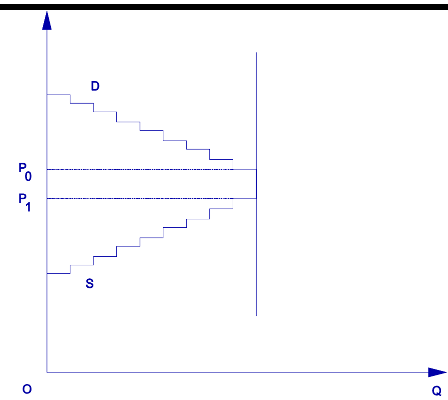

silver, but it costs the other sellers 40 cents more to produce gold. Thus, to maximize gains from exchange, the odd-numbered buyers should buy gold from sellers A, C, E, G, and J, with the even-numbered buyers buying silver from the remaining sellers.

What the Students Should Know

	Experiment #1		Experiment #2		Experiment #3	
	Gold	Silver	Gold	Silver	Gold	Silver
Buyer Values						
1	95	65	96	61	92	52
2	94	64	95	60	82	62
3	93	63	94	59	100	60
4	92	62	93	58	90	70
5	91	61	92	57	98	58
6	90	60	91	56	88	68
7	89	59	100	65	96	56
8	88	58	99	64	86	66
9	87	57	98	63	94	54
10	86	56	97	62	84	64
Seller Costs						
A	65	45	84	39	68	48
B	66	46	85	40	78	38
C	67	47	86	41	60	40
D	68	48	87	42	70	30
E	69	49	88	43	62	42
F	70	50	89	44	72	32
G	71	51	80	35	64	44
H	72	52	81	36	74	34
J	73	53	82	37	66	46
K	74	54	83	38	76	36

To participate in this experiment requires little economics preparation on the part of the students. However, previous experience in a double-oral auction setting will make the experiment move more quickly. A buyer must be able to observe the prices offered by sellers in the market and calculate whether buying silver or gold is more advantageous. A seller must make a similar calculation based on the bids of buyers for the two grades of widgets. Students who are already familiar with the basic double-oral auction will find this extension less difficult than those who are participating for the first time.

To understand and interpret the results of the experiment requires that students be able to evaluate gains from exchange from having individuals participate in one market or the other, and to recognize the relationship between gains from exchange and allocative efficiency. This concept is often introduced early on in an introductory course (e.g., with the introduction of comparative advantage), but a deeper understanding of the concepts of this experiment is usually obtained after discussion of general equilibrium under perfect competition. Again, participation and analysis of the earlier double-oral auction experiment should provide a solid foundation for analysis of the product-quality experiment.

Instructor Preparation and Materials

The instructor must prepare information sheets for each of the three experiments. If the students are accustomed to the double-oral setup, elaborate worksheets for them to use to record results and calculate profit are probably unnecessary. Three separate sheets should be used for the three experiments. (Do not give out all three sets of values at the beginning.)

The instructions must clarify the nature of the dual market in widgets. Students should be told explicitly that they should examine potential transactions in both markets to determine which has higher profit for them.

Running the Experiment

As always, the first step is to assign buyer/seller roles and hand out information sheets to the students. As noted above, it is desirable to have an equal number of buyers and sellers, so if an odd number of students are present one may be assigned to serve as the instructor's assistant.

Once sheets have been distributed, the instructor should read the instructions and respond to questions from the students. If the students are experienced in double-oral auctions, the experiment should progress smoothly. However, it is a good idea to remind them after the first trading period (and occasionally afterward if it seems necessary) to attempt to trade in the grade of widget that yields them the higher profit given the prices that prevail in the market for both grades.

The Report Assignment

The report assignment for the product-quality experiment can focus on the issue of allocative efficiency. Maximization of gains from exchange (profits in the experiment) can be assessed quite easily. Students should be asked to compare the actual result of each experiment (amount of each grade of widget traded, price, and gains to buyers and sellers) with the optimal (competitive equilibrium) result.

Having established the ability of the market to approximate optimal exchange (assuming that this result was achieved), students should be encouraged to think about "the optimal quality of widgets." In the first two parts, there is a unique optimal quality level, since either gold or silver is preferred (at equilibrium prices) by every buyer and seller. In the third, some buyers prefer gold enough to buy it while others do not, and some sellers are more easily able to supply the gold widgets.

Students may also be encouraged to consider whether the experiment would have been any different had the gold and silver widgets been completely different products: gadgets and gizmos, for example. The results should be the same, though a subtle assumption is being sneaked in: that the two commodities are close enough substitutes that a consumer would buy exactly one unit of the two commodities.

Reading and Discussing Student Reports

If students understand the basics of market equilibrium and gains from exchange, they are usually able to work out what the experimental market has done and what it means. Since this experiment usually comes later in the course, these concepts should be familiar and much of the analysis will be review.

If the students are having difficulty with the idea of an optimal level of quality, it may be helpful to present the analysis in terms of how much society is willing to pay for high quality (gold) rather than low quality (silver). Demand and supply curves for quality can be derived for each experiment. In experiments 1 and 2, these curves will be horizontal lines (extending out to the number of buyers or sellers) since each participant attaches the same value or cost to quality. For experiment 1 shown in Table 6-1, all buyers are willing to pay 30 cents extra for the extra quality represented by gold, so the demand curve for quality is horizontal at 30 cents out to the number of buyers, at which point it drops to zero. Similarly, since each seller can produce high quality at an additional cost of 20 cents, the supply curve for quality is horizontal at 20 cents out to the number of sellers, at which point it moves vertically upward. This analysis shows that the "equilibrium quantity of quality" is for all units exchanged to be of high quality. Repeating this analysis for experiment 2 gives the result that the supply curve for quality lies completely above the demand curve, so the equilibrium is zero high-quality units. In experiment 3, the demand and supply curves are two-level step functions that cross halfway out, showing that half of the widgets exchanged should be of high quality.

The main idea that comes out of this analysis is that competitive markets can balance buyers' willingness to pay for quality with sellers' costs of producing higher-quality goods. Once students have mastered this concept, they can be asked to apply it to more realistic settings and to consider conditions under which it might not work.

Most consumer goods come in a variety of forms with varying qualitative characteristics. Students should be able to extend this analysis to consider such quality problems as safety enhancements to autos and other products, durability properties of light bulbs, batteries, and other goods, strength of such items as plastic bags and paper towel, etc.

Perhaps the most crucial assumption that underlies the setup of this experiment is that of perfect information about quality. The students can think about how the experiment would have turned out if all sellers had to write down whether they were producing a gold or silver widget at the beginning of each period, but could then advertise widgets as gold even if they actually produced silver (without the buyer knowing). If buyers cannot realiably observe the higher quality, only silver widgets will be exchanged. Might some sellers develop a reputation for honesty that would convince buyers to buy gold widgets from them?

Discussion of the large issue of government regulation of quality and provision of quality information about commodities can be motivated by these questions. In particular, students should be convinced by these experiments that markets are capable, at least under the most favorable circumstances, of dealing in a desirable way with goods that vary in quality. When information about quality is imperfect, however, market failure can result.

Using the Experiment in Later Classes

By thinking of gold and silver widgets as a generic pair of distinct goods, this experiment can be used in a general discussion of allocative efficiency of competitive equilibrium. The idea that competitive markets induce each individual buyer to choose the "cheapest source of utility" and each seller to produce the good that satisfies the "most utility at the lowest cost" are effectively illustrated. Since the notion of general equilibrium and its efficiency implications are difficult ones to break down into simply understood pieces, the integrative, experiential learning provided by this experiment is particularly valuable.

Moreover, students sometimes enter the introductory economics class believing that the market-determined levels of quality for many goods are sub-standard. This experiment should help convince them that it is some market imperfection (such as lack of information), rather than simply quality variation, that upsets the optimal behavior of competitive markets. This understanding contributes to a more informed discussion of the government's possible roles in such markets: enforcing warranty contracts and laws regarding fraud, providing certified information about product quality, or intervening directly to enforce a specific minimum standard of quality.

Variations

Numerous variations of this experiment are possible. With a relatively large group of students (perhaps 30-50), one could introduce more than two levels of quality. This would make the student's chore a little more difficult, because he or she would have to compare three or more alternatives to find the one that yields the greatest benefit. However, as long as students are able to buy and sell rationally, the experiment should function as effectively as with two quality levels, and may be a more convincing demonstration of the power of markets.

A very promising variation, which gets at the regulation issues discussed above, is to alter the information transmission between buyers and sellers in the model. Several variations of this kind could be attempted. One would involve relaxing the "truth in advertising" rule that is implicit in the experiment.

In an experiment with unregulated advertising, sellers would write down before each period whether they were actually producing a gold or silver widget. Moreover, sellers are not allowed to show the slips on which they wrote their decision to buyers, so there is no reliable way for buyers to discern quality. Seller can say whatever they wish about the widget and buyers do not find our until after making the purchase what the true quality of the widget is. In the experiments of Lynch et al. (1986), this setup led to the familiar "lemons" result, in which the market for high-quality widgets disappeared completely because buyers could not be persuaded that any given widget was gold.

Another, rather trivial, variation would be to have the "government" enforce a quality standard prohibiting the sale of low-quality (silver) widgets. This might be most effective in the third experiment, where some buyers and sellers were already trading in gold, but others were better off exchanging silver. Because it is a very simple extension of the basic experiment, the time cost of running this variation is very low. This variation could easily be implemented quickly (for two or three trading periods) at the end of the regular experiment. Although this kind of regulation is, in the stark experimental setting, obviously inferior to the market outcome with reliable product-quality information, it is a common government response to quality concerns. Follow-up questions in the report could elicit student response on the social losses involved in such rules.

7 *Quality-Signalling Experiment*

The product-quality experiment described in Chapter 6 focuses on the ability of markets to deal with the pricing and resource allocation problems when two different grades of quality are available. However, products in naturally occurring markets usually vary in quality across more than one dimension—autos have characteristics of power, reliability, comfort, style, and size that all affect desirability. Moreover, some of these characteristics are highly subjective, so that it is impossible for sellers or a third party to provide a satisfactory measure of quality. This is particularly true of so-called experience goods—those that must be experienced in order to determine quality.

The quality-signalling experiment is an exercise evaluating whether various bases on which consumers may make quality judgments about products lead to accurate decisions.[2] Several groups of students are asked to rank various brands of a food product based on a single characteristic, each group using a different criterion. For example, if orange juice is the product being tested, one group may base judgments on appearance of the juice, another on packaging, another strictly on brand-names, another only on price, etc. At the conclusion, the rankings of the various groups are compared both with each other and with rankings based on a blind taste test, which is deemed to be the accurate measure of product quality.

Overview of the Experiment

As students arrive, they are divided into small groups. There should be as little visual contact or communication as possible between groups during the experiment. Each group is assigned one or more characteristics of the product as a basis on which to judge quality. They should have no more information about the products than is necessary for them to judge. For example, those judging on the basis of appearance should not only be unaware of which item is which brand, but should, if possible, not even know which brands are being tested.

Each student in the group prepares a quality ranking on the basis of the characteristic being evaluated by his or her group. The members of the group then discuss the product and prepare a consensus ranking. After all groups have completed their rankings, the identity of each product is revealed to the groups and the rankings developed by each criterion are posted on the blackboard.

---

[2]This implementation of this experiment in classes at two schools is described in Netusil and Haupert (1994).

Because this experiment tends to be relatively brief and the resulting analysis is not theoretically complex, it is often possible to discuss and evaluate the results of the experiment on the spot, rather than asking students to prepare a report.[3] There is usually considerable descrepancy between the rankings of the various groups. It is interesting to compare the rankings on actual taste, the relative prices of the products, and the rankings on the basis of other characteristics that are observable in the grocery store. It is sometimes possible to point out situations where a well-advertised brand-name, a colorful package, or an attractive-looking product allow a brand to sell for a higher price than would be justified based on the brand's actual quality as measured by the taste test.

Details of the Design

The first element of the experiment that must be chosen in the product to be examined. Orange juice, chocolate-chip cookies, and pumpkin pies have all been used successfully in classroom applications of this experiment. No doubt other products such as cola drinks or cheddar cheese could also be used. One should avoid products where the brand can be identified from the product itself, such as chocolate bars that have a name or logo embossed in them. One should also seek products with universal appeal—anchovies probably would not work too well. Ease of preparation is also a consideration. Since the instructor must prepare taste samples and samples of packaging, it is wise to choose something that is easy to work with and that can be stored briefly and served at room temperature. It can be interesting to choose a product that offers a choice between a small, local producer and a large, national brand. It is also interesting to compare ready-to-serve versions of an item with frozen, dehydrated, or concentrated forms, for example, including both never-concentrated and reconstituted frozen orange juice within the sample.

The next choice is the selection of information sets for the groups. At least one group must do a blind taste test.[4] One group can rank solely based on price, which is usually an uninformed ordering with most expensive assumed to be best. Other groups can use packaging, appearance, brand names, and (if available) advertising copy. Netusil and Haupert (1994) also have a group that ranks based on "word of mouth." This group bases its rankings on descriptions of each product given to them by a group doing a

[3]Netusil and Haupert (1994) claim that the experiment is typically completed in twenty to thirty minutes.

[4]Alternatively, everyone can do the taste test after ranking the products on other criteria, being careful not to allow the identification of taste samples with any other product characteristics that they used.

blind tasting. There is also a group in their experiment that gets all information about packaging, appearance, advertising, price, and word of mouth, but is not allowed to taste. This group is supposed to represent a well-informed first-time buyer. Another possible source of information is product rankings by publications such as *Consumer Reports*. If such rankings are available, they may be used by one group in preparing their rankings.

Another design consideration is how to deal with disagreements among members within a group. All of the rankings are highly subjective, so there is considerable room for dissent. While it is usual to solicit only a single group ranking, it may at times be interesting to see how much variation is present in the rankings of individuals within the groups. It is very easy to ask each individual to submit a ranking first, before any group discussion occurs, then to ask the group for a consensus ranking. This would allow one to examine the hypothesis that those criteria on which rankings are more unanimous are more highly correlated with quality or with price than those about which there was little agreement. The individual data are easily collected and may add information.

What Students Should Know

Students need no prior experience in order to participate in and understand the results of this experiment. In fact, they make decisions nearly every day that parallel those of the experiment. The experiment allows a systematic evaluation of various criteria for making consumer decisions about product quality.

Instructor Preparation and Materials

The most obvious instructor preparation is obtaining the good to be sampled and preparing taste samples and the exhibits for each group. Exhibits must be packaged in such a way that only the groups who are supposed to see the exhibits can do so. The room should also be set up in a way that allows maximum seclusion of groups from one another. In particular, the groups that are not allowed to use packaging information should be unable to see the packages that are being examined by another group. Although it may take more time to do the experiment this way, groups may be taken to another site one by one to perform their tests. The sites can be reset between groups to provide only the desired information.

Students and/or groups need slips of paper on which to report their rankings. Oral reporting to the instructor may be used, but public oral reporting is not advised, since later groups may possibly change their ranking based on the rankings of groups reporting earlier.

Running the Experiment

After the students arrive and have been separated into the desired number of groups, the instructor should provide instructions for the experiment. The instructions can be extremely simple: "You are going to be asked to rank the quality of various brands of a good based on information about one or more characteristics of the good." Students do not need to be told what information other groups will be getting or about how their rankings will be used.

Care must be taken to assure that groups do not share information. This can be done through physical or temporal separation of sampling activities, and may be reinforced by changing the identification of samples from group to group. For example, the tasting group may refer to the samples as 1 through 4, while the packaging group ranks them as A through D and the appearance group rates them as W through Z, where A, 1, and W are not necessarily the same product.

In cases where the members of a group are having difficulty reaching a consensus ranking, the instructor may suggest the averaging of individual rankings or some other means of breaking the deadlock. The separate reporting of individual rankings may make students more comfortable with accepting a group ranking that does not agree with their own.

If individual rankings are to be obtained, these should be submitted before any group discussion of the good commences. This will prevent group reactions from interfering with the independence of individual responses. Once all groups have reported their responses, the results can be written on the blackboard for general discussion.

The Report Assignment

As noted above, this experiment is less well-suited to a report assignment than some others. If a report is assigned, there are several issues that might be addressed. Among the most interesting issues are the correlation of the various rankings with actual quality as measured by the blind taste test and the correlation of rankings and actual quality with price.

Under perfect information, higher-quality goods should command higher prices. To the extent that quality and price rankings differ (and they often differ considerably), this may represent the effects of imperfect information.[5] Students may be asked to

[5]Recalling that quality perceptions are subjective and vary considerably across individuals, it may be expected that a ranking by any given group of people could vary considerably from the consensus ranking of the whole population, and therefore could differ from the price ranking even if prices accurately reflect the quality perceptions of the broader market. Since all of the subjects in the experiment are students at the same college, it is unlikely that they represent a broad sample of tastes in the overall market. Another possible source of discrepancy between student and market rankings is variation in the quality of the good itself. Brands of orange juice may have variations from year to year, region to region, or package to package in the quality of the juice. If the instructor happened to buy a

explain deviations between the quality and price rankings in terms of nontaste characteristics of the goods, including those being evaluated by groups in the experiment. For example, if product A is more expensive than its quality seems to warrant, does it have a recognizable brand name or attractive packaging that consumers (students in the experiment) take as a (false) signal of quality?

Students may also be asked to comment on the connection between the experiment and naturally occurring markets. What other kinds of goods have similar informational characteristics? To what extent are the judgments made by groups in the experiment similar to those made by actual consumers? How can the evidence of the experiment be used to make one a more rational consumer?

If the results show that acquiring more information leads to a better assessment of quality, students can be asked about the optimal acquisition of information. If information is costless, then they will want as much as possible. But a little reflection can make students realize that acquiring information usually entails explicit or implicit costs. This can lead them to the crucially important idea of acquiring information only as long as the marginal benefit of an additional piece of information exceeds the marginal cost.

Reading and Discussing Student Reports

Discussion of the experiment, whether a report was written or not, can focus on the issues outlined in the previous section. Emphasis can be placed on how information about product quality is obtained in natural markets. They can be asked to discuss the role of government or independent private certification, such as the Underwriters' Laboratory sticker on electrical devices or the Department of Agriculture rating of beef.

Information of this kind tends to be a public good—once the information has been collected, there is zero marginal cost to disseminating it and it is sometimes difficult to keep others from getting it whether or not they have paid. Many people use a library's copy of *Consumer Reports*. Students can be asked to consider under what conditions it would pay an individual to enter the quality-certification business and whether there are cases in which certification would be socially desirable but not privately profitable.

Warranties and guarantees are also methods that the market uses to communicate quality. Students may be asked how the availability of a warranty affects their perception of quality.

---

low-quality bottle of juice, students may rank it lower than the average ranking of that brand.

Variations

One variation on this experiment is to introduce explicitly a market in information. In order to motivate rational decisions, an objective of accurate quality measurement must be established. One way of doing this would be to provide (real or imaginary) cash awards for choosing the best product, say, $5 for picking the best, $3 for the second best, $2 for the third best, and only $1 for the worst. The quality of the product would be judged as before, by a blind tasting.

A market for information of various kinds must be established. This could be done in either of two ways. The simpler way is to establish a pre-set price for each possible source of information about the product (excluding tasting). Students could purchase a peek at the products' packages for, say, 75¢, the brand names for 25¢, or a look at the product itself for $1.00. Students could then move from station to station around the room paying for information as they choose, then make their quality choice when they have as much information as they desire. Perceptive students will recognize that they can save money by pooling information: "You go look at the packages and I'll look at the product, then we'll share information." This behavior allows the subsequent discussion to lead naturally into the public-good nature that information often has.

An alternative way of organizing the information market is to give certain individuals the right to sell their own rankings of the good based on one criterion. Students would be separated into two groups: information providers and information buyers. As an information provider, Jane might be assigned to establish and sell rankings based on packaging. More than one individual could sell rankings based on a single characteristic.

The students who are information buyers would decide what rankings to buy from whom, but would not be allowed to perform actual rankings themselves based on product characteristics. The price of any individual's information would be negotiated between the buyer and seller. As in the previous case, buyers may recognize the gains from pooling information, which may undermine the sellers' market.

In discussion of this variation, students should be able to focus clearly on the cost and benefit of information. They should also recognize the public-good aspect of information in the experiment. This can lead to excellent discussions of what kinds of information are effectively sold in the real world and what kinds are essentially nonproprietary. It can also lead naturally into how the government should be involved in providing or certifying product information.

8 A Search-by-Sampling Experiment

The experiment described in Chapter 7 deals with a particular kind of information problem: obtaining accurate information about multidimensional quality characteristics of a product. The present chapter describes an experiment to elucidate a different kind of information problem, namely, obtaining information on the price of a homogeneous good or service. Modern theories of price dispersion and unemployment are examples of theories based on repeated sampling by individuals of prices bid or offered by various counterparties in an attempt to obtain the most advantageous possible price. This experiment simulates this process, placing students in the role of a searcher who (usually) knows the overall distribution of prices, but must decide between choosing the best price so-far obtained and continuing to sample additional counterparties at an additional cost.[6]

Overview of the Experiment

For simplicity, the bulk of this chapter will describe the experiment in the context of a wage search by a supplier of one unit of labor. Students are searching to sell at the highest possible wage. They are initially given information about the distribution of the wage offers they are to encounter. The information is usually a complete specification of the wage distribution—uniform or triangular distributions are the simplest for students to understand—but more complicated incomplete information setups are discussed below as possible variations.

The main part of each period of the experiment consists of a sequence of wage offers for each student drawn from the given distribution. After each offer, each student decides whether to accept the offer or to continue searching. Each additional search (each wage offer received) imposes a marginal search cost (constant or increasing) on the student, which is deducted from the wage that is ultimately accepted to determine earnings for the experiment.

Most students accept an offer within the first four or five, so each period typically concludes rather quickly, allowing several periods to be run with varying offer distributions and search-cost specifications during one experimental session.

As discussed by Schotter and Braunstein (1981), the optimal strategy in this experiment is a reservation-wage strategy, with the level of the reservation wage depending on the wage distribution and the marginal cost of search. Langan's (1995) results suggest that most students in his classes eventually followed something

[6]The research roots of this experiment are in Schotter and Braunstein (1981). This classroom application is based on Langan (1995).

approximating a reservation-wage rule. Whether they ended up with the optimal strategy or not, students who have experienced the search process are likely to understand more easily the idea of a reservation wage. This puts them well on the way to understanding the nature of costly search and the balancing of marginal benefits and marginal costs of search. Performing a sequence of periods in which search costs and wage distributions are varied allow students to experience the effects of these changes on marginal benefits and costs, and on the optimal reservation wage.

A typical experimental session might begin with a uniform distribution from 100 to 200 and a search cost of 5. After running two or three periods of this setup to get the students thoroughly familiar with the process, the search cost could be raised to 10 for a couple of periods. Finally, the distribution could be changed (with cost of either 5 or 10) to a triangular distribution with range 100 to 200. The triangular experiment should probably run at least three periods since it may take a while for students to learn about the nature of the new wage distribution.

Details of the Design

One important initial design specification to be determined is the method by which wage offers are to be communicated to participants. Because each participant interacts only with the instructor, receiving wage offers and deciding whether or not to accept, this experiment can quite easily be implemented on individual personal computers—no networking is required.[7]

Communication Method. The communication method used by Langan (1995) is a table of random wage values, generated by the instructor prior to the class and printed on an overhead transparency. Each student in the class is randomly assigned to one column of the table—each student in a small class may have his or her own column; in larger classes, several students may have the same column (but they should not know one another's identity). At the beginning of the experiment, all rows of the table are carefully covered with paper so that they cannot be seen. The first wage offers are made by moving the paper to reveal the first row of the table. A student may then raise his or her hand to accept the wage offer shown in the assigned column. Those choosing to reject the first offer then receive a second offer as the instructor reveals the second row of the table, and so on until all students have accepted a wage offer.

The overhead-display method is simple and effective. Each different period of the experiment requires a separate table of random offers corresponding to the appropriate wage distribution, printed on a transparency. One disadvantage of this method is that all

[7]One advantage of a computer implementation is that students could participate on their own, outside of class, saving class time that would otherwise be used for the experiment. A method would then need to be devised for them to report their results either in written or electronic form.

wage offers are public. If a student sees a particularly high offer in another column of the table in the early rounds, it may cause her to "hold out" for a higher offer than she otherwise would expect. Another disadvantage is that transparencies cannot be reused for the same group of students and cannot be used if students accidentally are allowed to see the values while the instructor is setting up.

An alternative method that reinforces the idea of randomness in the students' minds is to draw from a prepared set of small cards or slips of paper with numbers reflecting the appropriate distribution. For example, if the wage distribution is uniform from 100 to 200, the instructor could prepare a set of cards with one (or two or three) cards having each number in the range. A triangular distribution would require a larger deck, since there would have to be one card with 100 and 200, two cards each with 101 and 199, three with 102 and 198, and so on up to fifty-one cards with 150. For each experimental period, the appropriate set of cards is placed in an opaque urn. Then students can file past the urn, drawing a card, noting the value on their record sheets, choosing whether to accept or reject, reporting that on the record sheet, and replacing the card.[8] The advantages of this method, in additional to the appearance of randomness, are that each student gets a unique value that is unknown to other students and that the experiment can be repeated at will without the need for additional overhead transparencies. The biggest disadvantage is the preparation of a large number of cards, but this can easily be done on a computer, by a teaching assistant, or by the students themselves. However, in a variation of the basic experiment where the distribution itself changes frequently, the management of randomly-drawn cards is probably impossible.

Another alternative that avoids elaborate instructor preparation is to have a set of random numbers (or a computer program to generate them) at hand and for the instructor to write the next sequence of random numbers down on slips of paper and hand them out to students. This is probably the best method for complex variations such as those in which inflation changes the distribution each period by random amounts that are unknown to students.

Units of Measure. The objective that is induced by the experiment in student-participants is maximization of experiment profit. Profit for each round is calculated as the accepted wage offer less the total search cost incurred in the period. For this difference to make intuitive sense, both must be expressed in the same units. Since students are likely to think of the wage in dollars (or other currency) per hour, while the

[8]Instructors who wish to maintain tighter control over the data generated by the experiment can maintain the record sheets themselves. The student shows the card to the instructor and reports her decision; the instructor writes it all down.

search cost will probably strike them as a one-time event (e.g., dollars per career), some care must be taken to reconcile the units of measure. The easiest way of doing this is to explain in the instructions that the job for which they are searching is for one year of work, and the wage offers are expressed in terms of earnings per year. That puts the wage numbers and the search-cost numbers on the same basis, giving a good intuitive basis for the difference as net profit.

Wage Distributions. Unless students can be assumed to have some background in probability theory, it is best to use very simple, finite distributions. The uniform distribution is the easiest for students to understand, so it should probably be used for most experiments. The triangular distribution is also quite easy to explain to students. By choosing a triangular distribution with the same mean as the uniform distribution used earlier, it is possible to test the effects on search of a mean-preserving change (reduction) in variance.

Search Costs. The simplest specification of search costs is for each wage offer sampled (perhaps excluding the first) to cost a fixed amount. Introducing rising search costs (i.e., the second offer costs more than the first, the third more than the second, etc.) makes the reservation wage change from offer to offer, which makes the theoretical problem being simulated much more complex.

The amount of the search cost should be relatively small. Too high a cost may lead some students to lose money if they get a few initial bad offers. For example, if each search costs 50, with wage distributions between 100 and 200, then three low wage offers can put a student in the position of losing money with certainty, inducing him to relax his reservation wage. A search cost of 5 will keep induce (presumeably) higher reservation wages and longer searches, but will make it less likely that students will truncate their searches to avoid losing money.

Recall. In search theory, the concept of "recall" refers to whether or not searchers can return to accept an earlier offer that was rejected at the time it was initially offered. A searcher following an optimal, reservation-wage strategy in a constant-cost experiment would not ever recall a previous offer: if it was above the reservation wage she would have accepted it immediately, if it was below, she would not accept it now.

Although it is not optimal behavior, students sometimes do recall earlier offers when this is allowed. Recall behavior is a relatively sophisticated phenomenon for introductory students. Unless the individual instructor wishes to focus extensively on search theory, it is probably best to either allow recall in all experiments or to exclude recall entirely.

What Students Should Know

This experiment is probably better to run *before* students have studied search theory. One of the pedagogical benefits of the experiment is giving students some "hands-on" experience with search and seeing if they come to a reservation-wage strategy naturally. Students who have already studied search theory and are familiar with the idea of a reservation wage are likely to use that strategy because they remember it from the textbook, not because it was a natural response to the situation they encountered in the experiment. If they come to the ideas of marginal costs and benefits of search and reservation wage naturally in the course of the experiment, they are likely to find search theory more believable.

Instructor Preparation and Materials

The materials to be prepared by the instructor depend on the method of communication chosen for the experiment. If the overhead-display method is used, then the instructor must prepare on transparencies sufficient tables of random draws from all distributions to be used. These should be carefully kept from view prior to the experiment.

If urns of cards are to be drawn, then sets cards must be prepared for each distribution to be used. This can be done fairly quickly be handing out blank cards to students and telling each one what cards to prepare. For example, to prepare a uniform distribution from 100 to 200 in a class of 20, ask the first student to prepare two cards each with numbers from 100 to 105, the second student two cards each from 106 to 110, etc. Triangular distributions are somewhat more complex. Each student can be given a number between 100 and 150 and asked to prepare a card for each number between 100 and his number, inclusive. Each can be given additional numbers until all numbers between 100 and 150 have been done once. Then each student is given a number between 151 and 200 and asked to make a card for each number between his number and 200, inclusive. Additional numbers can be assigned until all numbers have been done twice. (This distribution will require a set of about 2500 cards! It might be desirable to limit the range of wage distributions or only to use even-numbered wage values to reduce the amount of paperwork required.)

If individual random numbers are to be handed out by the instructor, then several preparations are necessary. First, there must be a method of generating random numbers. This could be (1) a pregenerated table of random numbers from each possible distribution, (2) a set of uniform random numbers with a method for converting them (via a table of the inverse cumulative distribution function) to the desired distribution, or (3) a computer programmed to generate random numbers on site from the desired distributions.

Running the Experiment

As noted above, a typical experimental session consists of about three experiments with differing distributions and/or search costs. Each experiment should have its own set of instructions that are carefully written to inform the students about changes in the parameters in a way that does not suggest how this should affect their behavior.

Instructions at the beginning should explain carefully two aspects of the search process that students will perform: how the wage offers are generated and how students' profits are calculated. For the uniform distribution, students can be told that wage offers are whole numbers between 100 and 200, with each number being equally likely to occur. For the triangular distribution, the instructor can give an example using the numbers from 1 to 5: it is as if the numbers are drawn from an urn with one ball each numbered 1 and 5, two balls each numbered 2 and 4, and three balls numbered 3. It may be desirable to point out in layman's terms that this makes outcomes near the mean of the distribution more likely than outcomes near the tails.

Profits are the difference between the wage offer accepted and the search cost of the wage offers sampled. Instructions should be clear enough that all students understand that they gain from getting a higher wage offer, but that given the wage offer they lose from additional searching.

Once the instructions are understood, the experiment can begin. The instructor must assure that adequate records are kept of students' offers and decisions. Students may keep their own records and turn them in to the instructor at the end of the experiment, or the instructor may keep records for all students. Since regret is common in this experiment (for example, a student wishing after five wage offers that he had accepted the first and not searched further), some care must be exercised to assure that student-maintained records are accurate, especially if a meaningful reward is attached to profits students earn in the experiment. A table for each student such as Table 8-1 can be maintained by the instructor or the student. Table 8-1 shows a record sheet for a student who received three wage offers of 153, 132, and 178, accepting the third. The top shows the participant and period numbers, the middle section records the offers and the participant's decisions, and the bottom part shows the calculation of profit for the period.

After every student's search is completed, it is time to begin the next period. Each experimental setup (specification of wage distribution, search cost, and recall option) should be repeated at least once so that each student gets more than one set of draws from the relevant distribution. Once the given experimental setup has been repeated as for as many periods as desired, students should be advised of the changes in the parameters for the next experiment. If students are keeping their own records, then the completed record sheets from the first experiment should be collected and new ones passed out. The next experiment can then be performed.

The Report Assignment

Record Sheet for Participant 14 **Period Number** 2

	1	2	3	4	5	6	7	8	9
Wage Offer	153	132	178						
Accept?			X						

Search Cost Calculation: Search Cost = 5 H Number of Offers = 15

Profit Calculation: Offer Accepted = 178
Less: Search Cost = 15
Equals: Profit = 153

Student reports on this experiment may focus on their individual search behavior (and perhaps their analyses of the behavior of others). In particular, students can be asked about their perceptions of the benefits and costs of continuing the search one more period and how (and whether) the benefits and costs changed with the arrival of each offer. They can also be asked whether they followed something similar to a reservation-wage strategy, how they chose their reservation wage, and how it varied with changes in the wage distribution and in search costs.

Since students are not likely to be theoretically equipped to derive the optimal search strategy, the report cannot usually link textbook theory with behavior. However, if the textbook gives examples of search behavior and reservation wages, they can be asked to connect their behavior in the experiment to these examples.

Another fertile area that can be explored in reports is to ask students to identify applications of the model other than wages. They should be able to recognize the symmetry of the experiment with buyer-search models in which, say, a buyer of a stereo system shops for the lowest price by sampling various sellers. Either in the context of wage search or price search, they can be asked about the nature of search costs, which are often implicit time costs. With a little prompting, students should be able to use the results of the experiment to argue that changes in institutions that lower search costs (advertising, job-placement centers, telephone price information from sellers, etc.) will lead to more searches, more complete price information, and higher efficiency.

Reading and Discussing Student Reports

Since this is reasonably difficult material that does not correspond directly to simple textbook theories, some students may be a bit lost in writing their reports. However, this can be an advantage: the absence of an obvious interpretation based on the textbook may encourage more creative responses on the part of the students.

Most or all of the points suggested above as possible focal points for the report can be explored in some detail in the discussion. Costs and benefits of search, reservation wages (prices), and comparative-static effects can all be discussed in the context of this experiment.

A good place to start the discussion is a comparison among students of the strategies that they used in conducting their searches, which may reveal interesting differences. After allowing students to explain their strategies and argue about which is best, the stage is set for the instructor to ask, "How would one go about determining which search stratregy is best?" Answering this question gives them the basic analytical framework of search theory.

The induced experimental objective of maximum profit corresponds to the assumed economic objective of the searcher: maximize wage earnings less search costs. This may be used to motivate this quantity as an appropriate maximand. Following the usual arguments of marginal analysis (with which students are probably, by now, familiar), marginal benefit and marginal cost of search can be introduced, with the idea that one should continue searching as long as MB > MC.[9]

In the basic experiment, it is easy for students to see that marginal cost is constant. Marginal benefit is more complicated and probably cannot be approached by direct analysis unless students are familiar with basic probability theory. However, students lacking the necessary math may still be able to understand the underlying economics. In intuitive terms, the marginal benefit of an additional search is the product of (1) the probability of receiving an offer higher than the current offer (or best offer to date if recall is allowed) and (2) the amount by which an improved offer (on average) exceeds the current offer.

In terms of a concrete example, if the distribution is uniform between 100 and 200 and the current offer is 150, then the probability of improving one's offer on the next turn is 0.5 and the (probability-weighted) average of the offers better than 150 is 175, which is an improvement of 25. Thus, the marginal (expected) benefit of search is 0.5 H 25 = 12.5. A searcher with an offer of 150 in hand should continue searching if the cost of an additional search is less than 12.5.

It should be easy for students to see that a higher current offer reduces both the probability of receiving a higher offer *and* the amount by which any higher offer exceeds the current offer. Thus, both factors going into the marginal benefit depend negatively on the current offer. To continue the numerical example, if the current offer is 160, then the probability of receiving a better offer is 0.4 and the average of offers better than 160

[9]Strictly speaking, we should talk of "expected marginal benefit," since MB is a random variable. Each instructor must decide how carefully to make this distinction based on the sophistication of his or her students.

is 180, for an improvement of 20. Thus, for a searcher with an offer of 160 in hand, the marginal benefit of search is 0.4 H 20 = 8 < 12.5. Perceptive students will see that for the uniform distribution, the marginal benefit curve is a parabola, depending on the square of the current offer.

Drawing the MB curve as a declining function of the current offer and the MC curve as a constant shows that there is a unique level of the current offer at which they cross. This is the reservation wage: offers at or above it should be accepted; offers below should be rejected.

The MB/MC apparatus explained above can also be used to motivate the comparative statics of the search model. Changing search cost moves the horizontal MC curve up or down, with the obvious effects on the reservation wage and duration of search. A change to a more concetrated wage distribution such as the traingular distribution lessens the probability of a higher offer for any current offer above the mean of the distribution, and also reduces the expected level of a higher offer, both of which tend to reduce the marginal benefit of search. An extreme example may illustrate this effect to students lacking an intuitive understanding of probability. Suppose that the probability distribution was such that 98% of all wage offers were of 150, with a 1% chance for each of 125 and 175. A searcher with an offer of 150 in hand is very unlikely to improve on that offer (only 1% of the time), and if she does improve it the improvement will be only 25, so the marginal benefit of search is very low (1% H 25 = 0.25).

Variations

One major variation on this experiment extends its use into the macroeconomics curriculum. The neoclassical theory of the Phillips curve developed by Friedman (1967) and Phelps (1967) relates the unemployment rate to the difference between the average wage and people's perception of the average wage. It is possible to simulate this model if searchers have imperfect information about the wage distribution.[10]

Searchers in the neoclassical model are assumed to know the distribution of real wages in the economy and the distribution of prices. The wage offers they receive are nominal wage offers that are the product of a drawing from the known real-wage distribution and a drawing from the known price-level distribution. A high nominal wage could be the result of high prices, high real wages, or both. Searchers must solve a "signal-extraction problem" to attempt to distinguish these and estimate the offered real wage.

This framework can be implemented quite easily in the present experimental setup. For example, the price level (which can be motivated as the cost of basic necessities) can be "high," "low," or "average," represented by 110, 90, and 100, respectively. Students are correctly informed that in each period the probabilities associated with these states

[10]One existing application of a related model in the experimental literature is ?????.

are 0.3, 0.3, and 0.4. A high (low) price level raises (lowers) the wage distribution by 10 relative to the U(100, 200) distribution that prevails when the price level is normal.[11] Each student's profit for the period is the accepted (nominal) wage offer minus search costs, less the price level—the wage net of search costs less the cost of basic necessities.

Before each period, the instructor randomly determines, but does not reveal, whether the price level is low, normal, or high. Nominal wage offers are then generated as in the basic experiment, adding or subtracting 10 if the price level is high or low. As soon as all students have accepted offers, the price level is revealed and students calculate their profit.

The result to be expected, based on the neoclassical model, is that in periods when prices are higher than expected, searches will be shortened as participants mistake high nominal wage offers for high real wages. The macroeconomic implication of shorter searches is lower unemployment. Similarly, lower-than-expected prices should lead to longer searches (higher unemployment).

Instructors who wish to present the neoclassical model of the Phillips curve and the principle of inflation neutrality may find this variation of the experiment a useful additional to the macroeconomics course.

[11]This specification based on adding and subtractive prices and wages corresponds formally to a uniform distribution of log-wages and log-prices. However, it is simpler to use the additive specification in this instuctional setting to avoid the necessity of students dividing by 110 or 90 to calculate profits.

9 Voluntary-Contribution Experiment

The voluntary contribution mechanism relates to free riding in the provision of public goods. This subject lends itself quite naturally to experiments and has been one of the most popular in the research literature of experimental economics.[1] Subjects in the experiment are confronted with the opportunity to spend on a public good, which yields a high payoff for society but a low payoff for the individual, or a private good, which yields a return that is higher to the individual but lower for society.

This experiment is a useful demonstration of the conflict associated with voluntary provision of public goods. Moreover, the results of the experiment seem to be quite sensitive to the degree of communication allowed and to the general relationship among the participants. This gives one an unusual opportunity to examine characteristics of the specific environment in which the experiment is conducted and the influence that they may have on the results.

Overview of the Experiment

All of the participants in the experiment have the similar role of "citizens." The experiment can, in principle, be designed for large or small groups. At Reed, classes have been divided into medium-sized groups of about 10 and small groups of 4 or 5. One interesting question to discuss after the experiment is whether and why the results of the experiment may have differed with the size of the group.

In each period, individuals within each group have the opportunity to invest a set of "tokens." (The number of tokens is arbitrary, but should be large enough to allow a reasonable continuum of outcomes, say, 10 or 20.) Each citizen chooses how many of his or her tokens for the period to spend on a public account or good; the remainder are spent on a private good. Tokens spent on the public good pay an equal amount to all members of the group; tokens spent on the private good pay only to the individual investor. The payoffs are constructed so that the total (social) payoff to all participants is higher for the public good, but the individual's (private) share of the payoff from the public good is lower than the private payoff from the private good. The citizens are told

---

[1]For an example from the sizable literature, a recent study pertaining to information issues discussed later is R. Mark Isaac and James M. Walker, "Communication and Free-Riding Behavior: The Voluntary Contribution Mechanism," *Economic Inquiry* 26:4, October 1988, 585-608. A general introduction to this literature is in chapter 6 of Douglas D. Davis and Charles A. Holt, *Experimental Economics*, (Princeton NJ: Princeton University Press, 1993).

at the beginning of the experiment what the payoffs are to the individual and to the group for each token spent on either good.

Citizens can communicate their decisions to the experiment monitor in several ways. Slips of paper are easy to arrange; networked computers are highly efficient if the necessary programming has been accomplished. Once the decisions for a period have been collected, the total contributions by all members of the group to private and public goods are posted and the next period begins.

The results of this experiment seem to be sensitive to many subjective and design factors. For example, participants seem more likely to contribute to the public good if they can effectively demonstrate to their peers (i.e., by displaying their decision slips) that they are doing so. Groups, such as those in classes of Reed students, who know each other well outside the classroom are probably also more likely to be "public spirited" than strangers. To the extent that one can control these factors in the laboratory, interesting comparisons can be made between the behavior of groups.

Because this experiment proceeds relatively quickly, it is possible to combine the experiment itself and discussion in the same 50-minute lab period. Thus, no report assignment is normally given for the voluntary contribution experiment.

Details of the Design

Specification of the design of the voluntary contribution experiment involves fixing group size and the payoff structure.[2] One must also think carefully about what, if any, restrictions are to be placed on participants' ability to communicate with one another or to reveal their contributions.

The payoff structure must satisfy the following inequalities: *SocPub* > *Priv* > *IndPub*, where *SocPub* is the social (group) payoff for each token spent on the public good, *Priv* is the payoff to the individual (only) for each token spent on the private good, and *IndPub* is the individual's share of the payoff for each token spent on the public good. According to the rules of the experiment, *IndPub* = *SocPub* / *n*, where *n* is the size of the group. In order to assure that students have the correct incentive structure, each of these inequalities should hold strongly, perhaps by a factor of two or more. However, notice that once *SocPub* and *Priv* are established (which must be done when the instruction sheets for the experiment are prepared), the ratio *Priv* : *IndPub* then depends on the number of participants.

If the total number of participants may vary (due to unpredictable student absences), it can be a little tricky to plan for this experiment. To see this, suppose that the instructor has planned for a group of ten participants, with payoffs set at 30 cents for the public good (3 cents to each individual) and 10 cents for the private good (all to the one individual). This make the ratio of *SocPub* : *Priv* : *IndPub* equal to 30 : 10 : 3, roughly a factor of three for each inequality. If only five students end up participating, the *Priv* : *IndPub* ratio becomes 10 : 6, and there is much weaker incentive to free ride. Obviously,

[2]For an analysis of the effects of group size and payoffs in a research experiment setting, see R. Mark Isaac and James M. Walker, "Group Size Effects in Public Goods Provision: The Voluntary Contributions Mechanism."

if the number of members in the group fell to three or less, there would be no free-rider incentive.

What the Students Should Know

Students typically pick up on the payoff structure very quickly, even if they have not been exposed to the theory of public goods or to the free-rider problem. However, the concept of "public account" or "public good" will make more sense to them if the basic theory has been covered.

Instructor Preparation and Materials

The most obvious preparation is the design of the experiment(s) and the preparation of information sheets and other materials. There are interesting potential differences between the behavior of larger and smaller groups in this experiment, so it is probably desirable to design one experiment for a group of about ten and one for a group of about four. Each experiment can be designed so that the ratio of *SocPub* : *Priv* : *IndPub* is approximately 10 : 3 : 1, as noted above.

Unless the number of students to participate in the experiment is known in advance, the instructor must be prepared to handle groups of varying sizes. A plan for handling all anticipated numbers is necessary preparation. For example, the instructor might set up a basic experiment design for about ten students and one for about four students. If exactly ten students show up, then all can be in one group of ten; if fourteen are present, then one group of ten and one group of four can be organized; if there are twenty students, then the experiment can be run with two groups of ten; and so on. For intermediate numbers of students, the size of the groups can be varied somewhat. Eleven could be accommodated within the ten-person design; twelve could be a stretch of the ten-person experiment or in two groups of eight and four. The instructor should plan ahead to determine the desired groups for each likely number of participants.

The materials that must be prepared include a mechanism for distributing the information about *Priv* and *IndPub* to the participants, a way for them to communicate their choices, and a record log on which the instructor can tabulate results. Since there is no private information in the simple version of the experiment, the payout information can be given orally or written on a blackboard as an alternative to the information sheets that are used in other experiments. Care must be taken in deciding how the students' investment decisions are to be communicated to the instructor. As noted above, there may be significant differences in behavior depending on whether students can (or must) reveal their contributions to one another. Having the students announce their contributions is likely to lead to a higher public-good contribution rate than if they write them on a slip of paper, especially if they are not allowed to show the paper to anyone else. In large classes, it is a useful variation to use different communications schemes for different groups and to test whether groups in which revelation is automatic or easy contribute more than those in which secrecy is enforced.

Running the Experiment

The initial task is dividing the students into groups of appropriate sizes according to the prearranged plan. This may be done randomly in order to break up friends who tend to sit together, or friends may be allowed to collect into groups. (Will more be contributed among friends than among less familiar groups?) Depending on the amount of communication to be allowed among group members, the members of each group may be set up together in one section of the room (the normal case) or dispersed around the room to make identification and communication among group members difficult.

As always, instructions must be given slowly and carefully with ample opportunity for students to ask questions. An example of individual and group behavior and the resulting outcome/payoff may clarify the nature of the experiment.

Once the experiment gets going, the instructor (and any assistants) can move among the groups, assuring that rules about communication are followed and collecting each round of results. The aggregate results of each round should be posted for the group to see, showing the total contribution of tokens to the public good. This will enable each member of the group to calculate his or her payoff from public and private tokens during the round as *Priv* times the number of tokens to the private good plus *IndPub* times the group's total contribution to the public good.

The Report Assignment

The economic theory underlying this experiment is probably clear to anyone who understands it well enough to be an informed participant. At Reed, we do not usually assign a report, opting instead to discuss the experiment immediately after performing it.

If a report is assigned, issues beyond the basic idea of the free-rider problem may be explored. For example, students may be invited to consider how the outcome of the experiment was or would have been affected by such factors as (1) whether participants know the identities of the other members of their group, (2) whether participants can observed one another's pattern of contribution, and (3) whether the participants know each other well and will have an ongoing relationship after the experiment is over. The students can be told about the results of research experiments in which random groups of individuals are isolated at computer screens to perform the experiment and asked to evaluate and explain any differences between the results achieved in class and those reported in the research literature. Finally, the students may be asked to explore differences among actual public goods or externality situations: Is it easier to elicit contributions for public projects at a local level than at a national level? Is it easier if individuals receive some identifiable token in exchange for their contributions (such as a public television bumper sticker)?

An interesting question with important implications that students may explore in a report is what would happen if the entire group were to vote on a contribution rate to the public good that would then be enforced on every member, i.e., a tax rate. Of course, in this experiment it is in everyone's interest to have a 100% tax, with all tokens contributed to the public good. Most of them will see this immediately. However, it raises important issues of tax enforcement, since if it is possible to cheat on taxes without being caught, the situation reverts to a standard voluntary contribution problem.

Reading and Discussing Student Reports

As noted above, Reed students do not typically write reports on this experiment. However, the issues discussed as potential report topics are discussed extensively within the lab session. The experiment can usually be completed in approximately 30 minutes, leaving plenty of time for discussion on the same day.

The subject of free riding is a familiar one that is easily grasped by most students. They may be able to come up with additional examples from their campus or home lives. These can contribute in an important way to a class discussion.

If the students have considered the optimal tax issue suggested above, then some discussion of why the optimal tax rate is 100% may be helpful. In this experiment, the marginal benefit of contributing to public or private goods is constant, no matter how much is contributed. This means that every token yields a higher social return if contributed to the public good. A more realistic setup would have the marginal return to tokens in either the private or public good declining as more tokens are invested. (This is discussed as a variation on the experiment below.) Under this alternative scheme, optimizing individuals would equate the marginal private return to private and public rather than necessarily achieving a corner solution in which all are spent on the private good. The social optimum would equate the marginal social returns of public and private expenditures instead of resulting in a corner solution with full contribution to the public good. In this case, the optimal tax rate would be the one that equated marginal social returns, not necessarily 100%.

Using the Experiment in Later Classes

This experiment is normally done rather late in the course, when the relationships between governments and markets are discussed. Because of this, there is often less opportunity to use the results later in the course. Discussions of public goods and externalities can both benefit from the illustration provided by the experiment. If public finance is a major focus of the course, then some of the issues raised in the reports section above could be used to explore issues of fiscal federalism, tax evasion and avoidance, and cost-benefit analysis.

Variations

Information. As noted above, one way in which this experiment can be varied is to change the degree to which participants know each other's identity and can find out or reveal contribution patterns. The first time this experiment was run at Reed, one class managed to get near-complete contributions to the public good. This seemed to be because one participant began showing his contribution sheet with full contribution to the public good to the others and challenging them to show theirs. Soon everyone was contributing fully and revealing this to the others, so that they could be confident that there were no free riders. The other class got much lower contribution rates, as they made their spending decisions in private and free riders were never revealed to their peers.

This difference could form the basis for a controlled experiment. One group could be provided with a mechanism for sharing contributions (the instructor could even publish them), while another is forbidden to discuss or reveal their contribution pattern. One could go even further be not informing them of the total number contributed to the public good on a round by round basis.

Declining Marginal Returns. With a sufficiently sophisticated pool of participants, one could design an experiment with declining marginal returns to both the private and public good. This could make both the individual and social optima interior solutions. Note, however, that students would have to make calculations about their optimal return. If the marginal return to the public good depends on *total* (as opposed to individual) contributions, then this calculation would be subject to uncertainty, since the marginal return on the public good would depend on the decisions of other participants.

Taxes. At the end of the experiment, it is interesting to propose a universal tax system under which they would "vote" on a tax rate that would apply to all individuals in the group. All tokens taxed would go to the public good and all tokens retained after taxes would go to the private good. In the simple setup, the optimal tax rate is 100%, which most students grasp quickly—even those who were free riders during the experiment.