Extracting Valuable Data from Classroom Trading Pits

Theodore C. Bergstrom and Eugene Kwok

Abstract: How well does competitive theory explain the outcome in experimental markets? The authors examined the results of a large number of classroom trading experiments that used a pit-trading design found in *Experiments with Economic Principles,* an introductory economics textbook by Bergstrom and Miller. They compared experimental outcomes with predictions of competitive-equilibrium theory and with those of a simple profit-splitting theory. Neither theory was entirely successful in explaining the data, although in the first rounds of trading there was significant profit splitting and, as traders became more experienced, outcomes were closer to those predicted by competitive theory.

Key words: classroom experiments, competitive equilibrium, pit trading, profit splitting, supply and demand
JEL codes: A1, A22, C9, D00

Experimental economics began in the 1940s in Edward Chamberlin's Harvard classroom. Chamberlin devised a classroom trading pit that served two purposes—instructing the participating economics students and testing scientific propositions. In both areas, Chamberlin's contribution was radically innovative.

Chamberlin (1948, 95) remarked that it is "a commonplace" that economics cannot resort to the laboratory techniques used in the natural sciences. He proposed to make "a tiny breach" in this position by describing a market experiment that he performed many times in his classroom. He devised his experiments to test a proposition that he had earlier advanced in abstract terms—that market outcomes differ significantly from competitive predictions, when conducted under realistic market conditions where recontracting is possible.

Chamberlin (1948) was also aware of the value of active learning as an educational strategy. He reported that his experiment

> . . . in my own experience has been found stimulating and instructive to students both (a) for actual participation as buyers and sellers in a market mechanism and (b) for the many comparisons afforded, both of similarity and contrast, between the laboratory market and its diverse counterparts in the real economic world. (p. 95)

Chamberlin (1948) induced market demand and supply by distributing cards that assigned each participating student a role either as a supplier or a demander.

Theodore C. Bergstrom *is the Aaron and Cherie Raznick Professor of Economics, University of California at Santa Barbara (e-mail: tedb@econ.ucsb.edu).* ***Eugene Kwok's*** *work on this article was performed while he was an undergraduate student at the University of California at Santa Barbara.*

Each supplier was assigned a seller cost at which he or she could supply a single unit, and each demander was assigned a buyer value for a single unit of the good. In any sale, the seller's profit was the difference between the price and the seller cost, whereas the buyer's profit was the difference between the assigned buyer value and the price. Students were asked to move about the room trying to make the best deal they could with a person of the other type. When a buyer and seller agreed on a price, the transaction was recorded on the blackboard. Trading continued until no more supplier-demander pairs were willing to make trades.

Chamberlin described his classroom experiments in the *Journal of Political Economy* in 1948, but this pathbreaking work received little attention until 1962,[1] when Vernon Smith recognized the merits of Chamberlin's experimental method and followed up with a remarkable series of experiments that ultimately persuaded much of the economics profession that economics can be an experimental science. Smith's early experimental work, like Chamberlin's, was conducted with students in economics classes. Smith's account of this work is found in a charming essay, "Experimental Economics at Purdue" (Smith 1991).

Most experimental economic research is now conducted with paid subjects outside the classroom. One good reason for using paid subjects is that if subjects are paid, one can subject them to repetitive activities that tuition-paying students would find boring and uninstructive. But the results of classroom experiments are a plentiful source of interesting data that researchers should not ignore. An advantage of using classroom experimental data is that the same experiment is often run year after year and at several different universities, generating large samples at low cost.

A SUPPLY AND DEMAND EXPERIMENT

Bergstrom and Miller (2000) in their introductory economics text, *Experiments with Economic Principles*, introduced each of the major ideas of introductory microeconomics through the use of Chamberlin-like classroom trading pits. Their first chapter is based on a demand and supply experiment similar to that of Chamberlin. This experiment has now been conducted in several hundred classrooms. For many classes, the results have been preserved and recorded in convenient form because experimental results are typically reported to the students as spreadsheets posted on the Web. For this study, we collected data on transaction prices and quantities from 31 classrooms at 10 universities.[2]

Each participant in the experiment was assigned a role as a supplier or a demander of apples. A supplier could choose to sell either zero or one unit, and, similarly, a demander could buy either zero or one unit. Each supplier was assigned one of two possible "seller costs" and each demander was assigned one of two possible "buyer values" for a unit of apples. Buyers and sellers were asked to roam around the room and try to make as profitable a deal as possible. When a seller and a buyer agreed on a price, they wrote it on a sales contract, along with their identification numbers and the seller cost and buyer value. The market manager recorded transaction prices on the blackboard as the contracts were turned in. If a seller with seller cost C sold a unit at price P to a buyer with buyer value B, then the seller's profit was $P - C$, and the buyer's profit was $B - P$.

Participants were reminded that they did not have to trade if they could not find a profitable trading opportunity.

This experiment included two sessions with different distributions of buyer values and seller costs. Each session consisted of two rounds of trading. In each session, after the first round of trading was completed and students had observed the results, students were asked to play again with the same buyer values and seller costs as in the first round. The only thing that was new in the second round was the experience that participants had gained from the first round.

Competitive Demand and Supply Curves

The number of persons with each buyer value and seller cost differed between classrooms, depending on the number of students in the class. However, the distribution of types was chosen so that equilibrium prices and the qualitative features of supply and demand were the same in all classes. In every class, there were two types of suppliers, high-cost suppliers with seller cost of $30 and low-cost suppliers with seller cost of $10 per unit. There were also two types of demanders, high-value demanders with buyer values of $40 and low-value demanders with buyer values of $20. In session 1, there were approximately twice as many low-value demanders as high-value demanders and about twice as many low-cost sellers as high-cost sellers. The competitive supply and demand curve and the competitive equilibrium price and quantity for a class of 47 students in session 1 are shown in Figure 1; the competitive equilibrium price was $20, and the competitive equilibrium quantity was 15 units sold.

Session 2 had the same two types of buyers and sellers, but in this session, there were twice as many high-cost suppliers as low cost and twice as many high-value demanders as low value. This time, the competitive equilibrium price was $30. The competitive supply and demand curve and the competitive equilibrium price and quantity for a class of 47 students are shown in Figure 2.

Comparing Results to Theory

Average price: Predictions and outcomes. Participants in the market were told nothing about the distribution of buyer values and seller costs. They knew only their own values and whatever they learned from talking to other participants.[3] Given that individuals knew little about market conditions when they participate in the first round, we did not expect all transactions to take place at the competitive equilibrium price.

The distribution of classroom mean prices in rounds 1 and 2 of session 1 for the 31 classrooms in our study is shown in Figure 3. The competitive equilibrium price for this session was $20. Even in the first round, the average price in most classrooms was fairly close to the competitive equilibrium price. In the second round, as traders learned more about the prices at which others bought and sold, prices typically moved closer to the competitive equilibrium price.

The distribution of mean prices in the first and second round of session 2, where the competitive equilibrium price was $30, is shown in Figure 4. Session 2

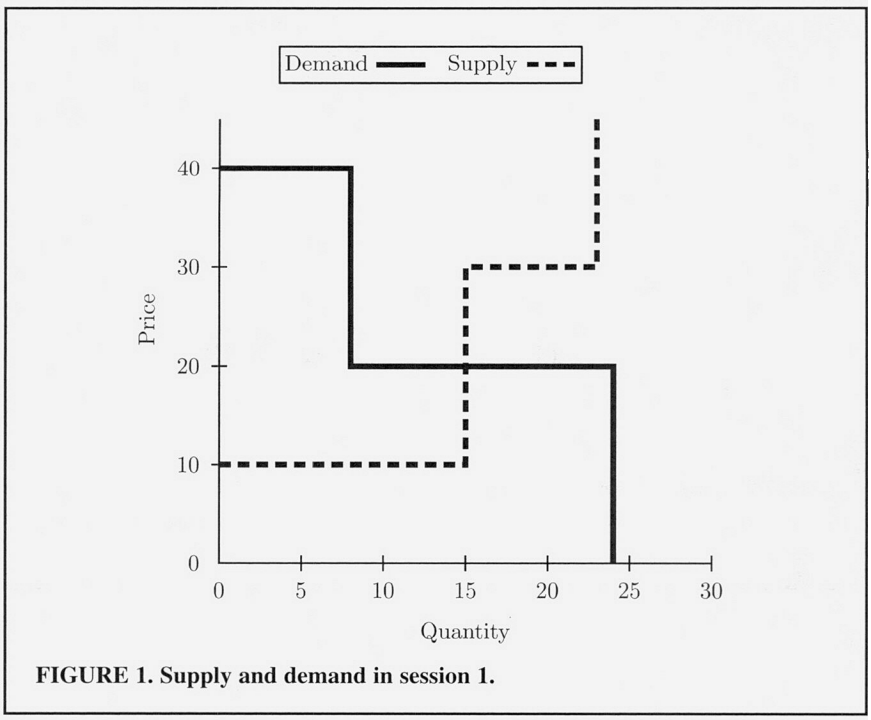

FIGURE 1. Supply and demand in session 1.

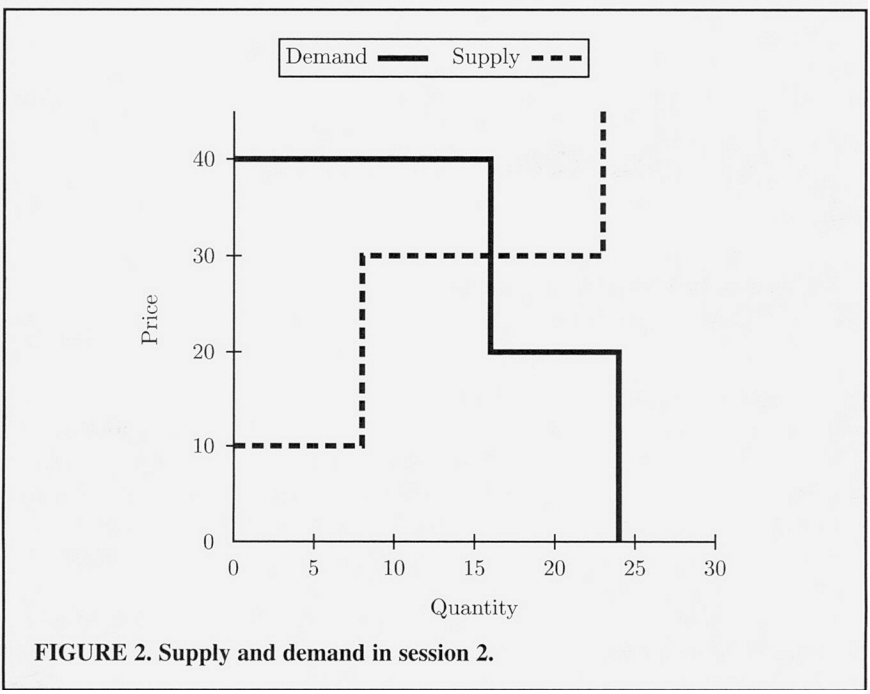

FIGURE 2. Supply and demand in session 2.

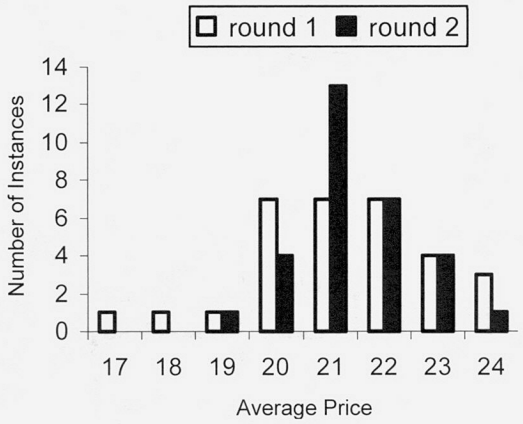

FIGURE 3. Average prices in classrooms: Session 1, rounds 1 and 2.

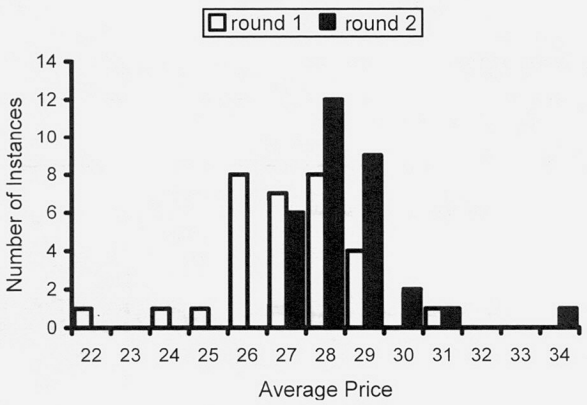

FIGURE 4. Average prices in classrooms: Session 2, rounds 1 and 2.

took place immediately after the close of round 2 of session 1, in which the equilibrium price was $20. Students were not told that the equilibrium price in session 2 would be higher than in session 1 and, as shown in Figure 4, in round 1 of session 2, students seemed to have expected prices to be lower than the $30 equilibrium price. In round 2, having experienced the outcome of the previous round of trading, students seemed to adjust their expectations, and the average price in most classrooms moved closer to the equilibrium price. It is interesting that although the modal price in round 2 of session 1 was $21, which was $1 above the equilibrium price, the modal price in round 2 of session 2 was $28, which was $2 *below* the equilibrium price.

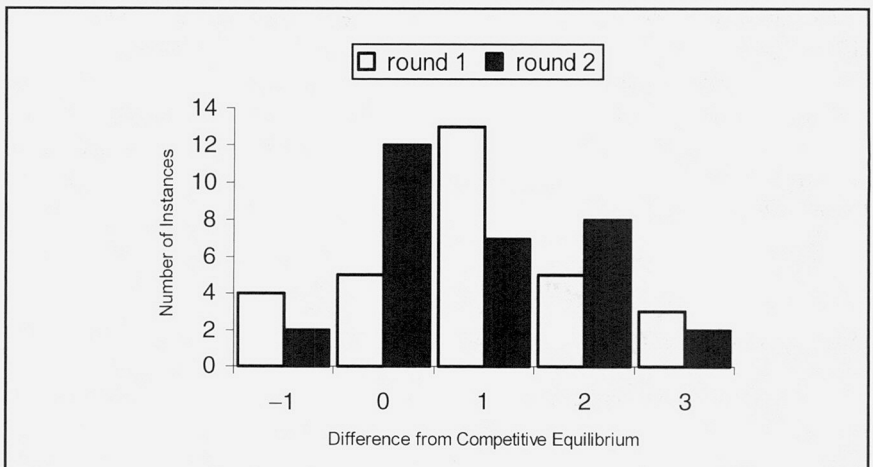

FIGURE 5. Quantity differences in session 1.

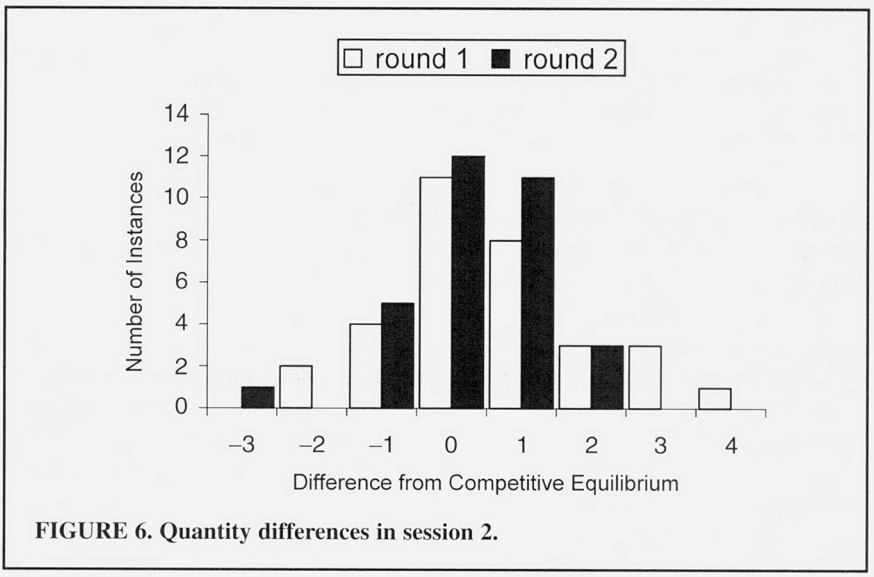

FIGURE 6. Quantity differences in session 2.

The average price in a classroom was usually fairly close to the competitive prediction. In both sessions, in the second round of trading, the mean price was within \$3 of the competitive equilibrium price in 31 of the 32 classes.

Quantity: Predictions and outcomes. In most classrooms, the number of units sold was close to the competitive equilibrium quantity. In the second round of trading, the quantity sold was within one unit of the competitive equilibrium quantity for 21 of the 31 classrooms in session 1 and for 27 of the 31 classrooms in session 2. The distribution across classrooms of the differences between the quantity actually sold and the competitive equilibrium quantity is shown in Figures 5 and 6.

The number of trades exceeded the competitive outcome more often than it fell short. This tendency for excess trade was also remarked by Chamberlin (1948). Chamberlin used a numerical example to make a plausible case that nontatonnement pit trading is likely to lead to too much rather than too little trading. He did not provide a proof of this assertion, nor did he spell out exactly what was to be proved. Bergstrom (2004) stated and proved a result that supported Chamberlin's conjecture.

EARLIER EXPERIMENTS

Chamberlin's Results

In most classroom market experiments today, experimental results are used to instruct students about how *well* competitive theory works to explain market outcomes. It is interesting that Chamberlin (1948), who originated classroom market experiences, emphasized the differences rather than the similarities between experimental results and the predictions of competitive theory. He reported the following results of 46 experiments conducted over the years in his Harvard classroom:

> the actual volume of trade was higher than the equilibrium amount forty-two times and the same four times. It was never lower. The average price was higher than the equilibrium price seven times and lower thirty-nine times. (Chamberlin 1948, 97)

Chamberlin (1948) did not supply further details about his results.[4] Thus, we do not know whether or not the experimental results were usually close to those predicted by competitive theory. We only know that the predictions were rarely exactly right and were biased upward in the case of quantity and possibly downward in the case of price.

Chamberlin (1948) saw no reason to expect that the outcome in his experiment would approximate competitive equilibrium. He pointed out that in his classroom experiments, as in real-world trading, there was no recontracting. Traders do not experience a single equilibrium price but must trade on the basis of their own limited information in encounters with others. Thus, there will be some trading at "false prices." In contrast, the standard accounts of competitive equilibrium posit a *tatonnement* mechanism such that no actual trades occur until an equilibrium price is found.

Chamberlin (1948) explained that

> my own skepticism as to why actual prices should in any literal sense tend toward equilibrium in the course of a market has been increased not so much by the actual data of the experiment before us ... as by failure, upon reflection stimulated by the problem, to find any reason why it should do so. (p. 102)

Chamberlin's (1948) experimental results correspond only to the first round of our experiment because he did not offer a second round of trading. Our experiments also differed from his in that he had many distinct buyer values and seller costs, whereas our experiment had only two possible buyer values and two possible seller costs. Although our classroom experiments showed a tendency toward excess trading in the first round, this tendency was not as strong as that

found by Chamberlin. In the first rounds of our two sessions, the volume of trade exceeded the competitive prediction 36 times, was equal to the competitive prediction 16 times, and was smaller than the competitive prediction 10 times. Unlike Chamberlin, even in the first round, we did not find a systematic tendency for the average observed price to be less than the competitive price. In session 1, the average observed price was usually higher than the competitive price, and in session 2, the observed price was usually lower.

Smith's Results

Vernon Smith (1991) decided to revise Chamberlin's procedures to give the competitive model a better chance. Smith explained that

> The thought occurred to me that the idea of doing an experiment was right, but what was wrong was that if you were going to show that competitive equilibrium was not realizable . . . you should choose an institution of exchange that might be more favorable to yielding competitive equilibrium. Then when such an equilibrium failed to be approached, you would have a powerful result. This led to two ideas: (1) why not use the double oral auction procedure, used on the stock and commodity exchanges? (2) why not conduct the experiment in a sequence of trading "days" in which supply and demand were renewed to yield functions that were daily flows? (p. 155)

Smith's (1962) first published discussion of the results of his classroom experiments appeared in the *Journal of Political Economy,* where he reported that

> The most striking general characteristic . . . is the remarkably strong tendency for exchange prices to approach the predicted equilibrium for these markets. As the exchange process is repeated . . . , the variation in exchange prices tends to decline and to cluster more closely around the equilibrium. (p. 16)

In Smith's (1962) view, real markets normally renew themselves periodically with buyers and sellers bringing new output and renewed needs to the marketplace in each trading day. In this process, traders gain knowledge of market conditions as they move from one day's trading to the next. Smith's (1962) experiments included three to five "trading days," corresponding to the rounds in our experiment.

Smith (1962) reported the results of 10 classroom experiments with differing shapes of supply and demand curves. Some of these experiments had additional differences that made them hard to compare either with Chamberlin's (1948) experiments or our own. Smith's first four experiments differed from Chamberlin's experiments only in the shape of the demand curves, the use of a double-oral auction, and the use of multiple rounds. In each of these experiments, the variance of prices decreased from the first round to the second and again from the second round to the third. Across these four experiments, the average of the variances in the second round was about 55 percent of that in the first round, and the average of the variances in the third round was about 60 percent of that in the second round. Variances in the fourth round were little different from those in the third.[5]

DO CLASSROOM RESULTS SUPPORT COMPETITIVE THEORY?

The Bergstrom and Miller (2000) experimental design followed Smith (1962) in adding a second round of trading. To save classroom time, most instructors do not conduct a third or fourth round, as did Smith. These experiments follow Chamberlin's (1948) open trading-pit design rather than Smith's double-oral auction.

We have seen that the quantities and average prices found in our classroom experiments were reasonably close to the predictions of competitive theory. Thus, it is usually easy to convince credulous undergraduates that competitive theory has impressive predictive power. Does this conclusion withstand closer scrutiny? Might there be a competing theory that works as well or better?

Although competitive theory makes reasonably good predictions of average prices, we show that a simple and plausible alternative theory predicts average prices in our classroom experiments at least as well. A better test of the competitive theory and of competing alternatives requires an examination of the detailed predictions of each theory. It is important to recognize that competitive theory makes many predictions beside those of total quantity and average price. This theory predicts that *all* transactions take place at the same price. It predicts not only the total number of sales but also the numbers of trades that will be made by each type of supplier and demander.

Profit-Splitting and Average-Price Outcomes

A natural candidate for an alternative to competitive equilibrium theory is the profit-splitting hypothesis. At the beginning of trade, suppliers and demanders are paired at random. For any pair, if the demander's buyer value exceeds the supplier's seller cost, then the two of them will agree to a price halfway between the demander's buyer value and the supplier's seller cost. If the demander's buyer value is less than the supplier's seller cost, then no mutually profitable deal can be struck, and they fail to trade. Those who do not trade with their first partner may search for another partner and, if mutually profitable trade is possible, split the profits with this partner.

The profit-splitting theory and the competitive theory make similar predictions about the average prices paid in both sessions. In session 1, approximately 2/3 of the demanders had low values of $20, and 1/3 had high values of $40. About 2/3 of the suppliers had low costs of $10, and 1/3 had high costs of $30. If encounters were random, then on average, 4/9 of the encounters would be between low-value demanders and low-cost sellers, 2/9 of the encounters would be between low-value demanders and high-cost sellers, 2/9 would be between high-value demanders and low-cost sellers, and 1/9 would be between high-value demanders and high-cost sellers. The profit-splitting hypothesis predicts that for those matchings in which the buyer's value exceeds the seller's cost, a sale would take place at a price midway between. The only individuals who did not make a trade with the first person they met were the low-value demanders with $20 buyer values who met high-cost suppliers with $30 seller costs.[6]

In session 2, about 1/3 of the demanders had low values, and 2/3 had high values, whereas 1/3 of the suppliers had low costs, and 2/3 had high costs. As with session 1, we could calculate the fraction of all matchings of each possible

TABLE 1. Outcomes by Pair Type Under Profit Splitting

Buyer value (BV) and seller cost (SC)	Predicted price	Fraction, session 1	Fraction, session 2
BV: Low: $20 SC: Low: $10	$15	4/9	1/9
BV: Low: $20 SC: High: $30	No trade	2/9	2/9
BV: High: $40 SC: Low: $10	$25	2/9	2/9
BV: High: $40 SC: High: $30	$35	1/9	4/9

TABLE 2. Average Prices, Predicted and Experimental, in Dollars

Variable	Session 1	Session 2
Competitive prediction	20.00	30.00
Profit-splitting prediction	20.70	29.30
Round 1 outcome	21.20	27.00
Round 2 outcome	21.20	28.50

combination of types and calculate the price predicted for such a matching. The expected fraction of each possible pairing of types of buyers and sellers and the price at which such a pair would transact under the profit-splitting hypothesis are shown in Table 1.

We used the entries in Table 1 to calculate the expected average price under profit splitting. In each session, transactions took place at three distinct prices, $15, $25, and $35, with differing probabilities in the two sessions. In session 1, the expected average price in each classroom was $15 × 4/7 + $25 × 2/7 + $35 × 1/7 = $21.2. A similar calculation showed that in session 2, the expected average price in each classroom was $15 × 1/7 + $25 × 2/7 + $35 × 4/7 = $29.3.

The experimental outcomes were closer to the predictions of the profit-splitting theory than to those of competitive-equilibrium theory (Table 2). It is especially interesting that the profit-splitting hypothesis correctly predicted that the average price in session 1 would be higher than the competitive-equilibrium prediction, and the average price in session 2 would be lower than the competitive prediction. Chamberlin (1948) observed that in his classroom experiments, the average price usually exceeded the competitive prediction, although he was unable to produce a theoretical explanation for this outcome.

Predicted and Actual Price Distribution

Competitive theory and the profit-splitting theory both make detailed predictions about the distribution of prices paid in the market. In a competitive equilibrium, *all*

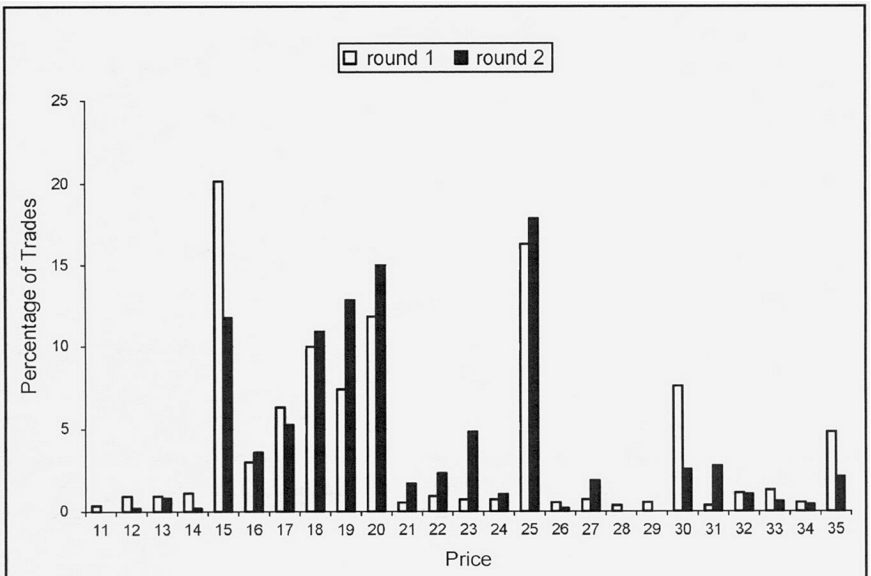

FIGURE 7. Price distribution in session 1.

trades take place at the same competitive price. The profit-splitting hypothesis predicts that all trades occur at one of three prices, $15, $25, or $35. Because our data included the prices from several hundred individual transactions, we could compare the distribution of actual prices in each round with the distributions predicted by the two theories. This was a much more stringent test than simple comparison of predicted to actual *average* prices. The detailed distribution pattern of transaction prices in each round of session 1 and session 2 are shown in the histograms of Figures 7 and 8.

In Tables 3 and 4, we show the predicted and actual percentages of transactions that were within $1 of each relevant price for the profit-splitting theory and also those within $1 of the competitive price in sessions 1 and 2, respectively. In both rounds of both sessions, the profit-splitting theory predicted far more trades at the "extreme" prices $15 and $35 than were actually observed. Although the data appeared to favor rejecting the hypothesis that all (or even a majority of) traders were profit splitters, the spikes observed at $15, $25, and $35 in Figures 7 and 8 suggested that at least a few traders did behave like profit splitters.

Does the competitive theory fare any better in explaining this data? In competitive equilibrium, all trades take place at a single competitive price. But in session 1, only 20 percent of all trades were within $1 of the competitive price in round 1 and 30 percent in round 2. This performance improved in session 2, where 32 percent of all trades were within $1 of the competitive price in round 1 and 42 percent in round 2.

In both sessions, trades tended to take place at prices closer to the competitive predictions in round 2, after traders had observed the round 1 prices at which

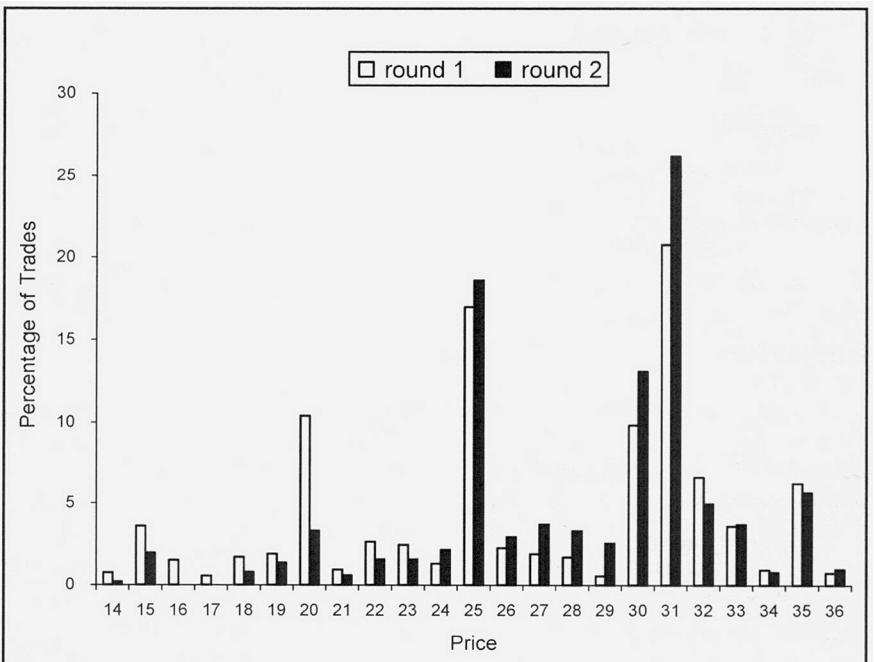

FIGURE 8. Price distribution in session 2.

TABLE 3. Percentage Actual and Predicted Prices in Session 1

	Price range			
Variable	$14–16	$24–26	$34–36	$19–21
Competitive prediction	0	0	0	100
Profit-splitting prediction	57	29	14	0
Actual shares, round 1	24	18	6	20
Actual shares, round 2	16	19	2	30

others traded. Prices in both rounds of session 2 were closer to the competitive predictions than in session 1. In the second round of session 2, the modal price was the competitive equilibrium price of $30 (Figure 7). Session 1 was the first market experiment that most participants had ever experienced. It appears that with the trading experience that they gained in the two rounds of session 1, participants acted more like competitive traders when they reached session 2. Given Smith's (1962) findings about the continued convergence of prices toward the competitive equilibrium through the first three rounds of trading, it is likely that a third round of trading in our experiments would also have pushed prices closer to the competitive values. It would be interesting to extend these classroom experiments to a third round in each session.

TABLE 4. Percentage Actual and Predicted Prices in Session 2

Variable	Price range			
	$14–16	$24–26	$34–36	$29–31
Competitive prediction	0	0	0	100
Profit-splitting prediction	14	29	57	0
Actual shares, round 1	7	20	8	32
Actual shares, round 2	2	24	8	42

Predicted and Actual Quantity Distribution

The competitive theory and the profit-splitting theory predict not only the total number of transactions, but they also predict the number of trades between each possible pair of types of trading partners. From the demand and supply schedules in Figure 1, we see that in competitive equilibrium, every low-cost supplier and no high-cost suppliers will trade. We also see that all high-value demanders and some low-value demanders will trade. Therefore in competitive equilibrium, the number of trades between high-value demanders and low-cost suppliers was equal to the total number of high-value demanders, and the number of trades between low-value demanders and low-cost suppliers must equal the difference between the total number of low-cost suppliers and the number of high-value demanders. In competitive equilibrium, no trades involved high-cost suppliers.

In competitive equilibrium for session 2, every high-value demander and no low-value demanders would trade, and every low-cost supplier and some high-cost suppliers would trade (Figure 2). In equilibrium, the number of trades between high-value demanders and low-cost suppliers equaled the number of low-cost suppliers, and the number of trades between high-value demanders and high-cost suppliers equaled the difference between the number of high-value demanders and the number of low-cost suppliers. No trades would involve low-value demanders.

The profit-splitting theory predicts that suppliers and demanders meet at random and trade on their first encounter, if the demander's buyer value exceeds the supplier's seller cost. If the number of suppliers equals the number of demanders, then everyone would meet somebody of the other type on a first encounter. The only pairs who did not trade on their first encounter were low-value buyers matched with high-cost sellers. It follows that those who failed to trade on a first encounter would not find anyone with whom they could make a profitable trade on later encounters. Given the fractions of low- and high-cost suppliers and low- and high-value demanders, we could calculate the expected number of pairings of each type.[7]

Enough detailed data were collected about trading outcomes so that we could compare detailed quantity outcomes with the theoretical predictions of these two competing theories. Predictions from each theory and actual results for sessions 1 and 2 are shown in Tables 5 and 6.

It is interesting that although price outcomes changed substantially between round 1 and round 2, there was relatively little change in the number of trades

TABLE 5. Predicted and Actual Quantities by Pair Type, Session 1

Buyer value (BV) and seller cost (SC)	Comp. equil.	Price splitting	Actual shares, round 1	Actual shares, round 2
BV: Low: $20 SC: Low: $10	197	290	221	218
BV: Low: $20 SC: High: $30	0	0	9	0
BV: High: $40 SC: Low: $10	241	145	207	209
BV: High: $40 SC: High: $30	0	73	34	38
Total No. trades	438	508	471	465

TABLE 6. Predicted and Actual Quantities by Pair Type, Session 2

Buyer value (BV) and seller cost (SC)	Comp. equil.	Price splitting	Actual shares, round 1	Actual shares, round 2
BV: Low: $20 SC: Low: 10	0	74	26	18
BV: Low: 20 SC: High: 30	0	0	6	2
BV: High: 40 SC: Low: 10	241	148	218	218
BV: High: 40 SC: High: 30	201	296	211	213
Total No. trades	442	518	461	451

taking place between matched pairs of each type. In almost every case where the two theories made different predictions, the outcome was closer to the prediction of competitive equilibrium than to that of the profit-splitting theory. In each case, however, the actual outcomes were between the two predictions.

CONCLUSIONS

In Chamberlin's (1948) experiment, demanders and suppliers traded only once in a decentralized pit-trading environment. In Smith's (1962) experiment, trading was by a public double-oral auction, and traders acted in three or more "trading days" in which participants faced the same market conditions on each new trading day as on previous days but with the common experience of the previous days' trading. Chamberlin found his experimental results to be far from the predictions of competitive-equilibrium theory, whereas Smith found that after three rounds of trading, prices were closely concentrated around the competitive price.

The Bergstrom and Miller (2000) classroom experiments shared some features with those of Chamberlin (1948) and some with those of Smith (1962). They used Chamberlin's pit-trading method rather than Smith's double-oral auction, but, like Smith, they included more than one round of each session. The demand and supply curves in these experiments were simpler than those of Chamberlin and Smith, having only two types of demanders and two types of suppliers, whereas Smith and Chamberlin had many distinct buyer values and seller costs. For these classroom experiments, competitive theory predicted total quantities and average transaction prices quite well. But even in the second round, a significant number of trades took place at prices substantially different from the equilibrium price. In the first round of either session, many participants appeared to split profits equally with the trading partner with whom they happened to be paired. In the second round, profit splitting became less common, but did not disappear. Although neither the competitive theory nor the profit-splitting theory satisfactorily explained detailed outcomes in the two rounds of each session, it appeared that as traders became more experienced with market conditions, their behavior became more like that predicted by competitive theory and less like that predicted by profit splitting.

We suspect that the real world has some markets that are better described by Smith's (1962) institutions and others by Chamberlin's (1948). Organized commodity markets and stock markets, where experienced traders engage repeatedly in public trading and where market fundamentals change only gradually, seem better approximated by Smith's design. In some markets, even if trading is repeated, the fundamentals of demand and supply change so rapidly that prices paid in previous periods offer little information to traders about the prices they can expect in the current period. Those markets may behave more like Chamberlin's experiment with a single round of trading. The classroom experiments that we studied lie somewhere in between.

NOTES

1. According to the *Social Science Citation Index* (Institute for Scientific Information), Chamberlin's (1948) article was cited only four times between 1948 and 1962.
2. Data from classroom experiments have also been used to study the effectiveness of classroom experiments as a teaching device. Studies by Emerson and Taylor (2004), Dickie (2000), and Mullin and Sohan (1999) compared improvements in scores on standardized economics tests for classes that used experiments from the Bergstrom and Miller (2000) text with those of classes that used standard lecture-oriented textbooks. Emerson and Taylor and Dickie found in their respective studies that the use of experiments significantly increased the improvement in test scores, whereas Mullin and Sohan found only a weak positive effect in their study.
3. Most participants had not yet studied the theory of supply and demand. Even if they understood competitive equilibrium theory, they would know neither the demand curve or the supply curve and thus could not deduce the equilibrium price.
4. Chamberlin (1948, 97) reported that "no statistical computations for the entire sample of forty-six experiments have been made."
5. In the Bergstrom and Miller (2000) classroom experiments, the average variance of prices decreased by about 59 percent between the first and second round.
6. Those who failed to trade in their first encounter might seek another trading partner, but the only traders who did not find a partner in the first round would be low-value demanders and high-cost sellers, who could not make mutually profitable deals with each other.
7. Because of variations in the number of students who come to class, the number of suppliers and demanders were not equal in all classes. We took a simplified approximation by calculating

appropriate fractions of the minimum of the total numbers of demanders and suppliers across all experiments. Numerical simulations indicated that this simplification made little difference.

REFERENCES

Bergstrom, T. C. 2004. *A proof of Chamberlin's excess trading conjecture.* Technical report, University of California, Santa Barbara, June.

Bergstrom, T., and J. Miller. 2000. *Experiments with economic principles: Microeconomics.* New York: McGraw-Hill.

Chamberlin, E. 1948. An experimental imperfect market. *Journal of Political Economy* 56 (April): 95–108.

Dickie, M. *Experimenting on classroom experiments: Do they increase learning in microeconomics?* Forthcoming, *Journal of Economic Education*, 2006.

Emerson, T., and B. Taylor. 2004. Comparing student achievement across experimental and lecture-oriented sections of a principles of microeconomics course. *Southern Economic Journal* 70 (3): 672–93.

Institute for Scientific Information. 2004. *Social Science Citation Index.* http://www.isinet.com/products/citation/ssci/.

Mullin, D., and G. Sohan. 1999. *Benefit assessment of classroom experimental economics.* Technical report, United States Air Force Academy, Colorado Springs, Colorado.

Smith, V. L. 1962. An experimental study of competitive market behavior. *Journal of Political Economy* 70 (April): 11–37.

———. 1991. Experimental economics at Purdue. In *Papers in experimental economics*, ed., V. L. Smith. Cambridge: Cambridge University Press. Originally appeared in G. Horwich and J. P. Quirk, eds., *Essays in contemporary fields of economics.* Purdue, IN: Purdue University Press.