

UC Santa Barbara
Ted Bergstrom Papers

Title

Extracting Valuable Information from Classroom Trading Pits

Permalink

<https://escholarship.org/uc/item/7893c1mz>

Author

Bergstrom, Ted

Publication Date

2004-07-05

Abstract

Edward Chamberlin, who initiated classroom market experiences, used the results of his experiments to argue that competitive equilibrium performs poorly in explaining the outcomes of real markets. Vernon Smith altered the design of Chamberlin's experiment so as to increase the amount of price information available to traders and in classroom experiments with this design found that trading outcomes were close to those predicted by competitive theory. This paper examines results of classroom trading experiments using the design found in *Experiments with Economic Principles* [2], an introductory economics text by Ted Bergstrom and John Miller. The procedure in this experiment is intermediate between that of Chamberlin and that of Smith. We have collected data on transaction prices and quantities from a large number of classroom experiments using this design. We compare the experimental outcomes with the predictions made by competitive equilibrium theory and by a simple profit-splitting theory. Evidence suggests that neither theory is entirely successful, though in the first rounds of trading there seems to be a significant amount of profit-splitting and as traders become more experienced, outcomes are closer to those predicted by competitive theory. profit-splitting.

Extracting Valuable Data from Classroom Trading Pits

Theodore C. Bergstrom and Eugene Kwok¹

July 5, 2004

¹Theodore C. Bergstrom is the Aaron and Cherie Raznick Professor of Economics, University of California at Santa Barbara, Santa Barbara, California. Eugene Kwok's work on this paper was performed while he was an undergraduate student at UCSB.

Experimental economics began in the 1940's 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. Chamberlin “induced” market demand and supply by distributing cards that assigned each participating student a role either as a supplier or a demander. Each supplier was assigned a seller cost at which she could supply a single unit and each demander a buyer value for single unit of the good. In any sale, the seller's profit is the difference between the price and her seller cost, while the buyer's profit is the difference between his 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 for all to see. Trading continued until no more supplier-demander pairs were willing to make trades.

Chamberlin described his classroom experiments in 1948 in the *Journal of Political Economy* [3], but this pathbreaking work received little attention¹ until 1962, 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 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,” in Smith's *Papers in Experimental Economics*. [5]

Most experimental economics research is now conducted with paid subjects outside the classroom. There are many good reasons for this, one of which is that if you are paying subjects, you can subject them to repetitive activities that tuition-paying students would find boring and uninteresting.

¹According to the Social Science Citation Index, Chamberlin's paper was cited only 4 times between 1948 and 1962.

But the results of classroom experiment 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.

This paper examines results of a classroom trading experiment designed by Ted Bergstrom and John Miller and published in an introductory economics text, *Experiments with Economic Principles* [2]. This experiment has been conducted in several hundred classrooms. For many of these classes the results have been preserved and recorded in convenient form, since experimental results are typically reported to the students as spreadsheets posted on the web. We have collected data on transaction prices and quantities from 31 classrooms at 10 universities for two sessions of a simple demand and supply experiments.

A Supply and Demand Experiment

In this experiment, each participant is assigned a role as a supplier or a demander of apples. Suppliers can sell at most one unit (a bushel) and demanders can buy at most one unit. Each supplier is assigned one of two possible “seller costs” and each demander is assigned one of two possible “buyer values” for a bushel of apples. Buyers and sellers are asked to roam around the room and try to make as profitable a deal as possible. When a seller and a buyer agree on a price, they write this price on a sales contract, along with their identification numbers and the seller cost and buyer value. The market manager records transaction prices on the blackboard for all to see as the contracts are turned in. If a seller with seller cost C sells a unit at price P to a buyer with buyer value B , then the seller’s profit is $P - C$ and the buyer’s profit is $B - P$.

This experiment included two sessions with different distributions of buyer

values and seller costs. Each session consists of two rounds of trading. In each session, after the first round of trading is completed and students have observed the results, students are asked to play again with the same buyer values and seller costs as in the second round. The only thing that is new in the second round is the experience that participants have gained from the first round.

Competitive Demand and Supply Curves

The number of persons with each buyer value and seller cost differs between classrooms, depending on the number of students in the class. But the distribution of types is chosen so that equilibrium prices and the qualitative features of supply and demand are the same in all classes. In every class there are two types of suppliers, high cost suppliers with seller cost of \$30 and low cost suppliers with seller cost of \$10 per bushel. There are 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 are approximately twice as many low value demanders as high value demanders and about twice as many low cost sellers as high cost sellers. Figure 1 shows the competitive supply and demand curve and the competitive equilibrium price and quantity in for a class 47 students in Session 1. As the graph shows, the competitive equilibrium price is \$20 and the competitive equilibrium quantity is 15 units sold.

Session 2 has the same two types of buyers and two types of sellers, but this time there are twice as many high cost suppliers as low cost suppliers and twice as many high value demanders as low value demanders. In this session, the competitive equilibrium price is \$30. Figure 2 shows the competitive supply and demand curve and the competitive equilibrium price and quantity for a class of 47 students.

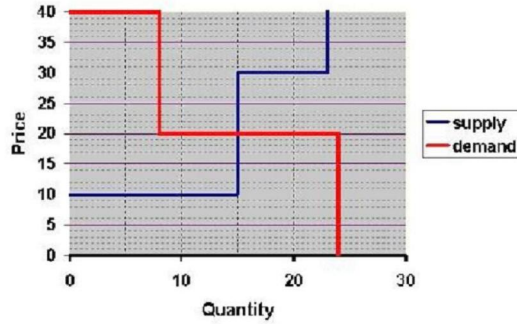


Figure 1: Supply and Demand in Session 1

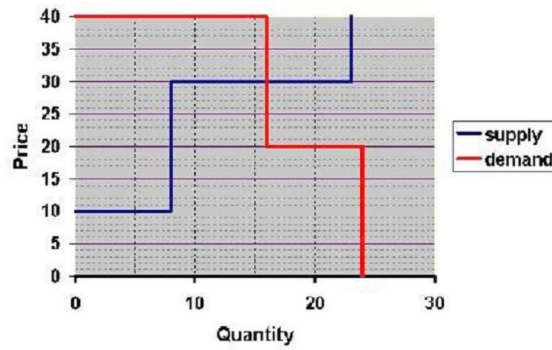


Figure 2: Supply and Demand in Session 2

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.² Given that individuals know

²Typically participants have not yet studied the theory of supply and demand. But 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.

little about market conditions when they participate in the first round, we would not expect all transactions to take place at the competitive equilibrium price.

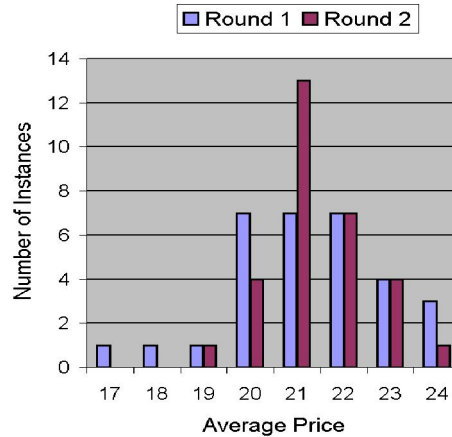


Figure 3: **Average Prices in classrooms: Session 1, Rounds 1 and 2**

Figure 3 shows the distribution of classroom mean prices in Rounds 1 and 2 of Session 1, for the 31 classrooms in our study. The competitive equilibrium price for this session is \$20. Even in the first round of Session 1, the average price in most classrooms is 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 prices.

Figure 4 shows the distribution of mean prices in the first and second round of Session 2, where the competitive equilibrium price is \$30. Session 2 takes 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 and, as we see from the figure, in Round 1 of Session 2 students seem to have expected that prices would be lower than

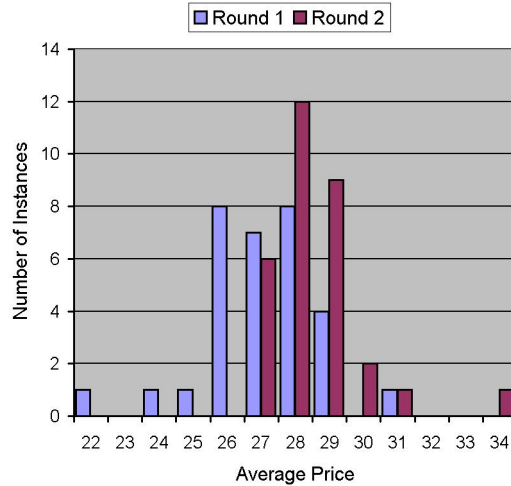


Figure 4: **Average price in classrooms, Session 2, Rounds 1 and 2**

the \$30 equilibrium price. In round 2, having experienced the outcome of Round 1, students appear to have adjusted their expectations upward and the average price in most classrooms moved closer to the equilibrium price. It is interesting to note that while 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 is \$2 *below* the equilibrium price.

The average price in a classroom is 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. Figures 5 and 6 show the distribution across classrooms of differences between the

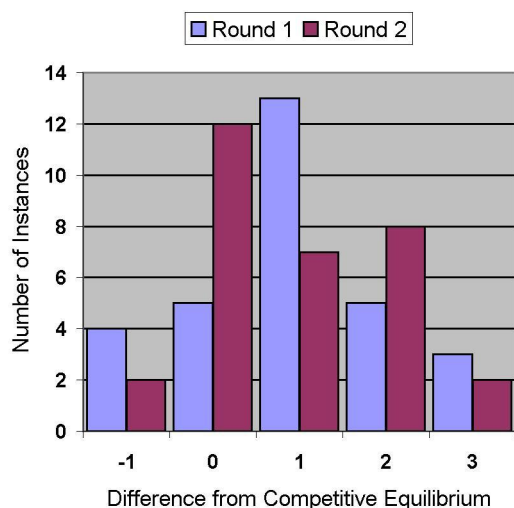


Figure 5: **Quantity Difference in Session 1**

quantity actually sold and the competitive equilibrium quantity.

As these figures show, the number of trades exceeded the competitive outcome more often than it fell short. This tendency for “excess trade” was also remarked by Chamberlin [3]. Chamberlin used a numerical example to make a plausible case that non-tatonnement pit-trading is likely to lead to “too much” rather than “too little” trading. He does not provide a proof of this assertion, nor does he spell out exactly what is to be proved. Bergstrom [1] states and proves a result that supports 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 to note that Professor Chamberlin, who originated

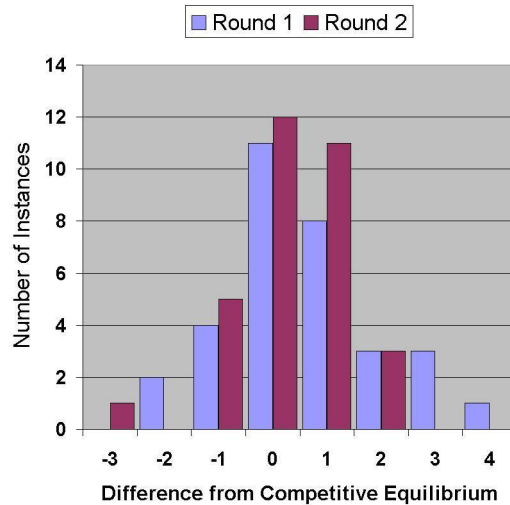


Figure 6: **Quantity Difference in Session 2**

classroom market experiences emphasized the differences rather than the similarities between experimental results and the predictions of competitive theory. He reports on the results of forty-six experiments conducted in his Harvard classroom over the years. According to Chamberlin [3]:

“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 did not supply further details about his results.³ 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 upwards in the case of quantity and possibly downwards in the case of price.

³He reports that “no statistical computations for the entire sample of forty-six experiments have been made.

Chamberlin saw no reason to expect that that the outcome in his experiment would approximate competitive equilibrium. He points out that in his classroom experiments, as in real-world trading, there is 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 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.

Chamberlin’s experimental results correspond to just the first round of our experiment, since he did not offer a second round of trading. Our experiments also differ from his in that he had many distinct buyer values and seller costs, while our experiment had only two possible buyer values and two possible seller costs. While our classroom experiments show a tendency toward excess trading in the first round, this tendency is 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, while in Session 2, the observed price was usually lower.

Smith's Results

Vernon Smith decided to revise Chamberlin's procedures so as to give the competitive model a better chance. Smith explains 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?” [5]

Smith's first published discussion of the results of his classroom experiments appeared in the *Journal of Political Economy* in 1962 [4], where he reports 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.

In Smith's view, real markets typically “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

experiments typically included three to five trading days, corresponding to the “rounds” in our experiment.

Smith [4] reports the results of ten classroom experiments with differing shapes of supply and demand curves. Some of these experiments have additional differences that make them hard to compare either with Chamberlin’s experiments or our own. Smith’s first four experiments differ 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 decreases from the first round to the second and again from the second round to the third. Across the four experiments on average the variance in the second round is about 55% of that in the first round and the average of the variance in the third round is about 60% of that in the second round. Variances in the fourth round are little different from those in the third. In our own sample of classroom experiments, the average ratio of the standard deviation of prices in the second round to that in the first was about 77%.

Do Classroom Results Really Support Competitive Theory?

The Bergstrom-Miller experimental design follows Smith in adding a second round of trading. To save classroom time, most instructors do not conduct a third or fourth round, as did Smith. We follow Chamberlin’s 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 are reasonably close to the predictions of competitive theory. Thus it is usually easy to convince credulous undergraduates that competitive theory has impressive predictive power. But does this conclusion

withstand closer scrutiny? Might there be a simple competing theory that works as well or better?

Although competitive theory makes reasonably good predictions of average prices, we will show that there is a plausible competing theory that predicts average prices in our classroom experiments more closely. A better test of the competitive theory and of competing alternatives requires us to examine the detailed predictions of each theory. It is important to recognize that competitive theory makes many predictions besides 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 number 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 following “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 seller’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.⁴

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 have “low” values of \$20 and 1/3 have “high” values of \$40. About 2/3 of the suppliers have “low” costs of \$10 and 1/3

⁴But as we will see, they will not succeed in finding a partner with whom they can trade.

have “high” costs of \$30. If encounters are random, then on average, 4/9 of the encounters will be between low value demanders and low cost sellers, 2/9 of the encounters will be between low value demanders and high cost sellers, 2/9 will be between high value demanders and low cost sellers and 1/9 will 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 cost, a sale will take place at a price midway between. The only individuals who do not make a trade with the first person they meet are the low value demanders with \$20 buyer values who meet high cost suppliers with \$30 seller costs.⁵

In Session 2, about 1/3 of the demanders have low values and 2/3 have high values, while 1/3 of the suppliers have low costs and 2/3 have high costs. As with Session 1, we can calculate the fraction of all matchings of each possible combination of types and calculate the price predicted for such a matching. Table 1 reports 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.

Table 1: Matching and Prices under Profit-Splitting Hypothesis

Buyer Value	Low: \$20	Low: \$20	High: \$40	High: \$40
Seller Cost	Low: \$10	High: \$30	Low: \$10	High: \$30
Predicted Price	\$15	no trade	\$25	\$35
Fraction, Sess 1	4/9	2/9	2/9	1/9
Fraction, Sess 2	1/9	2/9	2/9	4/9

We use the entries in Table 1 to calculate the expected average price under profit-splitting. We see that in each session, transactions take place

⁵Those who fail to trade in their first encounter may seek another trading partner, but the only traders who didn’t find a partner in the first round will be low value demanders and high cost sellers, who can not make mutually profitable deals with each other.

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 is $\$15 \times 4/7 + \$25 \times 2/7 + \$35 \times 1/7 = \21.2 . A similar calculation shows that in Session 2, the expected average price in each classroom is $\$15 \times 1/7 + \$25 \times 2/7 + \$35 \times 4/7 = \29.3

Table 2: Predicted Prices and Experimental Average Prices

	Session 1	Session 2
Competitive Prediction	\$20	\$30
Profit-Splitting Prediction	\$20.7	\$29.3
Round 1 Outcome	\$21.2	\$27.0
Round 2 Outcome	\$21.2	\$28.5

Table 2 shows that the experimental outcomes are closer to the predictions of the profit-splitting theory than to those of competitive equilibrium theory. It is especially interesting that the profit-splitting hypothesis correctly predicts that the average price in Session 1 will be higher than the competitive equilibrium prediction and the average price in Session 2 will be lower than the competitive prediction. Chamberlin [3] observed that in his classroom experiments, the average price usually exceeded the competitive prediction, though 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* 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. Since our data includes the prices from several hundred individual transactions, we can compare the distribution of actual prices in each round with the distributions predicted by the two theories. This is a much more stringent test than simple comparison of predicted and actual average prices across classrooms. Tables 3 and 4 show the predicted and actual percentages of transactions that are 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. The histograms in Figures 7 and 8 display the detailed distribution pattern of transaction prices in each round of Session 1 and Session 2.

Table 3: Actual and Predicted Prices, Session 1

Price Range	\$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%

Table 4: Actual and Predicted Prices, Session 2

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%

The profit-splitting theory predicts far more trades at the “extreme” prices \$15 and \$35 than are actually observed. In Session 1, this theory predicts that 57% of all trades will be at \$15 and about 14% will be at \$35.

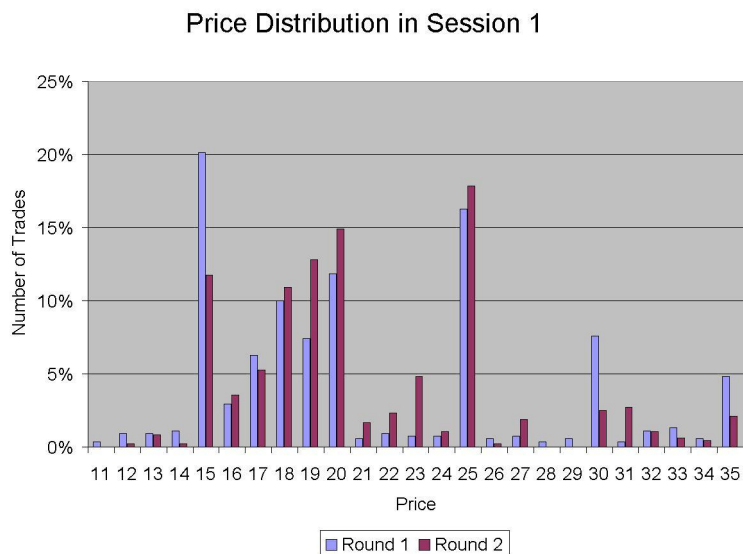


Figure 7: Price Distribution in Session 1

In Round 2 of this session, the fraction of prices that were within \$1 of these prices are respectively 16% and 2%. Profit-splitting theory predicts that in Session 2 about 57% of trades will be at \$35 and about 14% at a price of \$15. In Round 2 of Session 2, the actual percentages of trades at these prices are 8% and 2% respectively. Although the data appears to reject the hypothesis that all (or even a majority of) traders are profit-splitters, the spikes observed at \$15, \$25, and \$35 in Figures 7 and 8 suggest that at least a few traders do behave like profit-splitters.

Does the competitive theory fare any better in explaining this data? The competitive theory predicts that all trades take place at a single competitive price. But in Session 1, only 20% of all trades are within \$1 of the competitive price in Round 1 and 30% in Round 2. This performance improves in Session 2, where 32% of all trades are within \$1 of the competitive price in Round 1 and 42% in Round 2.

In both sessions, we see that trades tend to take place at prices closer

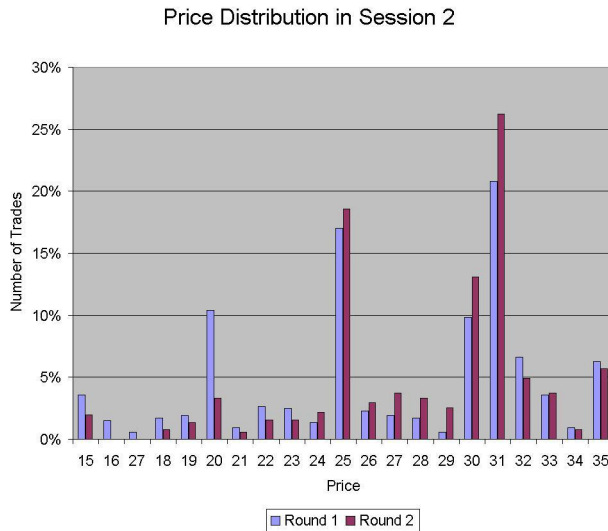


Figure 8:

to the competitive predictions in Round 2, after traders have observed the Round 1 prices at which others traded. We also see that prices in both rounds of Session 2 are closer to the the competitive predictions than in Session 1. As Figure 7 shows, in the second round of Session 2 the modal price is the competitive equilibrium price of \$30. Session 1 is the first market experiment that most participants have ever experienced. It appears that with the trading experience that they gained in the two rounds of Session 1, participants act more like competitive traders when they reach Session 2. Given Smith’s findings about the continued convergence of prices toward the competitive equilibrium through the first three rounds of trading, it seems likely that if we had a third round of trade in each session, prices would approach equilibrium more closely. These results suggest that it might be worthwhile for this experiment to be run with three rather than two rounds per session.

Predicted and Actual Quantity Distribution

The competitive theory and the profit-splitting theory predict not only the total number of transactions, but 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 the number of trades between high value demanders and low cost suppliers is 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 there are no trades involving high cost suppliers.

From Figure 2 we see that in competitive equilibrium for Session 2, every high value demander and no low value demanders will trade, and that every low cost supplier and some high cost suppliers will trade. In equilibrium the number of trades between high value demanders and low cost suppliers equals the number of low cost suppliers and the number of trades between high value demanders and high cost suppliers equals the difference between the number of high value demanders and the number of low cost suppliers. There will be no trades involving 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 and of demanders are equal, then everyone will meet somebody of the other type on a first encounter. The only pairs who do not trade on their first encounter are low value buyers matched with high cost sellers. It follows that those who fail to trade on a first encounter will not find anyone with whom they can make a profitable trade on later encounters. Given the fractions of low and high cost

suppliers and low and high value demanders, we can calculate the expected number of pairings of each type.⁶

Enough detailed data was collected about trading outcomes so that we can compare detailed quantity outcomes with the theoretical predictions of the two competing theories. These predictions and actual results for Sessions 1 and 2 are shown in Tables 5 and 6.

Table 5: Predicted and Actual Quantities in Session 1

Buyer Value Seller Cost	Low: \$20 Low: \$10	Low: \$20 High: \$30	High: \$40 Low: \$10	High: \$40 High: \$30	Total No. Trades
Comp. Equil.	197	0	241	0	438
Price-Splitting	290	0	145	73	508
Actual, Rd 1	221	9	207	34	471
Actual, Rd 2	218	0	209	38	465

Table 6: Predicted and Actual Quantities in Session 2

Buyer Value Seller Cost	Low: \$20 Low: \$10	Low: \$20 High: \$30	High: \$40 Low: \$10	High: \$40 High: \$30	Total No. Trades
Comp. Equil.	0	0	241	201	442
Price-Splitting	74	0	148	296	518
Actual, Rd 1	26	6	218	211	461
Actual, Rd 2	18	2	218	213	451

It is interesting to notice that although price outcomes change substan-

⁶Because of variations in the number of students who come to class, the number of suppliers and demanders were not exactly equal in all classes. We take a simplified approximation by calculating appropriate fractions of the minimum of the total numbers of demanders and suppliers across all experiments. Numerical experiments suggest that the magnitude of the differences implied by such an elaboration is small.

tially between round 1 and round 2, there is relatively little change in the number of trades taking place between matched pairs of each type. In almost every case where the two theories make different predictions, the outcome is closer to the prediction of competitive equilibrium than to that of the profit-splitting theory. In each case, however the actual outcomes are between the two predictions.

Conclusions

In Chamberlin's experiment, demanders and suppliers traded only once in a decentralized pit-trading environment. In Smith's experiment, trading was by a public double oral auction, and traders acted in three or more "trading days" where in each new trading day, participants faced the same market conditions 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, while Smith found that after three rounds of trading, prices were closely concentrated around the competitive price.

In the real world, organized commodity markets and stock markets seem to be best approximated by Smith's design in which trading is public and experienced traders trade repeatedly in an environment where market fundamentals change little from day to day. In some markets the fundamentals of demand and supply may change so rapidly that the prices paid in previous periods offer little information to traders about the prices that they can expect in the current period. These markets may behave more like Chamberlin's experiment with a single round of trading.

The design of the classroom experiments studied here is intermediate between that of Chamberlin and that of Smith. These experiments use Chamberlin's pit-trading method rather than Smith's double oral auction. These

experiments include two rounds of each session rather than Chamberlin's one or Smith's three or more.⁷ For these classroom experiments, competitive theory predicts total quantities and average transaction prices quite well. But even in the second round, a significant number of trades take place at prices substantially different from the equilibrium price. In the first round of either session, many participants appear to split profits equally with the trading partner they happen to be paired with. In the second round, profit-splitting becomes less common, but does not disappear. Although neither the competitive theory nor the profit-splitting theory satisfactorily explains detailed outcomes in the two rounds of each session, it appears that as traders become more experienced with market conditions, their behavior becomes more like that predicted by competitive theory and less like that predicted by profit-splitting.

⁷Our demand and supply curves were simpler than those of Chamberlin and Smith. We had only two types of demanders and two types of suppliers while Smith and Chamberlin each had many distinct buyer values and seller costs.

References

- [1] Theodore C. Bergstrom. Chamberlin's excess trading conjecture. Technical report, University of California Santa Barbara, 2004.
- [2] Theodore Bergstrom and John Miller. *Experiments with Economic Principles: Microeconomics*. McGraw-Hill, New York, 2000.
- [3] E.H. Chamberlin. An experimental imperfect market. *Journal of Political Economy*, April:95–108, 1948.
- [4] Vernon L. Smith. An experimental study of competitive market behavior. *Journal of Political Economy*, 70(2):111–37, April 1962.
- [5] Vernon L. Smith. Experimental economics at Purdue. In Vernon L. Smith, editor, *Papers in Experimental Economics*. Cambridge University Press, Cambridge, England, 1991. Originally appeared in *Essays in Contemporary Fields of Economics*, edited by G. Horwich and J.P. Quirk. Purdue University Press, 1981.