

Tax and Subsidy Incidence Equivalence Theories: Experimental Evidence from Competitive Markets

Bradley J. Ruffle*

Department of Economics, P.O. Box 653

Ben-Gurion University

Beer Sheva, Israel 84105

e-mail: bradley@bgumail.bgu.ac.il

October 2001

Abstract

A basic tenet in elementary microeconomics is tax incidence equivalence. This tenet holds that the burden of a unit tax on buyers and sellers is independent of who actually pays the tax. By contrast, policymakers and the public often mistake statutory incidence for economic incidence. Recent evidence of the empirical validity of tax incidence equivalence is mixed. In this paper, using competitive laboratory markets, I test both tax incidence equivalence and an analogous theorem I refer to as subsidy incidence equivalence. For sufficiently large markets, the results show strong support for both theories. In these markets, there is little to no evidence, even in the short run, of the popular misperception that statutory incidence equals economic incidence. In smaller markets in which competitive forces are weaker and relative bargaining strengths may play a role, the evidence for tax incidence equivalence is weaker as minor price discrepancies may persist between markets.

keywords: tax incidence equivalence, subsidy, competitive markets, experimental markets, framing effects

*I thank Suleiman Abu Bader, Ted Bergstrom, Juergen Bracht, Jeremy Clark, Alex Cukierman, Moshe Justman, Todd Kaplan, Marvin Schrieder, Uri Simonsohn, Avia Spivak, Bart Wilson and seminar participants at Ben-Gurion University, Osaka University, Tel Aviv University, the 2001 Annual Economic Science Association Meetings and the 2001 Far Eastern Econometric Society Meetings for comments. I also thank my research assistants, Inna Agroskin, Gil Attia, Michal Livne, Iris Mini, Lilach Sudak, Tsafi Vana and Liat Winner, as well as Noam Green for programming the software for the transaction graphs. Financial support from Ben-Gurion University is gratefully acknowledged. The incidence of any remaining shortcomings must be fully borne by me.

1 Introduction

A basic tenet in elementary microeconomics is tax incidence equivalence, alternatively referred to as tax liability-side equivalence. This tenet holds that the burden of a unit tax on buyers and sellers is independent of who actually pays the tax. In other words, the price paid by the buyers will be the same no matter whether the buyers or sellers pay the tax; similarly, the price that the sellers receive (which equals the price that the buyers pay minus the sum of the tax) will be the same no matter who pays the tax. Thus, the economic incidence of the tax is independent of the statutory incidence. Instead, the relative tax burden depends solely on the relative elasticities of supply and demand.

We teach this fundamental proposition in introductory economics (see, for instance, Mankiw, 1997, pp. 121-124) as well as public finance classes (e.g. Rosen, 1999, pp. 260-265). Tax incidence equivalence theory is robust to market structure and to the type of good or service upon which the tax is imposed. It also extends naturally to subsidies: just as the relative burden of a tax is independent of who pays the tax, the relative benefit of a subsidy is independent of who actually receives the subsidy. For want of a better phrase, I refer to this theorem as “subsidy incidence equivalence”.

Despite its central place in elementary microeconomics, tax incidence equivalence and the subject of tax incidence more generally remain poorly understood outside of academic economics. The media and the public often mistake statutory incidence (who pays the tax) for economic incidence (who bears its burden). This can be seen in debates over who should pay the tax rather than how the burden of the tax will be distributed between producers and consumers. The much publicized debate over the cigarette tax represents a case in point. Those who want to discourage smoking advocate taxing the consumer with the expectation that the price of cigarettes will increase. Those who believe the tobacco companies should be punished focus on taxing the producer, expecting the price received to decrease (Rosen, 1999, p. 265). While both sides are correct in their assessment of the impact on price of a tax, they both ignore that a

tax, no matter upon whom it is levied, will simultaneously increase the price paid by the consumer and decrease the price received by the producer (assuming neither supply nor demand is perfectly inelastic). Who pays the tax is inconsequential to the relative amounts by which these prices will change.

The irony of the cigarette tax debates is the simultaneous and long-running existence of a subsidy for tobacco farmers. In the U.S. for instance, the House of Representatives and Congress voted in 1997 to continue to subsidize tobacco farmers by voting down a bill to end the \$34 million subsidy program (Associated Press, 1997). That policymakers and the public can support the taxation of cigarettes and the simultaneous subsidization of tobacco (cigarettes' primary input) reflects the extent of their ignorance about the implications of tax and subsidy incidence. More specifically, the coexistence of cigarette taxes and tobacco subsidies suggests that non-economists tend to regard each product separately, even in the case of a product and its input. This compartmentalization leads to the misguided thinking that the taxation of cigarettes is unrelated to, and therefore could somehow be consistent with, the subsidization of tobacco.¹

An equally absurd display of public and media ignorance on the subject of tax incidence concerns the public outcry and political backlash from the Canadian Conservative Party's replacement of a manufacturers' sales tax with a consumer sales tax on January 1, 1991. From 1924 to 1990, goods and services made in Canada were subject to a tax at every stage of the manufacturing process. This form of multiple taxation placed Canadian exporters at a severe disadvantage and favored imported goods, not to mention inflated the prices Canadians paid for domestically produced goods and services. Sound economics notwithstanding, the shift in statutory incidence was perceived by the media and the public as another government tax grab. The already unpopular Conservative Party leader and Prime Minister, Brian Mulroney, was forced to resign within a year

¹Political considerations are needed to begin to understand this misguided economic thinking: for example, acquiescence to the powerful tobacco farmers' lobby and pandering to the vote from rural America underlie, in part, the tobacco subsidy, whereas cigarette taxes help curb rising health costs and are politically popular.

after introducing the consumer sales tax.

Policymakers do not always appear better informed than the public on the subject of tax incidence. We already noted the folly of the simultaneous taxation of cigarettes and subsidization of tobacco. In addition, Krugman (2000) castigates President Bush Jr.'s (then a presidential candidate) proposal to cut consumer gasoline taxes as a means to combat rising gasoline prices, arguing that, given the inelastic demand for gasoline, OPEC and U.S. oil refiners would respond by raising oil prices.²

Recent experimental evidence on the validity of tax incidence equivalence is mixed. Kerschbamer and Kirchsteiger (2000) would seem not to fault the fact that the media and policymakers repeatedly ignore the implications of tax incidence equivalence. Their results challenge its empirical validity, as indicated by their provocative title, "Theoretically robust but empirically invalid? An experimental investigation into tax equivalence". Consistent with popular perceptions, Kerschbamer and Kirchsteiger show that the side of the market that pays the tax shoulders a disproportionate share of the tax burden. Using a more complicated and less focused experimental design, an earlier study by Kachelmeier et al. (1994) is less conclusive. On the one hand, prices show signs of converging to the post-tax competitive equilibria, as predicted by tax incidence equivalence. However, they also report evidence of a finding contrary to that of Kerschbamer and Kirchsteiger, namely, that the side of the market that pays the tax actually bears a smaller tax burden relative to the equilibrium prediction.

Is tax incidence equivalence a mere theoretical construct with little empirical validity? Kachelmeier et al.'s test of tax incidence equivalence involves three sellers facing three buyers in a double-auction market. Kerschbamer and Kirchsteiger's test uses the ultimatum game, a bilateral monopoly with a particular timing and payoff structure. Borck et al. (2001) use the posted-offer market, a multi-player extension of the ultimatum game, with three buyers and three sellers. Their results show support for the theory.

²Borck et al. (2001) provide numerous additional examples of publicized controversies reflecting European policymakers' ignorance of tax incidence equivalence.

All three of these previous tests of tax incidence equivalence are very particular tests of the theory involving either a bilateral monopoly or precisely three buyers versus three sellers. In this paper, I offer a test of both the tax incidence equivalence and subsidy equivalence theorems involving a more competitive market structure in which I vary the number of traders between 16 and 30 in each market. My results show strong support for both theorems, particularly in larger (i.e. more competitive) markets. In addition, the results are robust to the composition of the individuals who make up the market. I devise a test for the popular misperception that statutory incidence equals economic incidence. I find little to no evidence of the existence of this misperception even in the short run. In smaller markets, however, my results indicate that the relative negotiating strengths of the parties may influence prices with the result that minor price discrepancies between markets may persist in the direction predicted by popular misperceptions. In such cases, the incidence theorems may not necessarily hold.

2 Background on Tax Incidence Equivalence

The logic underlying tax incidence equivalence is that the relative burden of a unit tax is solely a function of the relative elasticities of the supply and demand curves, not who pays the tax. Part of the appeal of the tax and subsidy incidence equivalence theorems lies in their robustness to different market structures: they hold in markets characterized by monopoly, bilateral monopoly, oligopoly and competitive conditions. I refer the interested reader to introductory economics and public finance textbooks for expositions of the theory. See, for example, Mankiw, 1997, or, for a particularly detailed exposition of tax incidence and the equivalence of various taxes, Stiglitz, 1988, chapter 17. Here I illustrate the analogous theory for subsidies. Figure 1 demonstrates subsidy incidence equivalence graphically for discrete supply and demand curves, each consisting of 12 units of production. (These same demand and supply configurations will also serve as the basis for the experimental design presented in section 3.) Figure 1a displays the pre-subsidy, competitive equilibrium price range between 31 and 33, as indicated by the

two-unit vertical overlap of the supply and demand curves at the quantity of 10.

The payment of a 10-unit subsidy to the sellers effectively lowers their costs by 10 units, thereby shifting outward the supply curve by this amount, as shown in Figure 1b. The symmetric supply and demand curves imply that the benefit of the subsidy is divided equally so that the competitive equilibrium price range (the price received by the sellers) increases by five units to 36 to 38, whereas, the price paid by the buyers (which does not include the subsidy) falls by five units to 26 to 28. On the other hand, if the subsidy is paid to the buyers, their valuations on each unit of demand increase by 10 units, thereby shifting outward the demand curve by the amount of the subsidy, as displayed in Figure 1c. As in the case of the subsidy to the sellers, the price received by the sellers increases by five units to 36 to 38 in equilibrium. Since the buyers received the 10-unit subsidy the actual price paid by them is between 26 and 28 in equilibrium, precisely as in the case in which the subsidy is given to the sellers. Furthermore, as shown in Figures 1b and 1c, the equilibrium quantity increases in both cases by two to 12 units with the addition of the subsidy. It follows from this analysis that the distribution of the benefit of the payment of a subsidy is independent of whether the buyers or the sellers actually receive the subsidy, but rather depends solely on the relative elasticities of supply and demand.

[insert Figure 1 here]

Despite the central place tax incidence equivalence occupies in microeconomics in general and public finance more specifically, there are no empirical tests of the theory using data from the wild. The difficulty lies in the absence of data: to test tax incidence equivalence requires price data on a commodity or service upon which the statutory incidence fell upon the producers, at some point in time and, at another point in time, the statutory incidence of the tax shifted entirely to consumers. Policymakers don't often accommodate economic research with such stark turnabouts in policy.

Perhaps the closest such shifts in policy occurred in Japan, New Zealand and Canada in the late 1980s and early 1990s. Faced with growing budget debts and increased foreign

competition, New Zealand and Canada replaced a narrowly based manufacturers' sales tax with a broader based consumption sales tax. Besides expanding the base of goods and services subject to the sales tax, the intent of these governments was to eliminate the double, or even multiple, taxation of business inputs and the resultant distortions in the effective tax rates of various manufactured goods. For instance, Michael Wilson, the Finance Minister of Canada responsible for replacing Canada's manufacturers' sales tax with the current consumption sales tax in 1991, estimated that "[a]pproximately one-half of the total sales tax collected [under the former manufacturers' sales tax was] derived from tax on business inputs such as transportation equipment, office equipment, and building materials. The tax on these inputs can then result in cascading, as these inputs may be used to produce goods that are subsequently taxed again" (Wilson, 1987, pp. 46-47). Furthermore, to continue with the Canadian example,³ levying the tax on the producer placed Canadian exporters at a disadvantage in countries like the U.S. (Canada's largest trading partner) in which the sales tax has always fallen on consumers. By the same reasoning, goods imported from the U.S. were at an advantage compared to domestic Canadian goods. The replacement of the manufacturers' sales tax with a consumption sales tax therefore neither impeded export performance nor favored imports.

Apart from these motivations, the replacement of a manufacturers' sales tax with a consumption sales tax is suggestive of the Canadian, and other governments', recognition that statutory incidence does not equal economic incidence. That is to say, by shifting the statutory incidence of the sales tax from the manufacturer to the consumer, the government's primary motivation was to replace an anticompetitive and distortionary sales tax with a more efficient and uniform one, rather than to shift the economic burden of the tax.

Still, these instances of shifts in statutory incidence policy do not easily afford opportunities to test tax incidence equivalence. For one, changes in other economic variables

³I thank Tom McGirr, Chief Economist of the Sales Tax Division, Department of Finance, Canada, for informative e-mail exchanges on the history of Canada's sales tax.

confound the measurement of tax incidence attributable to the shift in policy. Moreover, the sales tax is an *ad valorem tax*; that is, the tax is a fixed percentage of the product's value rather than a fixed amount per unit. For tax incidence equivalence to hold for *ad valorem* taxes, the underlying supply and demand conditions need to be known so that the precise tax rates needed for the theory to obtain can be set.

Controlled laboratory experiments do offer the opportunity to observe the division of economic tax incidence both when a tax is levied on the producer and when it is levied on the consumer. Kachelmeier et al. (1994) were the first to examine experimentally tax incidence equivalence. They compare the tax incidence of three different *ad valorem* consumption taxes in complex, multilevel market experiments involving consumers, retailers and wholesalers. Their results from all three tax regimes offer some support for the theory; deviations from the theory, however, indicate that those that pay the tax bear the least tax burden. One feature of their experimental procedure that may contribute to this finding is the avoidance of the word "tax" in their experiments. Instead, the taxpaying side of the market is told that a percentage charge would be added to each transaction. For this reason, Kachelmeier et al.'s paper is unable to deal with the framing issues surrounding tax incidence and dealt with in this paper.

A paper closer in spirit to my paper is the work of Kerschbamer and Kirchsteiger (2000) (hereafter K&K). To test tax incidence equivalence, they use the ultimatum game. Player 1, the proposer, is given 70 Austrian shillings to divide between herself and player 2, the responder. If the responder rejects the proposer's split, then both players earn zero and the game is over. In the standard ultimatum game, acceptance entails payoffs determined by the split of the available sum. K&K's variation consists of a 20 shilling tax levied in the case of acceptance only. Proposers pay the tax in one treatment, responders in another. K&K find that proposers' net offers were significantly higher if the tax is levied on the proposer. Stated simply, in accordance with popular perceptions, one is better off when the other side of the market pays the tax than when the tax is levied on one's own side.

I suggest that K&K's test of tax incidence equivalence is a very special case. For

one, the ultimatum game is a very particular representation of bilateral monopoly with sequential timing, a restricted strategy space and an all-or-nothing payoff structure. Second, the peculiarities of ultimatum game behavior are well known: in addition to not converging to the unique subgame-perfect-equilibrium outcome, even with repetition and monetary incentives, ultimatum-game offers and rejections have been shown to be sensitive to a plethora of contextual variables, framing and other perceptual effects, issues of fairness, intentions and property rights. In fact, K&K recognize the particularity of their results. They conjecture that if “trade takes place on competitive markets and if the characteristics of the good traded are completely specified, prices and quantities converge to the market clearing level rather quickly” (p. 733). Pursuing K&K’s conjecture, this paper offers a more general test of the theory with a larger and varying number of traders and a different laboratory institution.

Intrigued by K&K’s findings, Borck et al. (2001) conduct their own laboratory test of the theory. Using a posted-offer market (a multi-player extension of the ultimatum game in which buyers face take-it-or-leave-it posted prices), Borck et al. find that sellers are able to maintain prices above the competitive equilibrium range in both of their “SellerTax” and “BuyerTax” treatments. However, the amount by which prices differ across the two treatments is precisely the amount of the tax wedge. Thus, tax incidence equivalence is supported by their experiments. All of their posted-offer markets involve, nonetheless, exactly three buyers and three sellers.

In this paper, I test tax incidence equivalence in pit markets varying in size between 16 and 30 traders. In contrast to K&K’s and Borck et al.’s tests of the theory, buyers are not at a strategic disadvantage relative to sellers in these pit-market experiments: rather, the negotiating strategies available to buyers and sellers are symmetric and the timing of the negotiations is simultaneous.

3 Experimental Design

All experiments take place in a pit market. Price negotiations in the pit market are decentralized: buyers and sellers meet in an open space and freely negotiate with one another.⁴ At least eight pairs of buyers and sellers participate in all experiments.

Each buyer in these experiments possesses a single unit of demand while each seller has available one unit of production for sale. At the beginning of each period, each buyer receives a randomly drawn card from a set or distribution of cards; likewise, each seller receives a randomly drawn card from a different set of cards. The value (cost) of a buyer's (seller's) unit is given by the number written on his card. Buyers earn their consumer surplus, namely, their valuation minus the negotiated price in each period that they fulfill their unit demand. Sellers earn their producer surplus, namely, the negotiated price minus their cost in each period that they sell their unit of production. A subject's profit is zero for periods in which he does not trade.

All sessions consisted of 19 three-minute trading periods. During the first 8 periods, subjects participated in an ordinary pit market, with no tax or subsidy. Beginning in period 9, a single change was introduced, the payment of a 10-unit tax by either the buyers (to be referred to as "*taxb*" sessions) or the sellers ("*taxs*"), or the receipt of a 10-unit subsidy by either the buyers ("*sudb*") or the sellers ("*suds*"). Subjects remained in this same treatment for 11 periods, until the end of the experiment.⁵

⁴The pit market closely resembles the double-auction market. The two institutions differ in the way bids and asks are organized. In the pit market, traders negotiate directly with whomever they choose from the other side of the market, with participants freely exchanging bids and asks between them until a transaction price is agreed upon. In the double-auction market, bids and asks are publicly recorded on the blackboard or computer screen. Only improvements (i.e. increases in the standing bid or decreases in the standing offer) are recorded. A transaction occurs when a market participant accepts the standing bid or offer of the other side of the market.

⁵I conducted a single pilot experiment in which subjects participated in the baseline treatment followed by the two tax treatments back-to-back (7 baseline periods followed by 6 *taxb* periods followed by 6 *taxs* periods). The transaction price data reveals that buyers' initial resistance to the sharp 10-unit increase in the equilibrium price in going from *taxb* to *taxs* delayed convergence. To eliminate this

All subjects were inexperienced in this trading institution. Many market experiments suggest that eight periods may be insufficient to permit prices to stabilize. For this reason and since none of the research hypotheses relate to the baseline (i.e. pre-tax or pre-subsidy) treatment, the data from it will not be analyzed. Instead, the main purpose of the 8-period baseline treatment is to permit subjects to acquire some experience in this environment and to familiarize themselves with the trading rules, negotiation process and profit calculations, before complicating the design with the introduction of a tax or a subsidy.

One source of potential variation between experiments is the size of the market. I vary the number of traders in the market from 16 to 30 (i.e. between 8 and 15 pairs of buyers and sellers) in order to examine the robustness of the incidence equivalence theories to market size.

To compare results across experiments it is important to hold constant the equilibrium price for all market sizes. Table 1 displays the distributions of valuations and costs as a function of the number of pairs of buyers and sellers in the market.⁶ Constructing the supply and demand curves from these parameters, one observes that a two-unit competitive price tunnel exists between 31 and 33, where the supply and demand curves overlap. This competitive price tunnel is held constant across all sessions by maintaining two units at the marginally profitable (or second-to-last) demand step of 33 and two units at the marginally profitable (or second-to-last) supply step of 31. For sessions with 12 pairs of traders, two units of demand and supply exist at each step on the respective demand and supply curves. When the number of pairs of traders differs from 12, units are added to, or subtracted from, interior (intramarginal) steps on the supply and demand curves.

undesirable anchoring effect, I reduced the number of treatments in a session from three to two. This latter design has the simultaneous benefit of allowing more periods per treatment in a session in a given amount of time, thereby affording an improved chance of convergence.

⁶In approximately half of the sessions, six units were added to all valuations and costs. This shift in parameters provides a further robustness test of the results. In order to compare the results of the different sessions, I report the normalized prices, i.e. subtract six units from the prices in the shifted sessions.

mand curves, leaving the competitive price range and the number of traders outside of the competitive equilibrium unaffected.

The symmetry of the design implies that the addition of a tax increases the equilibrium price paid by the buyers by five units (half the amount of the tax) to between 36 and 38, whereas the post-tax equilibrium price received by the sellers falls by five units to between 26 and 28. This holds for both the *taxb* and *taxs* treatments and for all market sizes explored herein. Furthermore, the post-tax competitive equilibrium quantity falls by two units for both tax treatments and all market sizes.

The addition of the subsidy to the baseline treatment has a similar effect: it increases the equilibrium price received by the sellers to between 36 and 38, and decreases the price paid by the buyers to between 26 to 28. The payment of the subsidy permits the trading of two additional units in both the *sudb* and *suds* treatments compared to the baseline treatment for all market sizes.

[insert Table 1 here]

The experimental design, an adaptation of Bergstrom and Miller (1997, p. 107), includes several appealing features. First of all, the supply and demand curves intersect not at a unique point, but share a vertical overlap. A horizontal overlap would mean that marginal traders earn zero profit from trades at the unique competitive price. The absence of economic profit from trading is likely to prevent the exchange of these units and therefore, at the very least, the convergence to the competitive quantity. Some researchers pay subjects a commission for each unit they trade; however, this effectively shifts the demand (supply) curve upward (downward) by the amount of the commission. The end result is that the vertical overlap of the supply and demand curves is reestablished of height twice the amount of the commission.

Second, to eliminate possible psychological effects associated with round numbers, I avoid the use of round numbers in choosing buyers' valuations, sellers' costs and, most importantly, the competitive price range in all treatments.

4 Experimental Procedures

4.1 Procedures

Upon arrival subjects were seated and divided into two equal groups of buyers and sellers. In the event that an odd number of subjects showed up, I solicited a helper from among the subjects, offering to pay him the average payment earned by subjects in the experiment. The instructions were distributed along with a personal record sheet allowing the subjects to keep track of their own transactions and profits.⁷ The subjects read the instructions and then I read them aloud answering any questions about the experiment. The buyers' valuation and sellers' cost cards were distributed for the first three-minute period.⁸ Subjects were told the period had begun at which point they rose from their seats to begin negotiating in the center of the classroom. Whenever a buyer and a seller reached an agreed upon price, together they approached one of the transactions booths. They submitted their cards, face down, to one of the experimenters and reported the transaction price. They then returned to their seats, filled in their record sheets and awaited the commencement of the next round. Meanwhile, the experimenter recorded the transaction and passed the transaction sheet onto another experimenter who stood at the blackboard and wrote down the transaction price for all subjects to observe. Finally, the transaction sheet was forwarded to another experimenter seated in front of a notebook computer who entered the transaction in an Excel spreadsheet.

At the end of the eighth period, subjects were instructed verbally of the introduction of a tax or a subsidy to begin in period 9 until the final period 19. They were informed of the meaning of the change for their profit calculations. Specifically, in the *taxb* treatment, for example, they were told:

⁷A much more detailed discussion of the experimental design, procedures, including instructions, and features of the data typical in pit markets can be found in Ruffle (2001).

⁸The card each subject received was his own private information. The distributions of cards were unknown, although buyers (sellers) learned about the distribution of buyers' valuations (sellers' costs) through receiving a new card each period.

Now we begin the second part of the experiment which involves a single change: on each transaction, the buyer must pay a tax in the amount of 10 units. This amount, 10 units, is simply subtracted from the buyer's profit on each transaction that occurs. If the buyer does not purchase a unit in a given period, he does not pay the tax. The sellers' profits are unaffected by the tax.

Note that buyers in this treatment, as in all other treatments, were *not* explicitly told that the addition of the tax is equivalent to subtracting 10 units from the values written on their cards or that the tax amounts to a downward shift in their unit demand by 10 units. The intentional omission of this equivalence and the framing of the downward shift in demand as a tax provides full scope for the observance of psychological effects possibly associated with statutory tax incidence or statutory subsidy payment such as the sense of duty or moral obligation to bear the burden of the tax or the sense of right or moral entitlement to reap the benefit of the receipt of the subsidy. A second possible effect following the introduction of the tax (subsidy) suggests that subjects perceive the bargaining position of the side of the market that pays the tax (receives the subsidy) as weakened (strengthened) compared to the baseline treatment. No matter the reasoning, both of these perceptions push prices apart in the same direction in the *taxb* and *taxs* sessions as well as the *sudb* and *suds* sessions. Convergence to the competitive equilibrium in spite of the latitude yielded to such framing effects would attest to the strength of the competitive outcome and the tax and subsidy incidence equivalence theorems.

During each period, subjects were free to negotiate with whomever they chose from the other side of the market. Also, I did not prevent subjects from discussing negotiating strategies or colluding. My belief was that even if subjects attempted to collude, competitive forces would foil such attempts. This belief turned out to be supported by the numerous failed attempts at collusion observed in these experiments, with one important exception discussed in detail in section 5.3.

4.2 Subjects and Payments

An effort was made to vary the number of traders in different sessions while, at the same time, balancing the composition of sessions across treatments in terms of the numbers of traders, in particular, as well as their fields of study.⁹ Thus, for instance, the six sessions in which the 10-unit tax was imposed on the buyers (*taxb*) consisted of 8, 10, 11, 12, 13, and 14 subject pairs compared to 8, 9, 10, 12, 13 and 15 subject pairs in the *taxs* sessions. Similarly, the six *sudb* sessions consisted of 9, 10, 11, 13, 13, 15 subject pairs compared to 9, 10, 11, 12, 13, 15 subject pairs in the six *suds* sessions.

Tables 2a and 2b list the 24 sessions conducted. The first and third columns of each table indicate the names of the experiments. To illustrate my labeling system, “tax8b5” indicates that this session was a tax (“tax”) experiment with eight (“8”) pairs of buyers and sellers. The tax was imposed on the buyers (“b”) and this was the fifth (“5”) tax experiment conducted in the series of 12. Within each treatment, the experiments are listed in ascending order to the number of pairs of buyers and sellers in the market.

[insert Table 2a here]

[insert Table 2b here]

The effort to balance sessions across treatments in terms of the subjects’ academic backgrounds can be seen in the second and fourth columns of the tables, labeled “Pool”. Thus, for instance, three *taxb* sessions and three *taxs* sessions were conducted entirely on first-year economics majors during a two-hour class tutorial (econ 1). Similarly, one *sudb* and one *suds* session consisted entirely of first-year economists. For the remainder of the tax and subsidy sessions, we recruited from the overall student body (mix), targeting in particular, 2nd and 3rd year economics students (econ 2,3), as well as engineers (eng) and business majors (bus), again with an eye to balancing the subject pool composition across treatments.

⁹Minor differences in the composition of the treatments remain due to the usual unpredictability of no-shows, despite telephone confirmation the day before each experiment.

In total 552 subjects participated in one of 24 tax or subsidy experiments at Ben-Gurion University. The seemingly large number of sessions per treatment reflects an effort to allow a better chance at rejecting the theory through increased statistical power.

Subjects' earnings were the sole function of their realized profits from trading. That is, no showup fee was paid: my conjecture is that the absence of a security profit encourages subjects to trade even for a single unit of profit. Profit maximization is a necessary condition for convergence to the competitive equilibrium quantity.

The new Israeli shekel (NIS) to experimental currency exchange rate was set at 1:3 and 1:4 for the tax and subsidy experiments, respectively, for the eight sessions conducted on first-year economics majors during their two-hour tutorials. When I turned to general recruiting, the respective exchange rates were increased to 1:2 and 1:3.5 (to compensate students for showing up specifically to participate in the experiment).

The experiments took approximately one hour and 40 minutes. The average payment for the eight sessions conducted during the compulsory classroom tutorial was 44.75 NIS (s.d. 12.7), compared to 63.81 NIS (s.d. 13.4) for other sessions in which subjects showed up especially for the experiment.¹⁰

5 Results

5.1 Tax Experiments

5.1.1 Convergence and Tax Incidence Equivalence

Figure 2 displays the median transaction prices by period for each of the twelve tax sessions. Table 3 shows the median transaction price data for these tax sessions. The first price indicated in bold font in each session indicates the initial period in which the

¹⁰At the time these experiments were conducted \$1 US equaled approximately 4 NIS. Almost all subjects earned considerably more than their outside opportunity cost of 18 NIS per hour, the minimum wage in Israel.

median price entered, and remained within, the competitive price tunnel. Recall that the 10-unit tax wedge results in competitive equilibrium price ranges of 36-38 for the *taxs* sessions and 26-28 for the *taxb* sessions since the latter range does not include the payment of the tax. Visual inspection of the data points to the paper's first main result.

[insert Figure 2 here]

[insert Table 3 here]

Result 1 *Prices converge to the post-tax competitive equilibrium in all 12 tax sessions.*

Using the median transaction price as our benchmark, all 12 tax sessions converge to the competitive price range within the 11 periods designated to the tax treatment. Price convergence is particularly uniform in the *taxb* sessions. All six sessions show the same median price of 27, the midpoint of the competitive tunnel, in the final period. Moreover, convergence to the competitive price range is notably rapid: in ascending order according to the size of the market, the six *taxb* sessions required 6, 1, 3, 2, 3 and 1 period to converge to, and remain within, the competitive price range (see Table 3).

The transactions graph for session *tax13b10*, typical in its pricing dynamics is displayed in Figure 3. Each dot represents a transaction, with the height of the dot indicating the negotiated price. Vertical lines separate between the periods. The dots in each period are arranged according to the order in which transactions occurred. The average price and quantity exchanged are displayed beneath the x-axis for all 19 periods.

[insert Figure 3 here]

Session *tax13b10* begins with moderate price dispersion in the first few periods of the baseline treatment. (Many sessions display considerably higher price variance in the initial periods.) Price variance generally decreases over time. By period 5, both the average price and quantity have converged to the competitive outcome. The introduction of the

tax in period 9 lowers prices by less than the full 5 units initially. However, competitive forces produce a gradual downward drift in transaction prices until prices enter the competitive range for good by period 3 of the treatment by both the median and mean price measures. Notice that the price variance at the beginning of the tax treatment in this session (as in every other tax and subsidy session) is relatively low, considerably lower than at the outset of the baseline treatment. The probable explanation is that by period 9 traders have accumulated both experience in this market environment and knowledge about their own valuation or cost distribution. As a result, uncertainty and price experimentation are reduced, and “mistakes” (in the form of buying a unit at a price considerably higher than the other publicly displayed transaction prices or selling a unit at a price substantially lower than other displayed prices) are eliminated.

Price convergence to the competitive range occurs in all six *taxs* sessions as well, albeit more slowly than the *taxb* sessions. Sorted by ascending order according to the size of the market, the six *taxs* sessions required 5, 5, 11, 4, 6, and 5 periods before settling into the competitive price range for the duration. The transactions graph for session *tax8s7* appears in Figure 4. This session displays the typical pattern of declining price variance over time in both the baseline and tax treatments. What distinguishes this session is the absence of convergence after eight periods in the baseline treatment. In period 9, prices over-react to the imposition of the tax on the sellers, but gradually fall to enter the competitive tunnel for the duration by period 15 according to the mean price, period 13 according to the median price.

[insert Figure 4 here]

Examining the two treatments together, we see that 11 of the 12 sessions converge to, and remain within, the competitive range by period 6 of the tax treatment (period 15 overall). Session *tax10s11*, discussed in Ruffle (2001) as an example of exceptionally slow price, but not quantity, convergence, required all 11 periods before the median price reached 36.

That prices converge to the competitive outcome in all 12 sessions does not rule out the possibility that they nonetheless differ across treatments. Given the competitive price range, prices in the two treatments may differ by as much as two units and still converge to the competitive equilibrium. Observation 2 suggests some variation in transaction prices as a function of the size of the market.

Observation 2 *Markets consisting of smaller numbers of traders and a small competitive equilibrium quantity (six units or less) leave room for relative negotiating strengths to influence transaction prices. The result is that minor price discrepancies from the midpoint of the competitive range may persist.*

We previously noted the rapid and uniform price convergence of all six *taxb* sessions to 27, the midpoint of the competitive tunnel in this treatment. Prices in the six *taxs* sessions, however, display greater heterogeneity. Differences in prices across these six sessions can be neatly characterized as a function of the market size. Close inspection of prices from the terminal period of the *taxs* sessions reveals that *both* the median and mean prices are lowest in the three smallest sessions, *tax8s7*, *tax9s1* and *tax10s11*. The period 19 median price in these three sessions is 36, one or one-and-a-half units lower than the three larger sessions in this treatment. The first two rows of Table 4 report summary statistics for the period 19 prices in the six *taxb* sessions, and the *taxs* sessions divided according to market size.

[insert Table 4 here]

The data reveal that when the tax falls on the buyers, they are able to bring down the price by the full 5 units prescribed by the midpoint of the competitive price range. However, when the sellers pay the tax and there are relatively few traders, buyers are able to resist the full five-unit price increase. This resistance may be seen in the form of slower convergence to the competitive range from below, as previously discussed, and convergence to the lower bound of the competitive range.

The observation that buyers tend to be stronger in experimental market negotiations is not new. Smith and Williams (1982, p. 115) provide a plausible explanation for disproportionate buyer negotiating strength: students have more lifetime experience as buyers and therefore are more adept at negotiating in the role of buyers than they are as sellers.¹¹

Hence, we observe that when the market is relatively thin (a volume of six units or less at the competitive equilibrium, according to the results of the six *taxs* experiments), buyers are able to exert their force and prevent prices from rising by the full 5 units after the introduction of the tax.

By contrast, all but one session with a competitive quantity of seven or more units (11 or more pairs of traders) converge *precisely* to the midpoint of competitive tunnel. (Session *tax13s3* converges to a median price of 37.5, a mere half a unit above the midpoint of the competitive tunnel, and a mean price of 37.13.)

These findings suggest that the degree of support for tax incidence equivalence depends on the extent to which the market is characterized by competitive conditions.

Result 3 *Tax incidence equivalence theory is upheld in these experiments for sufficiently large markets. When the competitive quantity exceeds six units (alternatively, when more than six pairs of market participants can trade profitably at the competitive price), prices in sessions in which the tax is imposed on the buyer are not (statistically) different from prices in sessions in which the tax is imposed on the seller.*

Taking into account the nature of the observed price discrepancies in the *taxs* sessions, we may first separate the six *taxs* sessions into the three large, “*taxs* (3 large)”, and the three small, “*taxs* (3 small)”, markets and compare their prices with each other and with the prices from the six *taxb* sessions.¹² The Kruskal-Wallis test checks for

¹¹One might counter that the majority of this experience is in a posted-price environment. To this I would respond that the acts of comparing products and prices and shopping around for the best value are psychologically and motivationally similar to negotiating a better price.

¹²Ten units (the transfer of the tax) are added to all prices in the *taxb* sessions in order to compare them with prices from the *taxs* sessions.

differences in the central tendencies of the price distributions among these three different sets of markets. We are able to reject the null hypothesis that the period 19 average price distributions of the *taxb*, *taxs* (3 large) and *taxs* (3 small) are the same ($\chi^2(2) = 6.29$, $p=.043$). The result is even stronger if we compare the period 19 median price distributions ($\chi^2(2) = 9.50$, $p=.009$).

The source of the difference in prices between these three groups is clear. When the period 19 average transaction prices from all 12 tax sessions are ranked from lowest to highest, the mean rank (upon which the Kruskal-Wallis test is based) from the *taxs* (3 small) sessions is significantly lower than that of the other two groups. See the third row of Table 4 for the details. Put differently, despite price convergence in the *taxs* (3 small) sessions, prices in these three sessions remain, on average, one unit below prices from *taxs* (3 large) and *taxb* (inclusive of the tax) sessions. Comparing period 19 average prices in the *taxs* (3 large) sessions with those in the *taxb* sessions, we find no difference between them (Kolmogorov-Smirnov $z=.471$, $p=.979$). The identical result obtains if instead we compare the period 19 median prices (Kolmogorov-Smirnov $z=.471$, $p=.979$).¹³

5.1.2 Direction of Convergence and Initial Framing Effects

The above analysis is based entirely on prices from the terminal period. Except for relatively small markets in which the tax is imposed on the seller, we found no evidence for a tax framing effect in the form of either the perception that the bargaining power of the side of the market upon which the tax was levied was weakened or a sense of moral obligation to bear a disproportionate share of the tax burden on the part of this side of the market. Instead, prices converged uniformly to the midpoint of the competitive

¹³Note that for all of these non-parametric tests, we treat the average or median price from the terminal period of a session as a single observation. We could instead compare the distributions of all period 19 transaction prices. The results for all reported Kolmogorov-Smirnov and Kruskal-Wallis tests remain qualitatively the same. However, the results of tests using the entire period 19 price distribution are dubious since, due to the open negotiation process and the public display of transaction prices, price observations within a session are not independent but rather highly correlated. Precisely for this reason, a relatively large number of sessions, six, was conducted in each of the four treatments.

price range regardless of who pays the tax.

It may nonetheless be that this sense of weakened bargaining power or moral obligation is present immediately after the introduction of the tax, and is “squeezed out” over time by competitive pressures. In fact, the protests and charges of unfairness at the beginning of period 9 of those required to pay the tax would suggest that such initial reactions exist.

Short-term, temporary framing effects may have substantial economic consequences. For example, the short term for infrequently renegotiated wage contracts may be several years. If management and labor unions behave as if statutory and economic incidence are related, then out-of-equilibrium wages will persist for years in countries in which the payroll tax is levied upon only one side of the market.

Allow me to illustrate the psychological mechanism by which such framing effects operate in these experiments. Consider the treatment in which the tax is levied on the buyers. Associating statutory incidence with economic incidence, the buyers will feel either obliged to bear most of the tax burden or unable to avoid it. Thus, they will be willing to accept relatively high prices, prices above the competitive equilibrium, in the initial periods of the *taxb* treatment. This leads to the testable hypothesis that prices in the *taxb* treatment converge to the competitive equilibrium from above. Similarly, if the sellers are prepared to shoulder a disproportionate share of the tax burden when they pay the tax, the prices in the *taxs* treatment should initially be below the post-tax competitive equilibrium and converge from below.

To test the convergence properties of the two tax treatments, I employ the following linear regression developed in Noussair et al. (1995, 1997) and referred to as the Ashenfelter-El-Gamal model.

$$P_{it} = B_{11}D_1(1/t) + B_{12}D_2(1/t) + \dots + B_{1k}D_k(1/t) + B_2\left(\frac{t-1}{t}\right) + u_{it} \quad (1)$$

The term D_i is a dummy variable that assumes a value of 1 for experiment i , $i = 1, \dots, k$, and 0 otherwise. The term P_{it} represents a transaction price in period t in experiment i . Notice that when $t = 1$, the price in experiment i equals B_{1i} . The term B_{1i} can thus

be thought of as the initial price origin. Its impact decays over time, as indicated by the term $1/t$. By contrast, the impact of B_2 increases over time, as indicated by the expression $\frac{t-1}{t}$. The term B_2 can thus be thought of as the common price asymptote. The random error term u_{it} is distributed normally with mean zero.

As Noussair et al. point out, this dynamic statistical model of pricing behavior has several desirable properties. First of all, the weights attached to the initial price estimate and the price convergence estimate sum to 1. In period $t = 1$, the price in experiment i equals B_{1i} . As the experiment progresses, the weight accorded to the initial price B_{1i} decreases and shifts to the convergence asymptote, B_2 . In the limit, as t becomes large, the price converges to B_2 since the weight attached to the initial price decreases to zero. Note that the subscript i on the B_{1i} term indicates that the model allows for different initial prices for each experiment. However, in accordance with the theory of competitive equilibrium, all experiments are assumed to converge to the common price asymptote B_2 .

Noussair et al. employ the Ashenfelter-El-Gamal regression model to test whether various price and quantity time series converge to one theoretical benchmark or another in more complex trading environments than the one examined herein. Price convergence in our comparatively simple environment has already been firmly established. Instead, we are interested in the direction of convergence, given by the relationship between B_{1i} and B_2 . If the fitted regression line reveals that $B_{1i} > B_2$, then the regression line is negatively sloped, indicating that experiment i converges to B_2 from above. On the other hand, if $B_{1i} < B_2$, then experiment i converges from below. Table 5a and the first row of Table 5b report the results from the separate *taxb* and *taxs* regressions and lead to the next result.

[insert Tables 5a and 5b here]

Result 4 *Individual experiments converge to the competitive price range from within the competitive price range, from above and from below. Neither the taxb nor the taxs experiments display any systematic direction of convergence. The implication is that from*

the initial period of the tax treatments, competitive market forces dominate any perception of weakened bargaining power, and any sense duty or moral obligation associated with the payment of the tax.

Table 5a reveals that the initial price of two of the six *taxb* experiments (*tax10b6* and *tax13b10*) does not differ significantly from the price asymptote; that is, these experiments converge from within the competitive price range. Three experiments converge from above (each significant at the 5% level), as predicted by the initial framing effect hypothesis. And, contrary to the framing effect, one experiment converges from below (significant at the 1% level).

The *taxs* experiments display no clear direction of convergence at all. Whether we examine the overall *taxs* regression (row 1 of Table 5b) or the separate regressions for the three large *taxs* experiments (row 2) and the three small *taxs* experiments (row 3), two of the six experiments show no direction of convergence, two experiments converge from below and two experiments from above.

On a different note, the separate estimates of B_2 for the *taxb* and *taxs* treatments reinforce two previous results. That both B_2 estimates (27.25 (not including the tax) and 36.23) fall within the competitive price range confirms Result 1. Nonetheless, if we compare the tax-inclusive price asymptotes from the *taxb* and *taxs* regressions a one-unit price discrepancy exists. Observation 2 and the discussion that follows it suggest that the smaller market sizes in the *taxs* treatment may account for the comparatively low convergence price of 36.23. To verify this earlier result, I report separate regressions for the three large *taxs* experiments (row 2 of Table 5b) and the three small ones (row 3 of Table 5b). These separate B_2 estimates reveal a 1.22 unit price gap between the large and the small sessions; whereas the *taxb* and *taxs* (3 large) sessions converge to prices that differ by a mere 0.59 units.

5.2 Subsidy Experiments

5.2.1 Convergence and Subsidy Incidence Equivalence

Immediately following the introduction of the subsidy there is considerably more price variance in the subsidy experiments than that observed in the comparable tax experiments. One explanation is that subjects are more familiar with the concept of a tax and therefore respond more appropriately and more uniformly than they do to the less familiar subsidy. In spite of initial price variance, convergence prevails and more quickly than in the tax experiments.

Figure 5 displays the transactions graph for a session, *sud10b3*, with high price variance at the beginning of both the baseline and subsidy treatments. That notwithstanding, already by the fourth period of the subsidy treatment, all but at most one trade occurs within the competitive price range in each period.

[insert Figure 5 here]

Convergence to the competitive price range is even more rapid in the same size session in the *suds* treatment, *sud10s2*. Figure 6 displays the transactions graph for this session.

[insert Figure 6 here]

Rapid convergence characterizes the subsidy experiments more generally, as stated in the next result.

Result 5 *Prices in 11 out of the 12 subsidy sessions converge to the competitive equilibrium. Convergence occurs more quickly than in the comparable tax experiments.*

[insert Figure 7 here]

[insert Table 6 here]

Figure 7 and Table 6 demonstrate this result. Both display the median transaction price by period for each of the twelve subsidy sessions. With the exception of one session, *sud9b7*, to be discussed in detail below, all sessions converge to the competitive price range no later than period 13 (only the fifth period of the subsidy treatment). What is more, five of the sessions converge to, and remain within, the competitive price tunnel in the first period of the subsidy treatment. Three additional sessions do so in the second period of the treatment. In short, convergence happens even more quickly in the subsidy experiments than the tax experiments. The similar composition of the four treatments both in terms of the size of the market and the identity of the subjects adds credibility to this observation.

Nonetheless, theoretically, the two-unit competitive price tunnel permits systematic price variation between the treatments. Result 6 rules this out.

Result 6 *These experiments offer unequivocal support for subsidy incidence equivalence theory. Prices in sessions in which the subsidy is given to the buyers are not (statistically) different from prices in sessions in which the subsidy is given to the seller.*

With the exclusion of the single outlier session, *sud9b7*, the distributions of all period 19 transaction prices in the *sudb* and *suds* treatments look very similar: for instance, the period 19 median price is 37.0 for both the *sudb* (n=60) and *suds* (n=67) treatments. Moreover, the period 19 mean prices are nearly identical, 36.92 for *sudb* and 36.69 for *suds*. Comparing the period 19 mean prices at the session level, we cannot reject the null hypothesis that the sample distributions are drawn from the same population distribution (Kolmogorov-Smirnov $z=.826$, $p=.503$). The exact same result obtains if instead we use the period 19 median prices (Kolmogorov-Smirnov $z=.826$, $p=.503$).

In light of the findings of the tax experiments, the absence of a difference between the two subsidy treatments should come as no surprise. We saw in the *taxs* sessions that if the competitive quantity is less or equal to six units, prices do not necessarily converge to the same price as they do in larger sessions. However, the competitive quantity in the subsidy treatments increases by two units compared to the baseline treatment, and

is four units larger than the competitive quantity of a tax treatment with an equivalent number of traders. As a result, the smallest competitive quantity in the subsidy sessions conducted was eight units. And indeed we do not find systematic price differences between the sessions as a function of market size. These markets are large enough so that competitive forces bring about strong convergence to the same price.

5.2.2 Direction of Convergence and Initial Framing Effects

The convergence to a common price in both subsidy treatments does not rule out the possibility of initial price discrepancies. More concretely, a framing effect in which individuals associate statutory incidence with economic incidence in the subsidy treatments would suggest that the receipt of a subsidy is treated as a right or moral entitlement to be the primary beneficiary of the subsidy. Alternatively, the receipt of a subsidy may be perceived by either side of the market as strengthening the bargaining power of the side of the market that receives the subsidy. In either case, the implication is that the direction of convergence should differ by treatment. Prices in the *sudb* experiments should converge from below; that is, prices start out below the equilibrium price range since buyers feel empowered or entitled to benefit disproportionately from the subsidy paid to them, and sellers acquiesce with low prices. By the same logic, prices in the *suds* experiments can be expected to converge from above.

To examine these hypotheses, I estimate an Ashenfelter-El-Gamal regression equation for each of the treatments. In the *sudb* regression, I exclude the outlier session since it does not converge and therefore follows a different data generating process. In the next subsection, a separate regression will be estimated for this session to examine its convergence properties.

[insert Table 7a here]

Table 7a reports the regression results from the *sudb* experiments. Two of the five *sudb* experiments converge from within the competitive tunnel and show no significant

direction of convergence. Only one experiment converges from below, as predicted by the framing effect hypothesis; whereas, counter to this hypothesis, two experiments converge from above.

The framing effect hypothesis performs more favorably in the *suds* experiments. The first row of Table 7b indicates that, in accordance with this hypothesis, four of the six experiments converge from above. Interestingly, the difference between the initial prices of these four experiments and the common convergence price (“the convergence distance”) shrinks as the market size increases. That is, the extent to which a sense of empowerment or entitlement exists in this treatment is inversely related to the size of the market. Intensifying competitive forces appears to dampen or drive out altogether such psychological sentiments even at the earliest stages of trade.

Contrary to the framing effect hypothesis, one *suds* experiment converges from above, and one displays no direction of convergence.

[insert Table 7b here]

Result 7 summarizes the convergence tendencies of the two treatments.

Result 7 *Evidence for an initial framing effect in the form of perceived empowerment or entitlement associated with the statutory payment of a subsidy is limited. The *sudb* experiments display no systematic direction of convergence, while 4/6 *suds* experiments converge to the competitive price range from above, the direction predicted by the framing effect hypothesis.*

5.3 Collusion

As noted in section 4.1, subjects were permitted to discuss and coordinate negotiating and pricing strategies between periods. In fact, non-binding, verbal, collusive agreements were reached in numerous sessions. Typically, between periods, a trader, say a buyer, proposed aloud to other buyers to remain in their seats and not to buy above

some stated price. After some discussion, all of the buyers would agree. At the opening of the next period, buyers indeed remained seated. Not surprisingly, sellers did not easily succumb to the buyers' price demands. Seconds ticked away. The experimenter announced that one minute remained in the period. At this point, those buyers with the highest valuations became nervous. Unwilling to forego a large profit, these high-valuation buyers were almost inevitably the first to rise from their seats in search of a seller willing to trade at a price established in the previous period, a price above the collusive agreement. At this point, the collusion unravelled: other, perhaps medium-valuation, buyers observed the defector and decided that they too preferred to cash in a positive profit before the period terminated. By the next period, there was usually little, if any, discussion of a collusive price. The following observation summarizes these findings and notes the exception to the rule.

Observation 8 *Efforts to collude are numerous. However, in the face of the profit maximization motive and competition, collusion typically breaks down in these markets soon after it is attempted – often in the same period. In the single instance in which collusion succeeded in maintaining out-of-equilibrium prices, the lost sales made it no more profitable.*

In the baseline, tax and subsidy treatments, numerous attempts at price collusion were observed. All such attempts were soon abandoned, often within the same period in which they were initiated. Session *sud9b7* provides a striking exception to this rule. It is deserving of closer investigation. Figure 8 displays the transactions graph for this session.

[insert Figure 8]

The first noteworthy observation about the results of this session is how well-behaved prices are in the baseline treatment: already in period 1 the median price (period 2 by the mean price measure) falls within the competitive price tunnel where it remains for all

eight periods of the baseline treatment. Prices display the typical feature of decreasing variance over time, as prices converge strongly to the midpoint of the competitive range. However, precisely between periods 8 and 9, after the announcement of the 10-unit subsidy to the buyers, the sellers confer and agree not sell below 40 (that is, two units above the top of the competitive price range). While some transactions do occur below 40 in period 9, a redoubled effort in period 10 succeeds in maintaining the prices of 6/7 trades above 40. This pattern in which all but one or two trades occur at a price of 40 or more continues through period 15. Beginning in period 16 there are signs that the collusion is beginning to break down as prices drift downwards.

However, as the regression reported in row 2 of Table 7a points out, this downward drift in prices is not highly significant. Namely, the initial price term of 40.24 differs by a mere 0.73 units from the convergence asymptote of 39.51 ($z=1.066$, $p\text{-value}=.145$). Similarly, according to Noussair et al. (1995, p. 474), data converge weakly if the convergence asymptote, B_2 , is closer to the model's prediction than to the initial price term, B_{1i} . By this definition, I reject weak convergence for session *sud9b7* since 39.51 is closer to 40.24 than to 38 (the upper bound of the competitive price range).

In summary, this session does not converge, even weakly, to the competitive price. What is more, prices show no significant downward trend over time. Finally, the fact that prices are *above* the competitive equilibrium contradicts the hypothesized framing effect in this *sudb* session.

What distinguishes this session from all of the others in which failed attempts at collusion were soon aborted? The answer lies largely in the familiarity of subjects in this experiment with one another. All nine of the colluding sellers (and 17/18 subjects in the experiment) were engineering students, six of them classmates in the second year. The fact that all of these sellers knew one another enabled them to trust each other enough to permit cooperation. Furthermore, the threat of continued interaction with one another after the experiment imposes a cost to those who defect from the collusive agreement. The relatively small size of the market further facilitated successful collusion.

A natural question to ask is, did the sellers profit from their collusion? There is no clear evidence that they did. The average seller profit from the ten subsidy sessions in which subjects were recruited ranges from 58.4 NIS to 64.4 NIS. At 61.4 NIS, the profits of sellers in this session rank fifth out of ten. Thus, the higher prices achieved from the collusion appears roughly to compensate the sellers for the foregone sales. Buyers' profits, however, are severely hurt. With average earnings of 51.0 NIS, buyers in this session earned substantially less than buyers in the other nine recruited, subsidy sessions who earned on average between 56.0 NIS and 69.3 NIS. On the whole, average subject payments from this session were the lowest among the ten recruited, subsidy sessions.

6 Discussion and Final Remarks

6.1 Related Literature

Vernon Smith's (1962) classic paper tests the robustness of the competitive equilibrium outcome to a host of variables, including the elasticities of supply and demand, the number of units possessed by each trader, an unannounced increase in demand, subject experience, and the number of intramarginal traders near the equilibrium price. In the one to six periods Smith allows for each of these tests, prices approach, or even converge to, the competitive outcome with respect to all of these variables, with the exception of an extreme design with a perfectly elastic supply curve. Smith's paper thereby establishes the empirical relevance of the competitive equilibrium concept.

Since Smith's pathbreaking work, many subsequent experimental studies have examined the competitive equilibrium's robustness to various changes in design. (Consult Davis and Holt (1993) for a survey.) Two studies are particularly pertinent to the current paper. Williams and Smith (1984) employ a design in which the competitive benchmark price alternates between a low and a high price, the high price being more than 50% above the low price. This change in the equilibrium price is achieved by an unannounced, simultaneous shift of the supply and demand curves by equivalent amounts so that the

competitive quantity is held constant. Their results show that prices track well the alternating equilibrium price from period to period. Davis, Harrison and Williams (1991) investigate prices in designs with either unannounced demand or supply cycles, with the other curve stationary. According to their results, prices track the shifting competitive equilibrium price quite well in the double-auction market experiments, less well in the posted-offer market.

6.2 Summary

The results from the experiments reported herein can be seen as a further extension of the bounds of applicability of the competitive equilibrium in market institutions. Prices and quantities in these experiments have been shown to converge to the competitive outcome following a one-time, announced shift in either the supply or demand curve, framed as either the collection of a tax on the buyers or the sellers or the disbursement of a subsidy to buyers or sellers. The precise shifts of the supply and demand curves are designed to test the tax and subsidy incidence equivalence theorems. The framing of the shift as a tax or a subsidy permits subjects to interpret the imposition of a tax on them as a moral obligation or duty or as having weakened their bargaining power and the payment of a subsidy to them as a moral entitlement or right or as having strengthened their bargaining power. In spite of the potential scope for such framing effects, these markets show little evidence of them immediately following the introduction of the tax or subsidy. However, for smaller markets (in which between four and six pairs of market participants can trade profitably at the competitive equilibrium), the relative bargaining strengths of traders may play a residual role, possibly resulting in minor price discrepancies between different markets. The direction of these price discrepancies is consistent with the framing effect hypothesis. On the other hand, when the market size is sufficiently large (namely, the ability of seven or more pairs of market participants to trade profitably at the competitive price), I find strong support for both incidence equivalence theorems. Thus, in accordance with the theories, the burden of a unit tax is shown to be independent of who pays the tax; by the same token, the distribution of

the benefit of a unit subsidy is the same regardless of who actually receives it.

These findings, along with the results of Kerschbamer and Kirchsteiger, suggest that the domain of policy debate concerning who should pay the tax or receive the subsidy should first be focused on determining to what extent the market under consideration is characterized by competitive forces. If it is determined that prices in the industry under consideration are determined by competitive forces, the tax (subsidy) should be administered in a way that minimizes collection and compliance (disbursement) costs and accords with other local or national priorities. For instance, so as not to impede exports of a particular domestically produced good, a consumption rather than a production tax should be preferred; alternatively, the production rather than the consumption of the same good should be subsidized. To the extent that the market is not perfectly competitive, discussions concerning who should pay or who should receive take on increasing relevance.

References

- [1] Associated Press. "Tobacco Subsidy Survives," *Washington Post*, July 25, 1997, A16.
- [2] Bergstrom, Theodore C. and John H. Miller. *Experiments with Economic Principles*, MacGraw Hill: New York, 1997.
- [3] Borck, Rainald, Engelmann, Dirk, Wieland Müller and Hans-Theo Normann (2001) "Tax Liability-Side Equivalence in Experimental Posted-Offer Markets," forthcoming in *Southern Economic Journal*.
- [4] Davis, Douglas D., Glenn W. Harrison and Arlington W. Williams (1993) "Convergence to Nonstationary Competitive Equilibria: An Experimental Analysis," *Journal of Economic Behavior and Organization*, 22:3, 305-326.

- [5] Davis, Douglas D. and Charles A. Holt. *Experimental Economics*, Princeton University Press: Princeton, 1993.
- [6] Kachelmeier, Steven J., Stephen T. Limberg and Michael S. Schadewald (1994) "Experimental Evidence of Market Reactions to New Consumption Taxes," *Contemporary Accounting Research*, 10:2, 505-545.
- [7] Kerschbamer, Rudolph and Georg Kirchsteiger (2000) "Theoretically robust but empirically invalid? An experimental investigation into tax equivalence," *Economic Theory*, 16, 719-734.
- [8] Krugman, Paul. "Gasoline Tax Follies," *New York Times*, March 15, 2000.
- [9] Mankiw, N. Gregory. *Principles of Economics*, Harcourt Brace: Fort Worth, 1997.
- [10] Noussair, Charles N., Charles R. Plott and Raymond G. Riezman (1995) "An Experimental Investigation of the Patterns of International Trade," *American Economic Review*, 85:3, 462-491.
- [11] Noussair, Charles N., Charles R. Plott and Raymond G. Riezman (1997) "The Principles of Exchange Rate Determination in an International Finance Experiment," *Journal of Political Economy*, 105:4, 822-861.
- [12] Rosen, Harvey S. *Public Finance*, McGraw Hill: Boston, 5th ed., 1999.
- [13] Ruffle, Bradley J. (2001) "Competitive Equilibrium and Classroom Pit Markets for the Early 21st Century," unpublished manuscript, Ben-Gurion University.
- [14] Smith, Vernon L. (1962) "An experimental study of competitive market behavior," *Journal of Political Economy*, 70:2, 111-137.
- [15] Smith, Vernon L. and Arlington W. Williams (1982) "The Effect of Rent Asymmetries in Experimental Auction Markets," *Journal of Economic Behavior and Organization*, 3:1, 99-116.

- [16] Stiglitz, Joseph E. *Economics of the Public Sector*, W. W. Norton: New York, 2nd ed., 1988.
- [17] Williams, Arlington W. and Vernon L. Smith (1984) "Cyclical Double-Auction Markets with and without Speculators," *Journal of Business*, 57:1, 1-33.
- [18] Wilson, Michael H. (1987) "Canada Considers a Business Transfer Tax," in *The Consumption Tax*, C.E. Walker and M. A. Bloomfield (eds.), Ballinger: Cambridge, MA, 41-48.




Figure 1a: Supply and demand parameters in the absence of a subsidy.




Figure 1b: The payment of a 10-unit subsidy to the sellers shifts outward their supply curve by this amount.




Figure 1c: The payment of a 10-unit subsidy to the buyers shifts outward their demand curve by this amount.

Distributions of Costs and Valuations

Cards		Number of pairs of sellers and buyers							
Sellers' Costs	Buyers' Valuations	8	9	10	11	12	13	14	15
11	53	1	2	2	2	2	2	2	2
16	48	1	1	1	2	2	2	3	3
21	43	1	1	1	1	2	3	3	3
26	38	1	1	2	2	2	2	2	3
31	33	2	2	2	2	2	2	2	2
36	28	2	2	2	2	2	2	2	2

Table 1: The distributions of cost and valuation parameters used in the baseline (that is, pre-tax or pre-subsidy) treatment as a function of the number of pairs of sellers and buyers in the session.

Balance of Market Size and Subject Pool for Tax Treatments

Tax on Buyer Session	Pool	Tax on Seller Session	Pool
tax8b5	econ 1	tax8s7	econ 2,3
Tax10b6	econ 1	tax9s1	econ 1
Tax11b4	econ 1	tax10s11	eng., mix
Tax12b8	econ 2,3	tax12s2	econ 1
Tax13b10	eng.	tax13s3	econ 1
Tax14b9	bus., eng., mix	tax15s12	econ 2,3

Table 2a: The six sessions in which the tax is imposed on the buyer ("taxb") and the six sessions in which the tax is imposed on the seller ("taxs") are listed in ascending order according to the number of pairs of traders in the market, as indicated by the number (from 8-15) after the three letters "tax" in the name of the session. The column "Pool" indicates whether the subjects in the experiment were entirely first-year economics majors (econ 1), or whether a significant number of them were second-year and third-year economics majors (econ 2,3), business students (bus), engineering students (eng) or a mixture of students (mix).

Balance of Market Size and Subject Pool for Subsidy Treatments

Subsidy to Buyer Session	Pool	Subsidy to Seller Session	Pool
sud9b7	eng.	sud9s11	econ 2,3, eng.
sud10b3	eng.	sud10s2	econ 1
sud11b6	econ 2,3, bus., eng.	sud11s4	econ 2,3
sud13b1	econ 1	sud12s5	econ 2,3, eng.
sud13b12	econ 2,3, social sci	sud13s8	eng., mix
sud15b9	econ 2,3, eng.	sud15s10	econ 2,3, eng.

Table 2b: A list of the 12 subsidy sessions sorted by treatment and in ascending order according to the number of seller and buyer pairs in the market (from 9-15). The column "Pool" reveals the academic major of the subjects in the experiment. The caption of Table 2a provides a description.

Tax Experiments




Figure 2: Median transaction prices for each of the 12 tax experiments by period.

Median Transaction Price by Period for Each Tax Session

period	tax9s1	tax12s2	tax13s3	tax11b4	tax8b5	tax10b6	tax8s7	tax12b8	tax14b9	tax13b10	tax10s11	tax15s12
1	30	28	30	28	30.5	32	27	29	31	29	32	30
2	30	30	30	27.5	29.5	30.5	30	33	30.5	30	34	32
3	32	31	33	29	30.5	32	29	33	31	31	32	31
4	32	31	32	30	30.5	30	29	33	31.5	31	33	32
5	32	30	31.5	29	30	30	28.5	33	31	31	33	30
6	33.5	32	30	30	29.5	30	29	32	31	32	34	31
7	32	32	31	30	30	30	29	33	32	32	32	31
8	32	32	32	32	30.5	30	30	32.5	32	31	32	31
9	36.5	38	36	26	27.5	28	42	27	26	28	34	36
10	34	37	35	28	29	26	37.5	28.5	28	29	36	36
11	35	38.5	36	27.5	27.5	27.5	39	28	27	28	34.5	36
12	35	37	35	27	28.5	27	35	27.5	27.5	27	34	35
13	37	37.5	35	27	30	27	36	27	28	27	34	36
14	36	37	37.5	27	28	27	36	28	28	27	35	36
15	37	37	36	27	28.5	27.5	36.5	27.5	27	27	34.5	36
16	36	37	37	27	28	27	37	27	27.5	27	34.5	37
17	36	37	37	27	27.5	27	36	27.5	27	27	34	37
18	36	37	38	28	27.5	27	36.5	27	27	27	35.5	36
19	36	37	37.5	27	27	27	36	27	27	27	36	37

Table 3: Median transaction price by period for each tax session. In period 9, the 10-unit tax was imposed. The first price in bold in each session indicates the first period from which the median price entered, and remained within, the competitive price tunnel in the tax treatment.

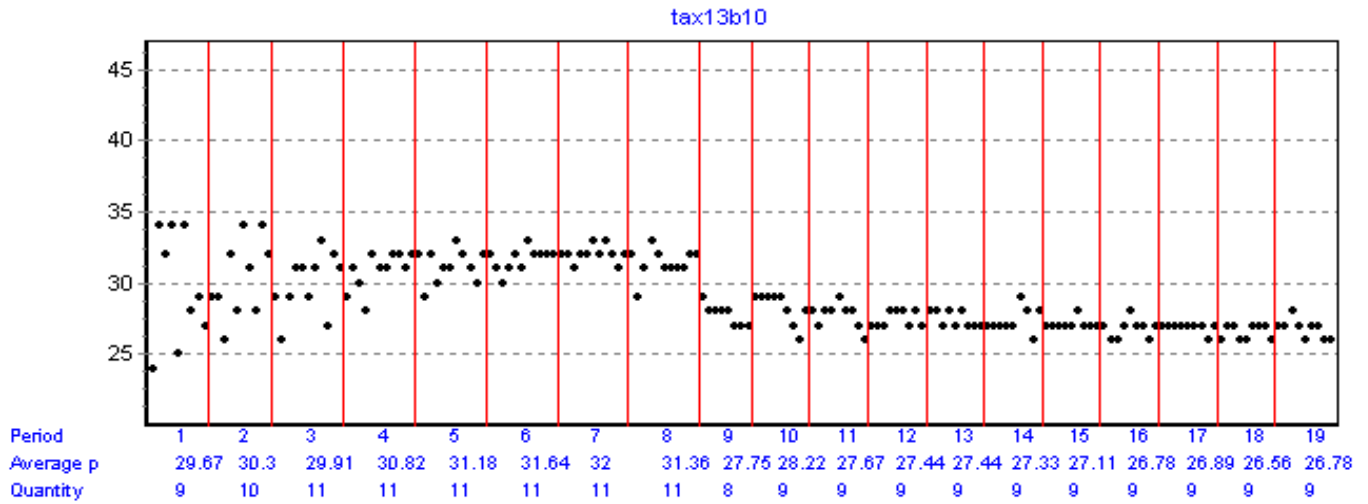


Figure 3: Transactions graph for session tax13b10. Each dot represents a transaction with the height of the dot indicating the negotiated price. The vertical lines separate the periods. The dots within each period are arranged in chronological order. The average price and quantity traded are displayed beneath the x-axis for each of the 19 periods.

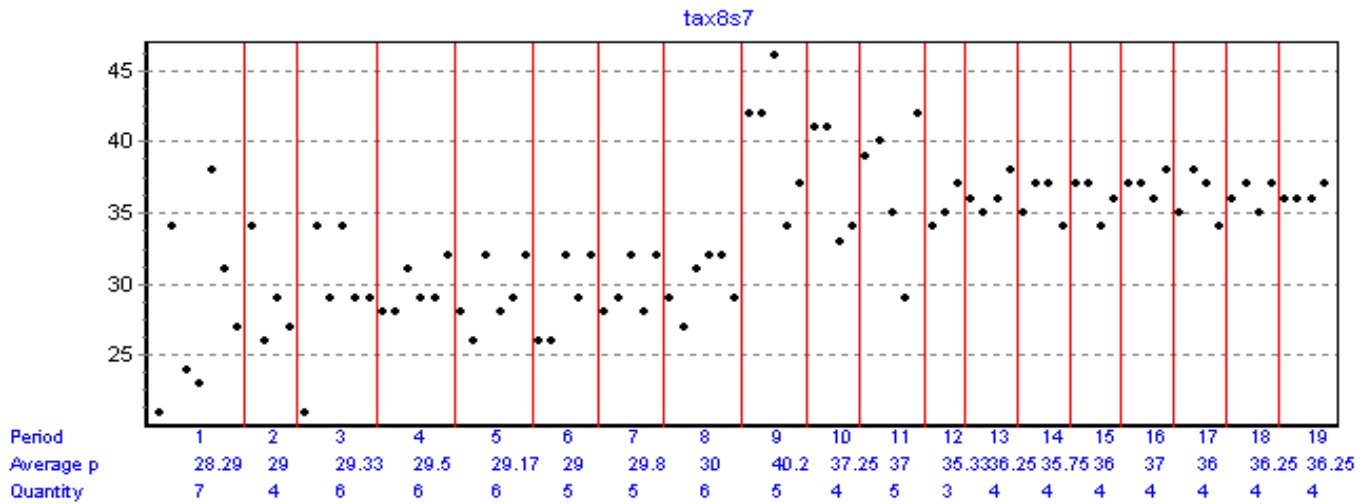


Figure 4: Transactions graph for session tax8s7.

Statistic	Sessions		
	taxb (all 6) (n=43)	taxs (3 large) (n=27)	taxs (3 small) (n=15)
average price, std. dev.	37.0, 0.69	37.0, 1.16	36.0, 0.85
median price	37	37	36
mean session rank for ave p	8.17	7.67	2

Table 4: Summary statistics for the period 19 prices for the *taxb* sessions, the three relatively large *taxs* sessions (*taxs* (3 large)) and the three relatively small *taxs* sessions (*taxs* (3 small)).

$$\text{Regression model: } P_{it} = B_{1i}D_1(1/t) + B_{12}D_2(1/t) + \dots + B_{1k}D_k(1/t) + B_2 \frac{(t-1)}{t} + u_{it}$$

<i>taxb</i> regression : dependent variable P_{it}							
$B_{1, \text{tax8b5}}$	$B_{1, \text{tax10b6}}$	$B_{1, \text{tax11b4}}$	$B_{1, \text{tax12b8}}$	$B_{1, \text{tax13b10}}$	$B_{1, \text{tax14b9}}$	B_2	n
28.71*	26.75	25.60**	28.11*	27.90	28.01*	27.25	461
(.606)	(.529)	(.536)	(.453)	(.426)	(.392)	(.102)	
above	none	below	above	none	above		

* The difference between B_{1i} and B_2 estimates is significant at the 5% level (t-test).

** The difference between B_{1i} and B_2 estimates is significant at the 1% level (t-test).

Table 5a: Results from the Ashenfelter-El-Gamal dynamic linear regression model for the *taxb* experiments. We estimate the initial price, B_{1i} , for each experiment i and a common convergence price asymptote, B_2 . (Standard errors in parentheses.) The difference between the initial price term and the convergence asymptote indicates the direction of convergence. Namely, if $B_2 > B_{1i}$, experiment i converges from below; whereas, if $B_2 < B_{1i}$, experiment i converges from above. “none” indicates no direction of convergence; that is, the difference between B_{1i} and B_2 is not significant. An initial sense of moral obligation among the buyers to bear the burden of the tax predicts convergence from above.

<i>taxs</i> regressions : dependent variable P_{it}							
$B_{1, \text{tax8s7}}$	$B_{1, \text{tax9s1}}$	$B_{1, \text{tax10s11}}$	$B_{1, \text{tax12s2}}$	$B_{1, \text{tax13s3}}$	$B_{1, \text{tax15s12}}$	B_2	n
39.30**	35.02	34.79*	38.66**	35.76	35.25*	36.23	459
(.786)	(.838)	(.712)	(.645)	(.603)	(.509)	(.143)	
above	none	below	above	none	below		
---	---	---	38.23*	35.35*	34.87**	36.66	299
			(.616)	(.577)	(.488)	(.167)	
			above	below	below		
39.94**	35.86	35.51	---	---	---	35.44	160
(.845)	(.909)	(.773)				(.254)	
above	none	none					

* The difference between B_{1i} and B_2 estimates is significant at the 5% level (t-test).

** The difference between B_{1i} and B_2 estimates is significant at the 1% level (t-test).

Table 5b: Three separate regressions for the *taxs* experiments using the Ashenfelter-El-Gamal model. The first row reports the results using all six *taxs* sessions. The second row uses the three large *taxs* sessions, while the third row uses the three small *taxs* sessions. An initial sense of moral obligation among the sellers to bear the burden of the tax predicts convergence from below.




Figure 5: Transactions graph for session sud10b3.

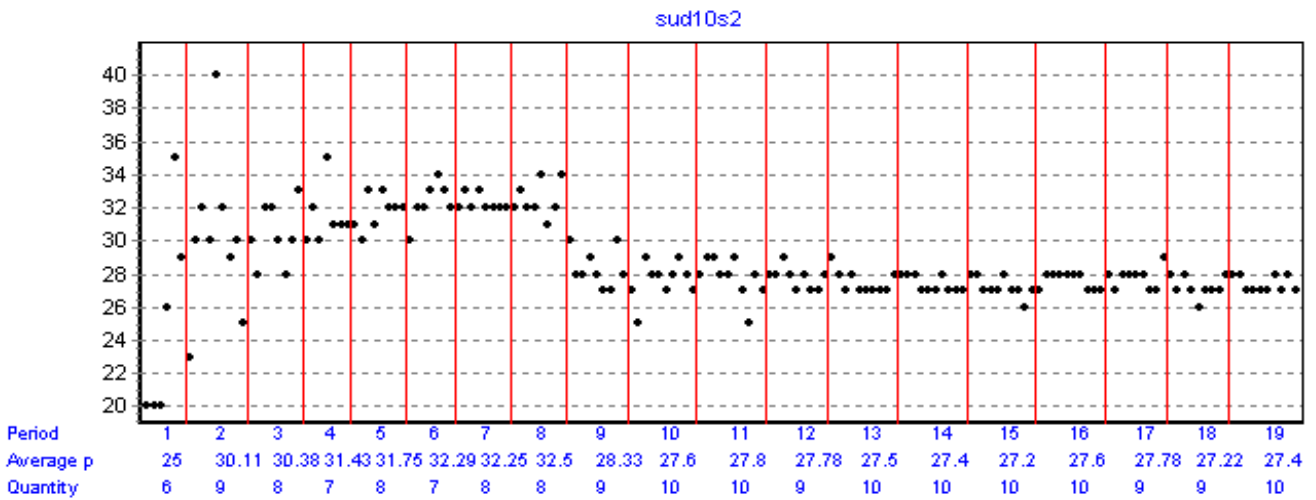


Figure 6: Transactions graph for session sud10s2.

Subsidy Experiments




Figure 7: Median transaction prices for each of the 12 subsidy experiments by period.

Median Transaction Price by Period for Each Subsidy Session

period	sud13b1	sud10s2	sud10b3	sud11s4	sud12s5	sud11b6	sud9b7	sud13s8	sud15b9	sud15s10	sud9s11	sud13b12
1	34	23	35	30	34	32	31.5	27	28	30.5	33.5	28.5
2	32	30	33	32	31	31	32	28.5	30	32	33	30
3	31	30	31.5	31	31	29	32	29	30	33	33	29
4	32	31	31	31	31.5	31	32	30	30	32	33	30
5	32	32	31.5	31	31.5	31	31	30	31	31	33	30
6	32	32	31.5	30	32	31	31	30	31	32	33	30
7	32	32	32	30	32	31	32	30	30	32	33	30
8	32.5	32	32	31	31	31	32	30	31.5	32	33	30
9	38	28	39	25.5	28	40	40	27.5	35.5	27	31	37
10	38	28	38	26	26.5	39	41	27	36	27	27.5	37
11	38	28	38	25.5	27	40	40.5	27	36	27	28	38
12	38	28	37	25	27	38	41	27	38	26	29	38.5
13	37.5	27	37	26	27	37	41	27	37.5	26	28	38
14	38	27	37	26	26	38	40	27	37	26	28	38
15	37	27	37	26	26	37	40	27	37	26	27	38
16	37	28	37	26	26.5	37	39.5	27	37	26	27	38
17	37	28	37	26	26	37	40	27	37	26	27	38
18	37	27	37	27	26.5	37	39.5	27	37	26	27	38
19	37	27	37	26	26	37	39	27	37	26	27	37.25

Table 6: Median transaction price by period for each subsidy session. In period 9, the 10-unit subsidy was added. The first price in bold in each session indicates the first period from which the median price entered, and remained within, the competitive price tunnel in the subsidy treatment.

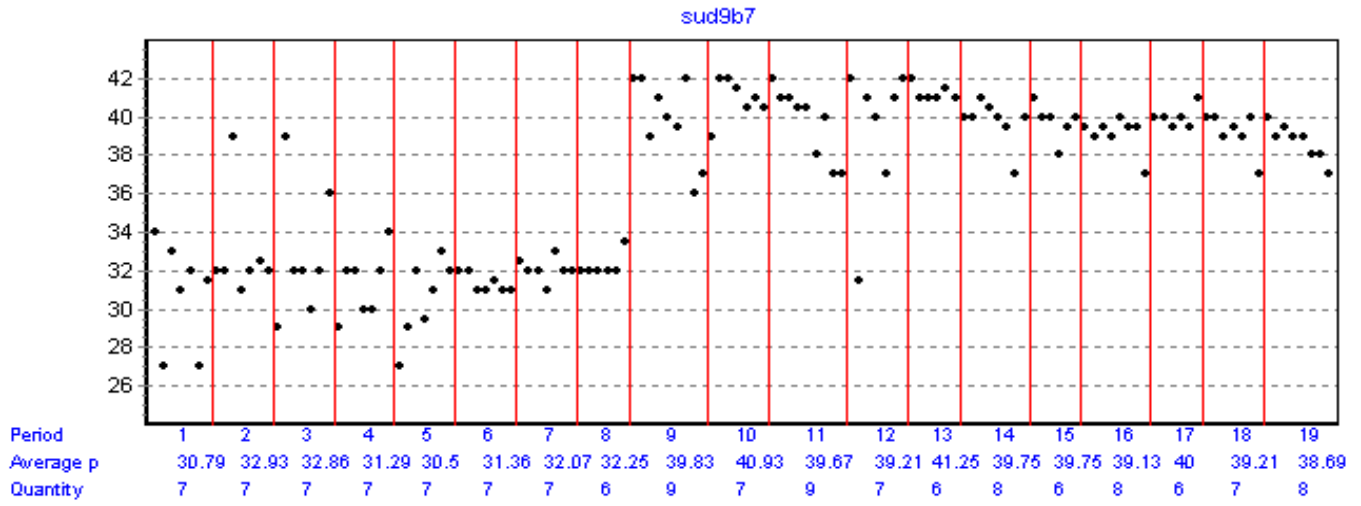


Figure 8: Transactions graph for session sud9s7, a case of sustained seller collusion. Sellers began to collude in the first period of the subsidy treatment (period 9) and maintained prices above the competitive range throughout the remainder of the experiment.

$$\text{Regression model: } P_{it} = B_{11}D_1(1/t) + B_{12}D_2(1/t) + \dots + B_{1k}D_k(1/t) + B_2(t-1) + u_{it}$$

<i>sudb</i> regressions : dependent variable P_{it}							
$B_{1, sud10b3}$	$B_{1, sud11b6}$	$B_{1, sud13b1}$	$B_{1, sud13b12}$	$B_{1, sud15b9}$	$B_{1, sud9b7}$	B_2	n
37.78 (.342) none	40.49** (.347) above	38.15** (.307) above	37.59 (.291) none	35.27** (.278) below	---	37.20 (.072)	659
---	---	---	---	---	40.24 (.526)	39.51 (.279)	81

* The difference between B_{1i} and B_2 estimates is significant at the 5% level (t-test).

** The difference between B_{1i} and B_2 estimates is significant at the 1% level (t-test).

Table 7a: Results from the Ashenfelter-El-Gamal dynamic linear regression model for 5/6 *sudb* experiments in row 1 and a separate regression for the outlier experiment *sud9b7* in row 2. I estimate the initial price, B_{1i} , for each experiment i and a common convergence price asymptote, B_2 . (Standard errors in parentheses.) The difference between the initial price term and the convergence asymptote indicates the direction of convergence. Namely, if $B_2 > B_{1i}$, experiment i converges from below; whereas, if $B_2 < B_{1i}$, experiment i converges from above. “none” indicates no direction of convergence; that is, the difference between B_{1i} and B_2 is not significant. An initial sense of moral entitlement among the buyers to benefit from the subsidy predicts convergence from below.

<i>suds</i> regression : dependent variable P_{it}							
$B_{1, sud9s11}$	$B_{1, sud10s2}$	$B_{1, sud11s4}$	$B_{1, sud12s5}$	$B_{1, sud13s8}$	$B_{1, sud15s10}$	B_2	n
30.52** (.311) above	29.02** (.292) above	24.75** (.291) below	27.29** (.265) above	27.73** (.257) above	26.71 (.233) none	26.60 (.059)	736

* The difference between B_{1i} and B_2 estimates is significant at the 5% level (t-test).

** The difference between B_{1i} and B_2 estimates is significant at the 1% level (t-test).

Table 7b: Regression results from the Ashenfelter-El-Gamal model for the *suds* experiments. An initial sense of moral entitlement among the sellers to benefit from the subsidy predicts convergence from above.