THE JOURNAL OF INDUSTRIAL ECONOMICS 0022-1821
Volume LXI March 2013 No. 1

WHEN DO LARGE BUYERS PAY LESS? EXPERIMENTAL EVIDENCE*

BRADLEY J. RUFFLE†

The rise of mega-retailers has precipitated a growing literature on large-buyer discounts. According to Rotemberg and Saloner [1986] and Snyder [1998], large buyers' ability to obtain price discounts depends on their relative size and the degree of seller competition. I test experimentally implications of this theory concerning the number of sellers and the sizes of buyers in the market. The results track the comparative-statics predictions to a surprising extent. Subtle changes in the buyer-size distribution or number of sellers can create or negate large-buyer discounts. The results highlight the previously unexplored role of the demand structure in determining buyer-size discounts.

I. INTRODUCTION

> 'Now what is your best price?' And if they told me it's a dollar, I would say, 'Fine, I'll consider it, but I'm going to your competitor, and if he says 90 cents, he's going to get my business. So, make sure a dollar is your best price.' (Claude Harris, Wal-Mart's first buyer, quoted from Walton and Huey [1993], p. 236)

> The negotiating prowess of Wal-Mart buyers is renowned in buying/vendor circles . . . Vendor relationships are important but the only guarantee of product continuity for vendors is to sharpen their pencils and to provide the best price, the first time it is requested. Failing to do so will lead to replacement in favor of a lower-cost provider. Competitive bids are encouraged and it is normal operating procedure for vendors to be pitted against one another . . . We didn't care what everyone else in the business world was paying for a product or service, we demanded the lowest possible Wal-Mart price. (Bergahl [2004], p. 30)

THE ACADEMIC LITERATURE AND BUSINESS PRESS (see, e.g., Fishman [2003]) often refer to two sources of buyer power: the buyer's size and the intensity

*I thank Shai Bernstein, Danny Cohen-Zada, Naomi Feldman, Hans Normann, Jean Beuve, Oren Rigbi, Oscar Volij, two anonymous referees, the Editor and seminar participants at Bar-Ilan University, Goethe University, Tel Aviv University and the APESA 2009 and IIOC 2009 meetings for useful comments. Ayala Waichman provided valuable research assistance. Tim Miller programmed the experimental software.
†Author's affiliation: Department of Economics, Ben-Gurion University, Beer Sheva, 84105, Israel.
e-mail: bradley@bgu.ac.il.

of seller competition. In fact, as the above quotes from Wal-Mart's first buyer and a former Wal-Mart executive reveal, the world's largest retailer extracts low prices from suppliers by threatening to buy from the competition. The opportunity cost of not agreeing to Wal-Mart's price demands is the loss of a voluminous sale to a more conciliatory competitor.

In this paper, I examine experimentally how large a buyer needs to be in order to obtain a lower price than its competitors. Size is measured as the number of units the buyer demands. Is a buyer's absolute size sufficient to determine whether a large buyer pays less? Or, is the decisive measure a buyer's size relative to other buyers in the market? What other aspects of the demand structure are relevant in determining the existence and magnitude of large-buyer discounts? What role does the degree of seller competition play in setting prices to buyers of different sizes? To answer these questions, I test experimentally a number of parameterized versions of a theoretical model that lay bare sellers' considerations in deciding how much to charge buyers of different sizes. In doing so, this is the first paper to bring to light the importance of the demand structure in the prices that buyers pay.

The theoretical framework of this paper has its origins in Rotemberg and Saloner [1986] (henceforth abbreviated as R-S). In their pioneering article, they demonstrate that periods of high demand (booms) are most conducive to price wars since current demand is high relative to future expected demand. To prevent price wars in booms, firms lower their prices just enough so that deviation is no longer profitable. Snyder [1996, 1998] reinterprets the macroeconomic shocks in R-S as buyers of different sizes arriving each period. Large buyers entice sellers to underbid one another aggressively in an attempt to supply the large quantities demanded. To prevent undercutting, sellers tacitly coordinate on lower prices when confronted with a large buyer. Thus, similar to Wal-Mart, large buyers in Snyder [1996, 1998] obtain a discount through their ability to force sellers to compete in price with one another. Notice that buyer size on its own is ineffective in this setup. Both in these models and in my experiments, buyers of all sizes behave identically, purchasing all of their available units of demand from the lowest-priced seller. As a result, buyers of all sizes would pay the same price against a monopoly seller. It is the interaction between buyer size and seller competition that yields, in some instances, buyer-size discounts.

In this paper, I parameterize the theoretical models of R-S and Snyder [1996, 1998] to make more precise the influences of buyer size and seller competition. For instance, a merger between two smaller buyers can negate the discount to the largest buyer in the distribution by reducing the latter's relative size. Similarly, removing a seller from the market against a given buyer-size distribution can eliminate large-buyer discounts. I test these and other predictions in controlled price-setting experiments.

Across a number of experimental treatments, I vary the demand structure (i.e., number of buyers in the market and size of each buyer) in a controlled manner that allows me to discern the influence of the demand structure on sellers' ability to collude tacitly as a function of buyer size. In each period of these indefinitely repeated games, a buyer of randomly determined size is drawn from a known distribution containing three or four buyer sizes. Symmetric sellers each simultaneously offer a price at which to supply the buyer's desired number of units, with the lowest-priced seller fulfilling these units. All other higher-priced sellers earn zero.

Overall, the results confirm the theory to a large extent. A large-buyer discount is predicted and observed in one treatment. In three others, no large-buyer discount is predicted and none is observed. Deviations from the theory are limited to one treatment in which a large-buyer discount is observed, but not predicted, and the presence of small-buyer premia. Yet, an analysis of buyer-size discounts over repeated games reveals that these deviations weaken and even disappear altogether over time. At the same time, those findings consistent with the theory are shown to be robust over the entire duration of the experiment.

In the next section, I review briefly the related literature. Section III summarizes R-S and Snyder's [1996, 1998] application of R-S to large-buyer discounts. In section IV, I detail the experiments designed to test comparative-statics predictions of these models. Section V presents the results. In section VI, I evaluate alternative explanations for the paper's main findings. Section VII concludes.

II. RELATED LITERATURE

A growing literature identifies a number of plausible sources of large-buyer discounts.[1] The starting point for the source of buyer-size discounts on which this paper focuses is Rotemberg and Saloner [1986]. In their setup, symmetric oligopolists face random *iid* demand fluctuations in each period of an infinitely repeated game. When demand is high, the benefit from undercutting the joint-profit-maximizing price is also high. To prevent undercutting and possible price wars, firms settle on the highest sustainable price, namely, the highest price at which the profit from unilaterally deviating from this price plus the subsequent punishment profit falls just short of the present discounted collusive profit. The implication is that collusive prices are lower in periods of high demand than in those of low demand.[2]

[1] See Ruffle [2009] and Snyder [2008] for surveys.

[2] Kandori [1991] shows that the R-S counter-cyclical pricing result is robust to serially correlated demand shocks as long as the discount factor is sufficiently high with respect to the number of firms. Bagwell and Staiger [1997] provide a general characterization of pro and counter-cyclical pricing as a function of the demand conditions.

Snyder [1996, 1998] interprets the random i.i.d. demand shocks in R-S as a buyer of varying size arriving each period. Large buyers tempt sellers to shade the monopoly price in order to capture the entire demand for themselves. To prevent such out-of-equilibrium deviations, sellers collude tacitly on lower prices against large buyers than against smaller ones, yielding the large-buyer-discount result. Snyder [1996] endogenizes buyer size: buyers may carry over unfilled units of demand to subsequent periods. By purchasing all of their backlogged demand at once, large buyers obtain low prices from sellers. To attenuate this off-equilibrium-path threat, sellers may reduce their prices, even if the buyer purchases in every period. In Snyder [1998] buyer size is exogenously determined with sellers facing a buyer of randomly determined size each period.

Notice that the source of buyer-size discounts in these models in no way relies on large buyers' superior negotiation skills—a trait commonly attributed to large buyers. In fact, buyers are essentially price-takers. It is the lucrative profit opportunity they represent that induces sellers to compete vigorously with one another to supply their large demand. Put differently, the interaction between buyer size and seller competition produces buyer-size discounts in this framework.

While this is the first test of this source of buyer-size discounts, Normann, Ruffle and Snyder [2007] test experimentally a theory of large-buyer discounts in a monopoly setting that depends on the monopolist's economies of scale. Consistent with the theory[3], their experiments reveal large-buyer discounts only if the monopolist has increasing marginal costs.

There are a few experimental tests that relate to R-S. Abbink and Brandts [2009] design experimental duopoly markets in which the total market demand follows a known evolution; demand persistently shrinks in each of the 27 periods in one treatment and expands in another. In contrast to R-S, Abbink and Brandts observe higher prices in the shrinking market; namely, sellers collude better when current demand is high relative to future demand.[4]

List [2009] conducts field experiments in open air markets. He employs confederate buyers to negotiate with sellers and an inside seller who supplies information on existing collusive agreements. Consistent with R-S, List finds that vendors cheat more frequently on these agreements on higher volume days.

Rojas [2012] examines the role of payoff information and monitoring in repeated two-person prisoners' dilemma experiments with random shocks

[3] In its most general terms, the theory conditions large-buyer discounts on the concavity of the aggregate surplus function (see, e.g., Stole and Zwiebel [1996], Inderst and Wey [2003], Chipty and Snyder [1999]).

[4] Their design is closer to Haltiwanger and Harrington (HH) [1991] in which demand follows a deterministic cycle. HH's prediction of lower prices when demand is declining is violated.

to payoffs. In one treatment designed to capture the essential features of R-S, players observe the payoff shock prior to choosing to cooperate or defect, and actions are *ex post* observable (i.e., perfect monitoring). In line with R-S, defection is most likely when the high-payoff matrix is played.

In the context of large-buyer discounts, Ellison and Snyder [2010] examine the wholesale prices of pharmaceuticals to buyers of different sizes and different substitution opportunities. They find that hospitals and HMO's—buyers able to substitute between off-brand drugs—pay substantially lower prices; whereas, despite their size advantage, chain drugstores pay similar prices to independents, even for off-patent drugs, because both are required to stock a range of substitute drugs in case a customer walks in with a prescription for a particular drug. Their results validate the importance of seller competition for volume discounts in addition to sheer buyer size.

In this paper, I test experimentally the conditions under which price-setting sellers willingly cede price discounts to large buyers. I make subtle changes in the demand and supply structures predicted either to induce or negate a large-buyer discount. I then investigate whether prices and, in particular, concessions to large buyers are responsive to these subtle parameter changes. All experimental treatments involve at least two sellers (as required by theory and supported by the results of Ellison and Snyder) facing a buyer drawn randomly from a known distribution of differently sized buyers.

III. MODEL

Using the buyer-size framework, I summarize the models of R-S and Snyder [1998] in this section. Consider the following infinitely repeated game with a time discount factor δ. An industry consists of N symmetric suppliers, each producing a homogenous product at constant marginal cost, which I normalize to 0. In each period t, the N suppliers each choose a price, p_{it}, and collectively face a single buyer of randomly determined size, s_t. Buyer size s_t is drawn from the distribution F, the support of which may be a continuous interval, $[\underline{s}, \overline{s}]$, or a discrete range, $\{\underline{s}, \ldots, \overline{s}\}$. The realization of F is observable to all players in the market. Buyer i is assumed to purchase from the lowest-priced seller and thus a buyer's strategy is completely determined by its demand function $D(p_t, s_t)$, where $p_t = min\{p_{it} | i = 1, \ldots, n\}$. If $n \leq N$ sellers set the same lowest price, each of these sellers supplies s_t/n units.

Like R-S, Snyder and others in this literature, I focus on the extremal equilibrium, namely, the subgame-perfect equilibrium that maximizes the suppliers' joint profit, and, also like the previous literature, I assume that sellers employ grim trigger strategies, namely, reversion to the static Nash equilibrium in all periods following any seller's deviation from the extremal equilibrium price.

Let $p^m(s_t)$ represent the monopoly price for buyer size s_t and $\Pi_i(p^m, s_t)$ seller i's resultant profit. By definition, this price maximizes the suppliers' joint profits given by, $\Pi^m(s_t) = [p^m(s_t) - c] \cdot D(p^m(s_t), s_t)$. For any seller i, it is profitable to deviate in period t against a buyer of size s_t if and only if,

$$\begin{aligned} N \cdot \Pi_i(p^m, s_t) + \frac{\delta}{1-\delta}\int_{\underline{s}_t}^{\overline{s}_t} \Pi_i(p^{NE}, s_t)dF(s_t) \geq \Pi_i(p^m, s_t) \\ + \frac{\delta}{1-\delta}\int_{\underline{s}_t}^{\overline{s}_t} \Pi_i(p^m, s_t)dF(s_t). \end{aligned} \tag{1}$$

The left-hand side of equation (1) represents the present discounted deviation profit composed of the one-period profit from infinitesimally undercutting the monopoly price plus the discounted punishment profit from setting the Nash equilibrium price for all future periods. The right-hand side of the equation represents the present discounted profit from maintaining the collusive price, p^m, and is equal to a single seller's one-period profit from monopoly pricing plus the continuation profit from this price.

Define s^* as the largest buyer for which the monopoly price is charged in equilibrium. Snyder makes the following two assumptions to rule out two trivial cases in which the same price is charged to buyers of all sizes:

$$N < \frac{1}{1-\delta} \tag{2}$$

$$\frac{\Pi^m(\overline{s})}{E_s(\Pi^m(s))} > \frac{\delta}{(1-\delta)(N-1)} \tag{3}$$

A violation of (2) yields marginal cost pricing as the only subgame-perfect equilibrium, while if (3) is violated, monopoly pricing for all buyer sizes is the unique subgame-perfect equilibrium. If both (2) and (3) hold, then Snyder establishes that $\underline{s} < s^* < \overline{s}$, implying that sellers deviate from the monopoly price for one or more buyer sizes greater than s^*.

The following proposition states this result formally.

Proposition I. (Snyder [1998]) In the extremal equilibrium, the price charged by all suppliers, $p(s_t)$, depends only on s_t and is independent of the pricing history. There exists $\underline{s} < s^* < \overline{s}$ such that (i) $p^*(s_t) = p^m(s_t)$ for $s_t \leq s^*$; (ii) $p^*(s_t)$ declines with s_t for $s_t > s^*$.

Proof. See R-S [1986, pp. 394–395] and Snyder [1998, p. 207].

R-S (p. 395) state two intuitive comparative statics results with respect to s^* : s^* decreases as either δ decreases or N increases. In the next section, I parameterize R-S and Snyder to test the theory and the sensitivity of seller pricing to the number of sellers and the demand structure. The goal is to understand the circumstances under which large buyers pay less.

IV. EXPERIMENTAL DESIGN AND PROCEDURES

IV(i). *Experimental Framework and Research Strategy*

To test the theory and the range of parameters under which large-buyer discounts obtain, I design a number of treatments. In each treatment, just like the theory, N symmetric sellers face a buyer of size s_t in each period. The sequence of buyer sizes is i.i.d. from a known distribution, F.[5] After observing the buyer-size realization, s_t, in round t, sellers simultaneously choose prices. The buyers are computerized and simply purchase from the lowest-priced seller. In the event that more than one seller sets the lowest price, the buyer divides its purchases equally among these sellers. After a buyer makes its purchases, sellers observe their own round profits and the other sellers' chosen prices. The round is over.

The theory applies to infinitely repeated games, the most credible laboratory implementation of which involves a repeated game with a known end-game probability, δ.[6] Thus, in any round t, the game continues to round $t+1$ with probability δ, and the game ends with probability $1-\delta$ and, time permitting, a new game with a different cohort of sellers begins.

The impact of the buyer demand structure on seller pricing has not been tested. Consequently, rather than commit *a priori* to an entire experimental design only to find that all of the treatments produce qualitatively similar results (e.g., no buyer-size discounts), I begin with a treatment predicted and behaviorally likely to yield large-buyer discounts. I let the results of this treatment guide my choice of parameters of subsequent treatments according to the following logic. If the initial treatment fails to produce large-buyer discounts, then I can, for instance, introduce an additional seller or increase the largest buyer's size, holding all other buyer sizes constant. On the other hand, if this initial treatment yields the predicted large-buyer discount, I can shrink the largest buyer sufficiently to negate the prediction of large-buyer discounts. If a discount to this buyer nonetheless persists, I can explore additional directions in the demand and supply structures expected to further weaken sellers' incentives to deviate against the largest

[5] In the absence of seasonality in sales, i.i.d. buyer-size sequences seem as plausible as any other assumption. Moreover, the assumption of stochastic i.i.d. demand is ubiquitous in the inventory control literature (see, e.g., Gallego and van Ryzin [1994]).

[6] See Roth and Murnighan [1978] for a classic laboratory study of infinitely repeated bargaining games and Feinberg and Husted [1993] for more recent infinitely repeated duopoly experiments.

buyer. In this way, I use the findings of one treatment to inform the design of the next with the goal of mapping out the boundaries of large-buyer discounts.

IV(ii). *Choice of Experimental Parameters*

In all treatments, I limit the number of sellers to two (the minimum number required by theory) or three. Collusion in experiments with more than two sellers has been shown to be unattainable (e.g., Isaac and Reynolds [2002]), even in the generally less competitive Cournot framework (Huck *et al.* [2004]). The inability of three or more sellers to maintain prices above the Nash equilibrium precludes buyer-size discounts or price variance more generally.

There are M equally likely buyer sizes. Each has inelastic demand for s_t units of production at time t, $s_t \in [\underline{s}, \overline{s}]$. A buyer of size s_t values each of his units of demand at 10, implying a collusive price $p^m = 10$.

In addition to the monopoly price, I permit two other prices, an undercutting price, p^{under} and a punishment price, p^{punish}, where $p^m > p^{under} > p^{punish}$.[7] Consistent with these interpretations, I place the following restriction on each of these prices.

$$(4) \qquad p^{under} \cdot s_t > \frac{p^m}{N} \cdot s_t \quad \forall s_t, N > 1$$

$$(5) \qquad p^{punish} \cdot s_t < \frac{p^m}{N} \cdot s_t \quad \forall s_t, N > 1$$

Equation (4) indicates that the individual seller earns more in round t from undercutting and supplying all of the buyer's units than charging the collusive price and earning $1/N$ of the total profit. Thus, p^{under} cannot be set too low relative to p^m to avoid the latter's being the equilibrium price for all buyer sizes. On the other hand, if p^{under} is too similar to p^m, then collusion may never be sustainable for any buyer size. Equation (5) requires a punishment price sufficiently low that deviation to this price is not profitable, even in the short run. I resolve these tensions by choosing p^{under} equal to 8 and p^{punish} equal to 2.

We can thus rewrite condition (2) for deviation against the largest buyer size, $\overline{s}$, from the monopoly price of 10 to the price below it, p^{under} as,

[7] In an indefinitely repeated-game experiment with imperfect monitoring, Feinberg and Snyder [2002] also restrict the price space to three. Three prices in my framework are meant to facilitate collusion (relative to more prices).

TABLE I
REPEATED GAME LENGTHS

Game	Periods (T)
1	4
2	17
3	8
4	5
5	12
6	4

The number of repetitions in each game. A perfect strangers design was employed between games.

$$8\overline{s}+\frac{\delta}{1-\delta}\cdot\sum_{s_t=\underline{s}}^{\overline{s}}\frac{2s_t}{M\cdot N}>\frac{10\overline{s}}{N}+\frac{\delta}{1-\delta}\sum_{s_t=\underline{s}}^{\overline{s}}\frac{10s_t}{M\cdot N}. \tag{6}$$

The left-hand side of (6) is the round profit from deviating to p^{under} plus the present discounted punishment profit from all N sellers charging p^{punish} indefinitely. The right-hand side of (6) is the seller's share of the present discounted collusive profit.

In selecting the continuation probability (equivalently, the discount factor), δ, Feinberg and Husted's [1993] duopoly experiments illustrate the importance of a high discount factor to achieve collusion. However, if the discount factor is too high, the monopoly price will be the equilibrium price for all buyer sizes. To achieve the desired balance, I choose the discount factor equal to 0.8, holding it constant across the initial and all subsequent treatments.

Prior to conducting the experiments and using the discount factor of 0.8, I randomly determined the length of each of the six repeated games, displayed in Table I. I applied these realizations to all sessions and treatments so that all seller cohorts regardless of session or treatment saw the identical randomly determined sequence of game lengths. Similarly, to permit a controlled comparison of sellers' behavior across treatments, I randomly drew one sequence of 70 buyer sizes from a uniform distribution over four buyer sizes and applied this sequence to all seller cohorts in all sessions and treatments with four buyer sizes.

IV(iii). *Initial Treatment*

The initial treatment I refer to as 'mega-retailer' and display its details in the first column of Table II. In this treatment, three sellers face the above constants and a buyer-size distribution $F=\{1, 2, 6, 14\}$. Deviation from the monopoly price is unprofitable against buyers of sizes 1, 2 and 6; whereas, a one-time deviation to p^{under} against the largest buyer of size 14 is profitable, yielding an excess present-discounted profit of 4 (as shown in the fourth-to-last row in Table II). Thus, the equilibrium prices (displayed in

TABLE II
GAME CONSTANTS, PARAMETERS, PROFITS AND EQUILIBRIUM PRICES

	Treatment				
Constants	*mega-retailer*	*large buyer*	*2 sellers*	*buyer merger*	*same mean*
p^m, p^{under}, p^{punish}	10, 8, 2	10, 8, 2	10, 8, 2	10, 8, 2	10, 8, 2
δ	0.8	0.8	0.8	0.8	0.8
Parameters					
Sellers (N)	3	3	2	3	3
Buyers (M)	4	4	4	3	4
Buyer Size 1 (LB)	14	10	10	10	8
Buyer Size 2	6	6	6	6	6
Buyer Size 3	2	2	2	3	3
Buyer Size 4 (SB)	1	1	1	—	2
net PDV deviation against LB	4	–4	–46	–20.89	–13.33
Deviate against LB?	Yes	No	No	No	No
range of r for which true	$(-\infty, 0.02]$	$[-0.02, \infty)$	$[-0.35, \infty)$	$[-0.14, \infty)$	$[-0.10, \infty)$
Eq'm Prices by Buyer Size	8, 10, 10, 10	10, 10, 10, 10	10, 10, 10, 10	10, 10, 10	10, 10, 10, 10

For all five treatments, the three available prices and discount factor are held constant. The five treatments differ by the number of sellers (N) and buyers (M) in the market and the distribution of buyer sizes. The net discounted profit from deviation against the largest buyer (LB) in the distribution (from equation (6)) and, as a result, whether sellers have an incentive to deviate from p^m against LB, the range of CRRA coefficients for which (not) to deviate against LB holds and equilibrium prices by treatment and buyer size (in descending order) are displayed in the last four rows.

the last row of the table) are 10 for the three smallest buyers and 8 for the largest buyer of size 14. Recall that the underlying theory assumes that sellers are risk neutral. For any parameterization, we can evaluate the robustness of the theory to alternative risk preferences. For example, if we replace each term in (6) with a typical CRRA utility function, say, $u(x)=\frac{x^{1-r}}{1-r}$, where x is the profit from a given buyer-size-price combination and r is the CRRA coefficient, the predicted large-buyer discount continues to hold in equilibrium for $r \in (-\infty, 0.02]$, as seen in the second-to-last row of Table II. In words, sellers offer a lower price to buyer size 14 for any degree of risk-lovingness, risk neutrality and only the slightest degree of risk aversion.

We don't anticipate prices to match the above equilibrium point predictions, in part because of the noted difficulty in achieving full collusion in laboratory experiments. Notwithstanding, if the theory captures all of the relevant behavioral considerations of the setup it aims to explain and my experimental design captures the behaviorally relevant aspects of the theory, then we would expect the experimental results to validate the theory's comparative statics predictions. In this treatment, the interesting qualitative prediction is that the largest buyer of size 14 will pay a lower price than the other buyers in the distribution.

IV(iv). *Experimental Procedures and Subject Pool*

The experiment consists of a number of stages, each to be discussed in turn below. Upon arrival, subjects were seated at a computer terminal.[8] They each read a printout of the instructions at their own pace before I read aloud the instructions for all to hear. Any questions were answered privately before proceeding to five practice rounds, not included in the calculation of subjects' profits. Subjects then gained experience in an identical game without discounting. That is, they played a finitely repeated game with a known terminal round 15. The randomly determined buyer size and terminal period and the reforming of seller cohorts at the termination of each game render the decision environment somewhat complex. Prior experience in a game of known length allows subjects to familiarize themselves with the role of buyer size on their decision task and profits. Only in section VI will I invoke the results from this finitely repeated game to evaluate alternative explanations for the findings.

At the completion of Part 1, subjects proceeded to Part 2 (see the Appendix for the instructions for Parts 1 and 2), the core of the experiment in which the terminal round in each game is randomly determined. A perfect strangers design is employed between games to minimize the dependence

[8] All of the experiments were conducted using z-Tree (Fischbacher [2007]).

between games. Specifically, at the end of the practice rounds, the 15 rounds from Part 1 and each repeated game in Part 2, seller groups are randomly reformed conditional on not having previously been in the same group as any of the other sellers in their present group.

Because of the economic nature of this experiment and the types of reasoning and computations required of subjects, I restricted recruitment to economics and business majors. In total, 172 subjects participated in one of five treatments. An entire session lasted about one hour and 45 minutes. Subjects were paid based on their cumulative earnings from Parts 1 and 2. The average payment was 71.6 NIS (*s.d.* = 15.1 NIS) or about $19 USD, well above the student wage of 19 NIS an hour.

V. RESULTS

To examine buyer-size discounts, I rely on two measures, sellers' offers and prices paid by buyers. In almost all cases, the qualitative results are identical regardless of the measure used. In some sense the price paid is more meaningful since prices, not offers, ultimately determine whether a buyer receives a discount. Also, outside the laboratory, prices, not offers, are observable, if at all. On the other hand, offers provide a fuller dataset with an observation for each seller in every round. Table III and Figure I display the mean offers and prices paid by buyer size for each treatment.

TABLE III
SUMMARY STATS AND FITTED OFFER PROBABILITIES BY TREATMENT AND BUYER SIZE

				Offer Probabilities		
Treatment	Buyer Size	Mean Offer	Mean Price	2	8	10
mega-retailer	1	7.65	7.08	0.082	0.796	0.120
	2	7.31	6.97	0.140	0.790	0.069
	6	6.91	6.15	0.207	0.751	0.040
	14	5.87	4.85	0.365	0.616	0.013
large buyer	1	7.52	6.54	0.122	0.681	0.192
	2	6.74	5.42	0.241	0.663	0.092
	6	6.41	5.23	0.301	0.625	0.064
	10	5.68	4.17	0.376	0.549	0.036
2 sellers	1	8.65	8.28	0.001	0.661	0.337
	2	8.33	8.11	0.006	0.824	0.170
	6	8.30	7.84	0.020	0.904	0.075
	10	8.03	7.87	0.022	0.907	0.064
merger	3	7.40	6.74	0.113	0.794	0.079
	6	7.09	6.46	0.163	0.778	0.054
	10	6.93	6.25	0.198	0.754	0.040
same mean	2	8.17	7.63	0.032	0.768	0.199
	3	7.76	7.38	0.081	0.821	0.092
	6	7.32	6.69	0.124	0.812	0.057
	8	7.19	6.44	0.137	0.795	0.043

Notes: Mean offers and prices by treatment and by buyer size and the fitted offer probabilities of the three prices derived from the second regression in each panel of Tables IV and V. The probabilities may not precisely sum to one due to rounding.

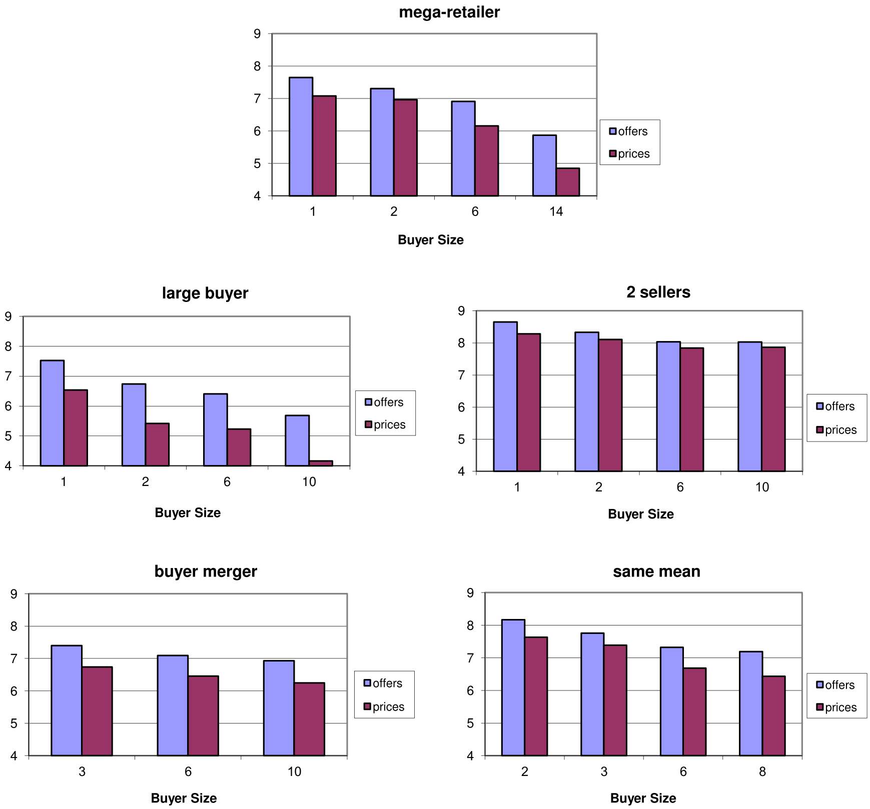

Figure 1
Mean Offers and Prices by Treatment and by Buyer Size

V(i). *Buyer-Size Discounts by Treatment*

Mega-Retailer Treatment. Thirty subjects participated in this initial treatment. The first noteworthy observation is that the mean offers and prices are well above the punishment price of 2 for all four buyer sizes, indicating that the measures taken to achieve collusion (e.g., only three available prices, a high continuation probability) succeeded to some degree.

Mean offers decrease monotonically in buyer size with the sharpest drop from buyer size 6 (6.91) to the largest buyer 14 (5.87). Mean prices reveal an even bigger large-buyer discount. The mean price is around 7 for the two smallest buyer sizes of 1 and 2, falls to 6.15 for buyer size 6 and swoons by more than a full point to 4.85 for the largest buyer of size 14. In percentage terms, the two smallest buyers each pays about 45% more than buyer size 14, while buyer size 6's price premium is 27%.

The left panel of Table IV reports ordered probit regressions for the *mega-retailer* treatment. The dependent variable is seller i's round t offer

TABLE IV
ORDERED PROBIT REGRESSIONS FOR MEGA-RETAILER, LARGE BUYER AND 2 SELLERS TREATMENTS

	mega-retailer			*large buyer*			*2 sellers*		
dependent variable regressor \ equation	$offer_{it}$ (1)	$offer_{it}$ (2)	$price_t$ (3)	$offer_{it}$ (4)	$offer_{it}$ (5)	$price_t$ (6)	$offer_{it}$ (7)	$offer_{it}$ (8)	$price_t$ (9)
BS6	0.47***	0.46***	0.61***	0.23***	0.19**	0.39***	0.05	0.04	0.08
	(0.12)	(0.12)	(0.18)	(0.09)	(0.10)	(0.14)	(0.09)	(0.09)	(0.12)
BS2	0.74***	0.78***	0.95***	0.41***	0.49***	0.56***	0.58***	0.66***	0.62***
	(0.13)	(0.15)	(0.23)	(0.16)	(0.16)	(0.17)	(0.13)	(0.14)	(0.17)
BS1	1.04***	1.08***	1.13***	0.87***	0.87***	1.02***	1.01***	1.11***	1.09***
	(0.15)	(0.15)	(0.25)	(0.15)	(0.16)	(0.21)	(0.13)	(0.14)	(0.19)
round	—	–0.02**	0.03**	—	–0.04***	–0.02	—	–0.02*	0.02
		(0.01)	(0.02)		(0.01)	(0.01)		(0.01)	(0.02)
game	—	0.19**	–0.31*	—	0.36***	0.21	—	0.19**	–0.12
		(0.08)	(0.19)		(0.09)	(0.17)		(0.08)	(0.16)
Part 1 subject control	No	Yes	—	No	Yes	—	No	Yes	—
Obs	1500	1500	500	1500	1500	500	2000	2000	1000
Mean Dependent Variable	6.95	6.95	6.28	6.60	6.60	5.36	8.26	8.26	8.03

Notes: [1]Dependent variable: seller *i*'s period *t* offer ((1), (2), (4), (5), (7), (8)) or the price paid in period *t* ((3), (6), (9)).
[2]Regressors: BS6, BS2 and BS1: indicators for buyer sizes 6, 2 and 1 (largest buyer omitted). game: the sequence of the repeated game. *round*: the overall round. *Part 1 subject control*: the number of rounds in which seller *i* chose the monopoly price in Part 1.
[3]Standard errors in parentheses are heteroskedasticity robust and clustered by seller (offer regressions) or by seller-cohort and game (price regressions).
[4]Coefficient significantly different from 0 at the 1% level ***, at the 5% level **, at the 10% level *.

(the first two regressions) or the lowest offer in each market in round t (namely, the price paid) (the third regression) with each of the three prices, 2, 8 and 10, as the possible realizations of the dependent measure. The standard errors reported are robust to heteroskedasticity and are corrected for possible non-independence of observations across rounds of play for the same seller by clustering on each seller's observations in the offer regressions[9] and on the seller-cohort-game cell in the price regressions. The basic ordered probit results in (1) of Table IV show that sellers offer significantly higher prices to all buyer sizes compared to buyer size 14.[10]

Regression (2) shows that the large-buyer discount is robust to the inclusion of controls for the overall round (*round*), the sequence of the repeated game (*game*) and an individual seller 'preference for collusion' parameter equal to the number of rounds from Part 1 in which the seller chose the monopoly price.[11] The coefficient on *round* is negative and highly significant, while the coefficient on *game* is positive and highly significant. Together these findings suggest that collusion erodes over time, but is restored in some measure at the beginning of each new game.[12] Wald tests on the coefficients of *BS1*, *BS2* and *BS6* in (2) indicate that, controlling for the above time trends, offers to buyer size 1 are higher than those to sizes 2 and 6 ($p = .01$ and $p < .001$, respectively). Similarly, offers to size 2 are higher than those to size 6 ($p < .001$).[13] In words, sellers graduate their offers to account for buyer size, leading to both a significant small-buyer premium and large-buyer discount.

I use my main offer regression of interest (2) to compute the mean fitted probabilities of choosing each of the three available prices by buyer size, with the *round* and *game* variables evaluated at their mean values. The right-hand columns of Table 3 display these fitted probabilities. The probability assigned to any of the three prices differs by no more than 0.06 between the two smallest buyer sizes. For instance, with probability 0.79

[9] I also tried clustering more finely on the seller-game cell. All results are identical.

[10] The significance and non-significance of all buyer-size effects are robust to GLS with random effects. The sole discrepancy between the two will be mentioned when discussing the *buyer merger* treatment. I choose to report the ordered probits because they also yield a fitted distribution of offers (displayed in Table III) that goes beyond simple mean offers.

[11] The coefficient on this latter variable is positive in each of the treatments, but significantly different from zero in this and the next *large buyer* treatments only.

[12] I observe this same pattern of declining offers over rounds, but offer increases with each new game in all subsequent treatments as well. This restart effect is commonly found in finitely repeated versions of another cooperation game, the public goods game, especially when the same cohort of subjects remains together for subsequent repeated games (see, e.g., Andreoni and Croson [2008]). Engle-Warnick and Slonim [2006] demonstrate that trusting behavior displays a restart effect in indefinitely repeated trust games.

[13] For perfectly inelastic demand curves, the theory does not admit such a possibility, instead predicting the same (monopoly) price for all buyers $s_t \leq s^*$. Snyder [1998, pp. 207–208] provides an example with individual downward-sloping demand curves, each with a vertical intercept of s_t, in which the extremal equilibrium price (equal to the monopoly price) actually increases with buyer size up to s^*.

sellers choose price 8 for both these buyer sizes with the remaining 0.21 split not too unequally between prices 2 and 10. However, the entries reveal a modest shift in the probability distribution toward a price of 2 for buyer size 6 and a larger shift for buyer size 14 with probability .365 placed on price 2 and a mere .013 on price 10.

I repeat the above regression analysis with the price in round t as the dependent variable.[14] The estimates in (3) reveal that the highly significant quantity discount for buyer size 14 with respect to all smaller buyers holds for prices as well. The smallest-buyer premium is less robust: the two smallest buyers pay similar prices to one another ($p = .31$), whereas both pay significantly more than buyer size 6 ($p = .01$ and $p = .03$ for *BS1* and *BS2*, respectively).

Having observed a significant large-buyer discount anticipated by the theory, I next design a treatment predicted to negate the discount.

Large Buyer Treatment. This treatment makes one subtle change to the *mega-retailer* treatment: the largest buyer is reduced in size from 14 to 10. As Table II shows, this single change renders unprofitable a deviation from the monopoly price against the largest buyer. The net discounted profit from deviation against the largest buyer of size 10 is −4, compared to 4 against size 14 in *mega-retailer*. Consequently, the range of CRRA coefficients for which deviation remains unprofitable is $[-0.02, \infty)$, the reflection about zero of the range for which deviation is profitable in *mega-retailer*.

Thirty subjects participated in this treatment. Once again, both mean offers and prices decline monotonically with buyer size. Mean offers are 7.52 to the smallest buyer, decreasing to 6.74, 6.41 and 5.68 to buyer sizes 2, 6 and 10, respectively. Like the *mega-retailer* treatment, the large-buyer discount is even more pronounced if we consider prices, instead of offers. The large buyer pays 4.17, 25.5% less than buyer size 6, with this discount ballooning to 57% compared to the smallest buyer.

The regressions in the center panel of Table IV confirm the significance of the large-buyer discount with respect to all smaller buyers in all specifications. Based on regression (5), the mean fitted offer probabilities in Table 3 reveal that as buyer size increases the probabilities with which sellers choose price offers of 10 and 8 progressively fall, while price offers of 2 steadily increase. The smallest-buyer premium (*BS1*) is significant compared to the second-smallest buyer (*BS2*) in both offers ($p < .001$) and prices ($p = .01$). On the other hand, the two intermediate-sized buyers (*BS2* and *BS6*) pay similar prices ($p = .16$).

[14] To save space, I report only the regression corresponding to (2). The results of the regression analogous to (1) are qualitatively similar to the offer regression in this and all subsequent treatments. Any differences will be noted as they arise.

To recap, the observed large-buyer discount was not anticipated by the theory. Consequently, one may begin to wonder whether large-buyer discounts are ubiquitous. Perhaps subjects are so accustomed to large-buyer discounts that they instinctively respond by offering a lower price whenever the large buyer appears.[15] To nullify the observed large-buyer discount, the theory suggests numerous directions. The remaining subsections present alternative treatments, each of which makes a single change to the *large buyer* treatment. Each change renders deviation from the monopoly price against the largest buyer even more unprofitable, implying no discount to large buyers.

2 Sellers Treatment. Removing one seller so that only two sellers remain produces a substantial deviation profit shortfall of –46, compared to only –4 in *large buyer* (see Table II). Intuitively, the set of CRRA coefficients for which deviation remains unprofitable expands to $r \in [-0.35, \infty)$, encompassing all degrees of risk aversion, risk neutrality and considerable risk-lovingness.

Forty subjects participated in this treatment. Both the mean offers and prices in Table III and the top panel in Figure I suggest that the treatment succeeded: the mean offers to buyer sizes 10 and 6 (8.30 and 8.03, respectively) and the mean prices (7.84 and 7.87) are nearly identical. The regressions in the right panel of Table IV confirm the absence of a large-buyer discount relative to the second-largest buyer. Moreover, the fitted offer probabilities in Table III are almost indistinguishable for buyer sizes 6 and 10, with just over 90% probability on the offer of 8 for both buyers 6 and 10, about .07 on the offer of 10 and .02 on the offer of 2.

Yet both of these buyer sizes obtain a discount compared to buyer size 2 whose mean offer received is 8.33 and mean price paid is 8.11.[16] Likewise, buyer size 2 receives a discount in both offers and prices compared to buyer size 1 whose mean offer (price) is 8.65 (8.28) ($p < .001$ and $p = .01$, respectively). In addition, offers and prices for all four buyer sizes in *2 sellers* appear to be substantially higher than their counterparts in *large buyer*.[17]

Buyer Merger Treatment. In this fourth treatment, I merge the two smallest buyers of sizes 1 and 2 into a single buyer of size 3, while holding

[15] Risk preferences are a highly improbable explanation for the large-buyer discounts observed thus far. The set of CRRA coefficients that is consistent with the discounts found in both treatments is $r \in (-\infty, -0.02)$, implying risk-seeking subjects. Yet across a variety of binary-choice lottery studies, Harrison and Rutström [2008] present robust evidence that subjects are risk averse, with estimates of r ranging from 0.1 to 0.8.

[16] The *BS2* coefficient is significant beyond the 1% level in all regressions in Table IV and Wald tests of *BS6* and *BS2* are highly significant ($p < .001$ based both on offers in (8) and on prices in (9)).

[17] Huck *et al.*'s [2004] meta-analysis of Cournot experiments (with symmetric buyers) along with their own results also find two sellers better able to collude than three.

TABLE V
ORDERED PROBIT REGRESSIONS FOR BUYER MERGER AND SAME MEAN TREATMENTS

	buyer merger			*same mean*		
dependent variable regressor \ equation	$offer_{it}$ (10)	$offer_{it}$ (11)	$price_t$ (12)	$offer_{it}$ (13)	$offer_{it}$ (14)	$price_t$ (15)
BS6	0.13*	0.10	0.15	0.097**	0.07	0.04
	(0.07)	(0.07)	(0.12)	(0.044)	(0.04)	(0.08)
BS3	0.39***	0.50***	0.44***	0.386***	0.49***	0.52***
	(0.11)	(0.12)	(0.14)	(0.120)	(0.12)	(0.15)
BS2	—	—	—	0.842***	0.85***	0.77***
				(0.153)	(0.15)	(0.19)
round	—	–0.03***	0.01	—	–0.03***	–0.04***
		(0.01)	(0.02)		(0.001)	(0.01)
game	—	0.38***	0.08	—	0.30***	0.44***
		(0.08)	(0.19)		(0.09)	(0.14)
Part 1 subject control	No	Yes	—	No	Yes	—
N	1650	1650	550	1950	1950	650
Mean Dependent Variable	7.15	7.15	6.49	7.61	7.61	7.04

Notes: See Table IV.

constant the two largest buyers of sizes 6 and 10. This subtle alteration in the demand structure effectively increases the mean buyer size from 4.75 to 6.33, a 33% increase stemming from two sources: i) the merged buyer is, by definition, larger than either of its component buyers; ii) the reduction in the number of buyers from four to three, which engenders an increased probability that sellers will face a buyer of size 10 or 6 in any given round. The end result is that deviation against the largest buyer of size 10 is now considerably more unprofitable (yielding a profit shortfall of –21) than in *large buyer* and the range of CRRA coefficients for which this result holds extends to $r \in (-0.14, \infty)$.

Thirty-three subjects participated in this treatment. The mean offers to buyer sizes 6 and 10 (7.09 and 6.93, respectively) and mean prices (6.49 and 6.25, respectively) differ by no more than 3% from one another. The ordered probits in the left panel of Table V reveal that offers to buyer size 6 are marginally significantly higher than offers to buyer size 10 in the absence of additional controls (10). Even this marginal large-buyer discount disappears when time controls are included: the coefficient on *BS6* in (11) is not significantly different from zero ($p = .16$). The fitted offer probabilities based on (11) support this finding with similar probabilities on price offers across all three buyer sizes. In fact, the probabilities of any given offer never differ by more than .035 between buyer sizes 6 and 10.

For prices, the regression results offer no trace of a significant discount to buyer size 10 compared to size 6: *BS6* is not significantly different from 0 ($p = .22$) in (12), which includes controls for *round* and *game*, nor in price regressions corresponding to (10) without these controls ($p = .38$) (not reported). Once again, as predicted, the largest buyer does not receive an unambiguous discount compared to the second-largest buyer.

The table also reveals a substantial small-buyer premium. Offers and prices to buyer size 3 are significantly higher than those to sizes 6 and 10 ($p < .01$ in all cases).

How do the prices paid by the merged buyer in this treatment compare to those of the pre-merger buyers of sizes 1 and 2 in the otherwise identical *large buyer* treatment? The model predicts the same monopoly price of 10 for all three buyer sizes. Notwithstanding, we've already anticipated and observed offers and prices well below this level. This allows for the behaviorally sensible outcome that the merged buyer pays less than the component pre-merger buyers.[18] On the contrary, Table II reveals that mean offers and prices to the merged buyer (7.40 and 6.74, respectively) are significantly *higher* than those to size 2 in *large buyer*.[19] In fact, mean offers and prices to buyer size 3 resemble more closely and are statistically indistinguishable from those of size 1 in *large buyer*.

One plausible explanation for this paradoxical finding is that prices in *buyer merger*, as in *2 sellers*, are sharply higher than those in *large buyer*. Here, fewer deviations and the less frequent use of punishment prices attest to sellers' ability to better collude against the largest buyer in *buyer merger*; whereas, when collusion breaks down on buyer size 10 in *large buyer*, sellers mete out punishment not only on this buyer, but on smaller buyers as well.

Same Mean Treatment. In this final treatment, I return to a buyer-size distribution with four distinct buyer sizes and a mean size of 4.75. Here, I shrink the size of the largest buyer from 10 to 8, while compensating with a two-unit increase in the smallest buyer's size from 1 to 3. Buyer sizes 2 and 6 remain unchanged. As equation (6) reveals, holding constant the mean buyer size while shrinking the largest buyer necessarily reduces the appeal of deviation. Table II shows that the net shortfall from deviating against the largest buyer drops from –4 in *large buyer* to −13.33 in *same mean*, while the set of CRRA coefficients supporting no deviation ranges from –0.10 to ∞.

Thirty-nine subjects participated in this treatment. To the largest and second-largest buyers of sizes 8 and 6, mean offers (7.19 and 7.32, respectively) and mean prices (6.69 and 6.44) once again appear similar to one another, as the theory predicts. Mean offers and especially prices rise substantially for buyer size 3 (7.76 and 7.38) and still further for buyer size 2 (8.17 and 7.63).

[18] Snyder [1996] shows that a merger benefits both the merged buyer and all other buyers in the form of lower prices in a setup in which multiple, variously sized buyers are simultaneously present in the market and each buyer is able to carry over unfulfilled units of demand to the next period.

[19] P-values are less than .01 for both non-parametric Mann-Whitney tests where, to minimize the dependence between observations, the average offer or price in a seller cohort and game is the unit of observation.

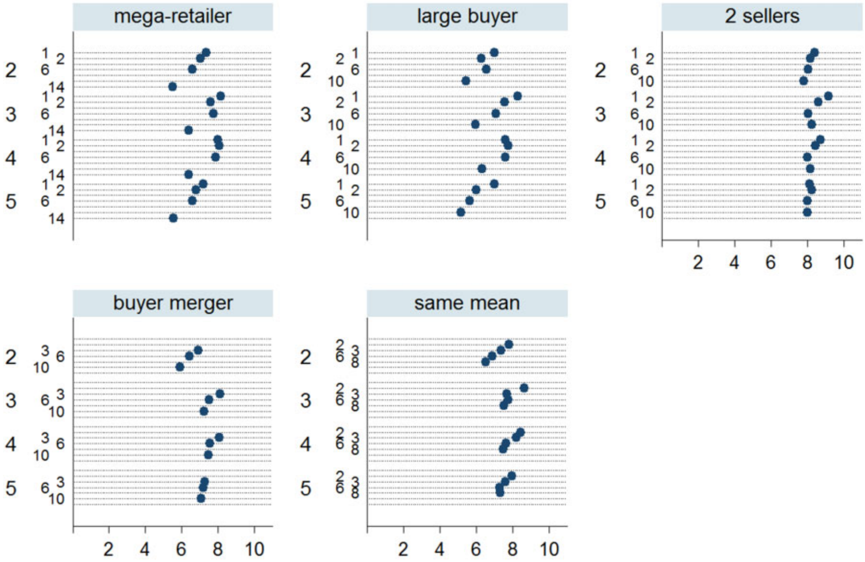

Figure 2

Mean Offers by Supergame and Buyer Size for Each Treatment

Notes: Mean offer for supergames 2–5 for each buyer size and treatment. Labels of adjacent buyer sizes are spaced alternately.

Except for the basic offer regression in (13) in which round and game controls are absent, offers and prices to buyer sizes 6 and 8 are not significantly different from one another in any of the other regressions (including the analogous price regression not reported).[20] The fitted offer probabilities based on the ordered probit estimates in (14) reinforce the lack of a largest-buyer discount with respect to the second-largest buyer: the offer probabilities differ by no more than 1.7 percentage points for any of the three available prices. As in previous treatments, the smallest buyers of sizes 2 and 3 in this treatment pay significantly more than their larger counterparts.

V(ii). *Buyer-Size Discounts over Time*

Thus far, the data analysis has focused on comparing prices to buyers of different size over the entire experiment. The panel nature of the design has yet to be exploited to address price differences between buyer sizes over time. For supergames in which each buyer size in the distribution appeared at least once (games 2–5, see Table 1), Figure 2 displays the mean offer by

[20] In a random-effects GLS regression parallel to (13), the large-buyer discount also is not significant.

supergame to each buyer size, with a separate graph for each of the five treatments. A number of striking observations emerge from the figure. First, for the *mega-retailer* treatment in which a significant large-buyer discount was predicted and observed, this discount is present already in game 2 and persists through game 5. Second, in *mega-retailer*, although prices decrease monotonically with buyer size (in contrast to the theory), the discount to the largest buyer of size 14 relative to the second-largest buyer appears to be substantially larger than the corresponding discount to any other buyer size in all of the games (in congruence with the theory). Third, in *2 sellers*, *buyer merger* and *same mean*, whatever scant evidence exists of a large-buyer discount in game 2 vanishes entirely by games 4 and 5. The *2 sellers* treatment is particularly notable for the degree of coordination sellers achieve on the same average price for all four buyer sizes.

The random-effects regressions reported in Table VI provide statistical support for the above observations and uncover additional ones.[21] In *mega-retailer*, the size of the large-buyer discount remains unchanged and highly significant throughout games 2 to 5 (0.96 units less than the second-largest buyer in games 2 and 3 and 0.98 units less in games 4 and 5, $p < .01$ in both cases). In *large buyer*, the magnitude of the discount (not predicted by the theory) is halved and only marginally significant ($p = .10$) in games 4 and 5. In the remaining three treatments in which no large-buyer discount was predicted, the relevant point estimates are close to zero and not significantly different from zero in games 2 and 3 nor in games 4 and 5.

Additional evidence of learning in the direction of the theory stems from the last two rows of Table VI. In 4/5 treatments, the point estimates for the small-buyer premium fall from games 2 and 3 to games 4 and 5. What is more, in *2 sellers* and *buyer merger*, the small-buyer premium shrinks to the point of no longer being significant.

Overall, the findings from this subsection establish further support for the theory. The small-buyer premia display signs of weakening over games. In treatments in which the large-buyer discount is not predicted, it is either entirely absent or shrinks substantially over time (*large buyer*). Only in the treatment in which the large-buyer discount is predicted (*mega-retailer*), does it appear robust and highly significant throughout.

VI. ALTERNATIVE EXPLANATIONS

In this section, I explore the extent to which other explanations are consistent with this paper's observed findings. To begin, my experimental design excludes more familiar explanations for buyer-size discounts, such

[21] Because the interest here is in the significance of the buyer-size effects and whether these effects change over time (i.e., interaction terms), I report random-effects regressions for ease of interpretation. The identical qualitative results hold for ordered probit regressions.

TABLE VI
RANDOM-EFFECTS REGRESSIONS TO TEST BUYER-SIZE DISCOUNTS ACROSS GAMES

regressor label	treatment regressor \ equation	*mega-retailer* (16)	*large buyer* (17)	*2 sellers* (18)	*buyer merger* (19)	*same mean* (20)
a	2nd BS	0.96***	0.91***	0.11	–0.03	0.16
	(6/6/6/6/6)	(0.31)	(0.34)	(0.09)	(0.22)	(0.11)
b	3rd BS	1.36***	1.05***	0.32***	0.65**	0.54**
	(2/2/2/3/3)	(0.34)	(0.43)	(0.09)	(0.28)	(0.24)
c	4th BS	1.77***	1.72***	0.69***	—	1.17***
	(1/1/1/–/2)	(0.36)	(0.33)	(0.10)		(0.27)
d	1st BS*games45	0.05	0.29	–0.14	0.36	0.34
		(0.39)	(0.33)	(0.10)	(0.27)	(0.25)
e	2nd BS*games45	0.08	–0.18	–0.31***	0.49**	0.14
		(0.38)	(0.37)	(0.12)	(0.22)	(0.28)
f	3rd BS*games45	0.09	0.31	–0.23*	–0.05	0.21
		(0.34)	(0.31)	(0.12)	(0.28)	(0.27)
g	4th BS*games45	0.05	0.30	–0.46	—	–0.01
		(0.35)	(0.29)	(0.15)		(0.22)
	round	–0.003	–0.021	0.011***	0.011	0.007
		(0.018)	(0.16)	(0.004)	(0.015)	(0.012)
	Part 1 subject control	Yes	Yes	Yes	Yes	Yes
	Obs	1260	1260	1680	1386	1638
	R^2	.11	.07	.07	.03	.06
linear combo. of regressors	interpretation	*mega-retailer*	*large buyer*	*2 sellers*	*buyer merger*	*same mean*
a + e – d	large-buyer (1st BS) discount	0.98***	0.44*	–0.06	0.10	0.05
	wrt 2nd BS in games 4, 5	(0.32)	(0.26)	(0.07)	(0.07)	(0.09)
b – c	small-buyer (4th BS) premium	–0.41***	–0.67***	–0.37***	–0.68***	–0.32***
	wrt 3rd BS in games 2, 3	(0.13)	(0.21)	(0.06)	(0.20)	(0.11)
b + f – (c + g)	small-buyer (4th BS) premium	–0.37***	–0.65**	–0.14	–0.14	–0.62***
	wrt 3rd BS in games 4, 5	(0.22)	(0.33)	(0.18)	(0.13)	(0.18)

Notes: [1]Dependent variable: seller *i*'s period *t* offer.

[2]Regressors: 2nd BS, 3rd BS and 4th BS: indicators for the 2nd largest, 3rd largest and 4th largest buyer sizes, respectively, in each treatment (largest buyer omitted); actual buyer sizes appear in parentheses below the variable name, ordered by treatment. 1st BS*games45: indicator for 1st largest buyer size in the treatment interacted with an indicator for repeated games 4 and 5; similar interactions follow for other buyer sizes in the distribution; *round*: the overall round; *Part 1 subject control*: the number of rounds in which seller *i* chose the monopoly price in Part 1.

[3]Standard errors in parentheses are heteroskedasticity robust and clustered by seller.

[4]Linear combinations of regressors to evaluate the large-buyer discount and small-buyer premium over time are reported below the regression estimates. For the *buyer merger* treatment (with only 3 buyer sizes), the small-buyer premium computations are based on a – c and a + e – (b + f), respectively.

[5]Coefficient significantly different from 0 at the 1% level ***, at the 5% level **, at the 10% level *.

as seller cost structure—all sellers in my setup have linear costs, which are normalized to zero; and large-buyer bargaining power—all buyers are computerized and act as price-takers. In addition, risk preferences cannot account for the observed large-buyer discounts in the first two treatments and the range of risk coefficients consistent with the negation of buyer-size discounts in the remaining three treatments is too wide (from modest risk-lovingness to infinite risk aversion) to be informative.

Can the mere labeling of buyers by quantity demanded account for the results? The ubiquity of quantity discounts, even in retail industries, makes them familiar to subjects. With the preconception that larger buyers pay lower prices, sellers set their prices accordingly. Alternatively, sellers may aim to reach some satisficing profit (Simon [1955]) or minimum acceptable profit for each buyer size. As a result, they tolerate lower prices for larger, more profitable buyers and aim to keep prices higher for smaller buyers.[22]

Both the labeling and the aspiration-profit explanations predict that prices decline monotonically in buyer size, regardless of the number and distribution of buyer sizes in the market. The last three treatments contradict these explanations for the largest buyers. Aspiration profits produce additional inconsistencies in comparing prices across treatments. For example, for a given price level, duopoly sellers earn more than triopolists, implying the former ought to accept lower prices than the latter. Yet offers and prices are highest in the *2 sellers* treatment. On the other hand, labeling and aspiration profits are at least consistent with the small-buyer premia, whereas these premia are not anticipated by R-S and Snyder.

Recall that according to R-S and Snyder, large-buyer discounts or the absence thereof are predicated on infinitely repeated play between sellers. One may question the necessity of infinitely repeated play for the theory to hold in practice. For example, perhaps even in a one-shot or finitely repeated game with a known terminal round, individuals forego undercutting if the potential profit from doing so is small (because the buyer is not too large). Sellers reserve undercutting for sufficiently large buyers.[23] Such a theory does not rely on infinitely repeated play and could be made consistent with the observed large-buyer discounts or their absence in each of the five experimental treatments if correctly parameterized.

To evaluate this alternative explanation, I examine pricing by buyer size in Part 1 of the experiments. Table VII reports mean offers and prices over the 15 rounds of play in Part 1 for all buyer sizes in each treatment. The most striking feature of this table is the presence of a substantial

[22] Applying aspiration profits to an I.O. context, Cyert and March [1956] present empirical evidence that firms are guided by some acceptable profit norm. More recently, Huck *et al.* [2007] find that, in violation of the merger paradox, bilateral mergers in an experimental Cournot setting are not unprofitable for the merged firm. A series of follow-up treatments reveals that aspiration profits account for this result.

[23] I am indebted to the Editor for suggesting this explanation.

TABLE VII
SUMMARY STATISTICS FOR PART 1

	mega-retailer		*large buyer*		*2 sellers*		*buyer merger*		*same mean*	
Buyer Size	Offer	Price	Offer	Price	Offer	Price	Offer	Price	Offer	Price
1	7.2	5	6.6	4.5	8.43	7.5	—	—	—	—
	(2.74)	(3.03)	(3.38)	(3.12)	(2.08)	(2.45)				
2	7.58	6.6	7.2	5.8	8.37	7.93	—	—	7.33	5.85
	(2.21)	(2.55)	(2.52)	(2.91)	(1.49)	(1.76)			(2.63)	(3.10)
3	—	—	—	—	—	—	7.54	6.18	7.18	6
							(2.28)	(2.77)	(2.35)	(2.84)
6	6.35	5.36	5.33	4.04	7.63	7.3	6.41	5.51	5.83	4.95
	(2.84)	(2.99)	(3.11)	(2.85)	(1.73)	(1.97)	(2.80)	(2.96)	(3.07)	(3.13)
8	—	—	—	—	—	—	—	—	5.37	4.37
									(3.13)	(3.02)
10	—	—	3.93	2.72	7.01	6.56	5.31	4.4	—	—
			(2.89)	(1.96)	(2.34)	(2.57)	(3.09)	(2.95)		
14	4.28	3.2	—	—	—	—	—	—	—	—
	(3.03)	(2.41)								

Notes: Mean offers and prices for Part 1 by treatment and by buyer size. Standard deviations in parentheses.

large-buyer discount in both offers and prices in all five treatments. In fact, regressions parallel to those in Tables 4–5 (available upon request) reveal that the largest buyer pays significantly less than the second-largest buyer (and all smaller buyers) in each of the respective treatments. This pertinacious large-buyer discount contradicts the more nuanced appearance and disappearance of the discount observed in the indefinitely repeated games.[24] We conclude that indefinitely repeated play and the accompanying threat of lengthy punishment are essential for deterring lower prices to buyers who aren't sufficiently large, such as the largest buyers in *2 sellers*, *buyer merger* and *same mean*.

VII. CONCLUSION

In this paper, I design a number of experimental treatments to evaluate a plausible and often-cited explanation for large-buyer discounts. According to this explanation and the accompanying theory, buyers pay less if they are sufficiently large relative to other buyers in the industry and there exists sufficient competition between sellers for their business. My experiments test this theory by varying the distribution of buyer sizes in an industry, the relative size of the largest buyer and the number of competing sellers in ways that either create or negate seller incentives to charge less to the largest buyer.

The results mostly support the comparative-statics predictions: seller pricing responds to subtle changes in the demand structure and number of sellers along the lines of the theory. Large buyers obtain discounts where predicted. To nullify the discount, the experiments elucidate several possibilities such as fewer sellers, the merger of smaller buyers and shrinking the largest buyer while maintaining the average buyer size in the industry. Moreover, an analysis of buyer-size discounts over time reveals that the above findings (all consistent with the theory) appear early on and remain unchanged for the duration of the experiment. At the same time, observed deviations from the theory (i.e., small-buyer premia and an unanticipated large-buyer discount in the *large buyer* treatment) diminish or even disappear entirely over time.

There are two senses in which the results underestimate the presence and magnitude of large-buyer discounts. First, the observed negation of the large-buyer discount centers on a comparison between the largest and the second-largest buyers. Compared to other smaller buyers in the market, the largest buyer pays significantly less in most of the treatments studied.

[24] For the first three treatments, the table also reveals a curious smallest-buyer discount. This finding is due to an endgame effect: the smallest buyer of size 1 appears in the terminal round in these treatments. Excluding round 15 from the analysis, prices that decline monotonically with buyer size are restored.

Second, these experiments focus on a single explanation, controlling for other explanations possibly present outside the laboratory. To the extent that large buyers also benefit from a credible threat of backward integration, supply-side scale economies, experience and superior negotiating skills, the price discrimination found here will be amplified.

REFERENCES

Abbink, K. and Brandts, J., 2009, 'Collusion in Growing and Shrinking Markets: Empirical Evidence from Experimental Duopolies,' in Hinloopen, J. and Normann, H.-T. (eds.), *Experiments and Antitrust Policy* (Cambridge University Press, Cambridge, England).

Andreoni, J. and Croson, R., 2008, 'Partners versus Strangers: The Effect of Random Rematching in Public Goods Experiments,' in Plott, C. R. and Smith, V. L. (eds.), *Handbook of Experimental Economics Results* (North Holland Press, Amsterdam, The Netherlands).

Bagwell, K. and Staiger, R. W., 1997, 'Collusion over the Business Cycle,' *RAND Journal of Economics*, 28, pp. 82–106.

Bergahl, M., 2004, *What I Learned from Sam Walton: How to Compete in a Wal-Mart World* (John Wiley & Sons, Hoboken, New Jersey, U.S.A.).

Chipty, T. and Snyder, C. M., 1999, 'The Role of Firm Size in Bilateral Bargaining: A Study of the Cable Television Industry,' *Review of Economics and Statistics*, 81, pp. 326–340.

Cyert, R. M. and March, J. G., 1956, 'Organizational Factors in the Theory of Oligopoly,' *Quarterly Journal of Economics*, 70, pp. 44–64.

Ellison, S. F. and Snyder, C. M., 2010, 'Countervailing Power in Wholesale Pharmaceuticals,' *Journal of Industrial Economics*, 58, pp. 32–53.

Engle-Warnick, J. and Slonim, R. L., 2006, 'Learning to Trust in Indefinitely Repeated Games,' *Games and Economic Behavior*, 54, pp. 95–114.

Feinberg, R. M. and Husted T. A., 1993, 'An Experimental Test of Discount-Rate Effects on Collusive Behavior in Duopoly Markets,' *Journal of Industrial Economics*, 41, pp. 153–160.

Feinberg, R. and Snyder, C. M., 2002, 'Collusion with Secret Price Cuts: An Experimental Investigation,' *Economics Bulletin*, 3, pp. 1–11.

Fischbacher, U., 2007, 'z-Tree: Zurich Toolbox for Ready-Made Economic Experiments,' *Experimental Economics*, 10, pp. 171–178.

Fishman, C., 2003, 'The Wal-Mart You Don't Know,' *Fast Company*, December: 77, pp. 68.

Gallego, G. and van Ryzin, G., 1994, 'Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons,' *Management Science*, 40, pp. 999–1020.

Haltiwanger, J. and Harrington Jr., J. E., 1991, 'The Impact of Cyclical Demand Movements on Collusive Behavior,' *RAND Journal of Economics*, 22, pp. 89–106.

Harrison, G. W. and Rutström, E., 2008, 'Risk Aversion in the Laboratory,' in Cox, J. C. and Harrison, G. W. (eds.), *Research in Experimental Economics: Risk Aversion in Experiments*, 12, (Emerald Publishing, Bingley, West Yorkshire, England).

Huck, S.; Normann, H.-T. and Oechssler, J., 2004, 'Two Are Few and Four Are Many: On Number Effects in Cournot Oligopoly,' *Journal of Economic Behavior and Organization*, 53, pp. 435–446.

Huck, S.; Konrad, K. A.; Müller, W. and Normann, H.-T., 2007, 'The Merger Paradox and Why Aspiration Levels Let it Fail in the Lab,' *Economic Journal*, 117, pp. 1073–1095.

Inderst, R. and Wey, C., 2003, 'Bargaining, Mergers, and Technology Choice in Bilaterally Oligopolistic Industries,' *RAND Journal of Economics*, 34, pp. 1–19.

Isaac, R. M. and Reynolds, S. S., 2002, 'Two or Four Firms: Does It Matter,' in Holt, C. and Isaac, R. M. (eds.), *Research in Experimental Economics*, 9 (JAI-Elsevier, Amsterdam, The Netherlands).

Kandori, M., 1991, 'Correlated Demand Shocks and Price Wars during Booms,' *Review of Economic Studies*, 58, pp. 171–180.

List, J. A., 2009, 'The Economics of Open Air Markets,' NBER working paper, available at http://www.nber.org/papers/w15420.

Normann, H.-T.; Ruffle, B. J. and Snyder, C. M., 2007, 'Do Buyer-Size Discounts Depend on the Curvature of the Surplus Function? Experimental Tests of Bargaining Models,' *RAND Journal of Economics*, 38, pp. 747–767.

Rojas, C., 2012, 'The Role of Information and Monitoring on Collusion,' *RAND Journal of Economics*, 43, pp. 78–109.

Rotemberg, J. and Saloner, G., 1986, 'A Supergame-Theoretic Model of Business Cycles and Price Wars during Booms,' *American Economic Review*, 76, pp. 390–407.

Roth, A. K. and Murnighan, J. K., 1978, 'Equilibrium Behavior and Repeated Play of the Prisoner's Dilemma,' *Journal of Mathematical Psychology*, 17, pp. 189–198.

Ruffle, B. J., 2009, 'Buyer Countervailing Power: A Survey of the Theory and Experimental Evidence,' in Hinloopen, J. and Normann, H.-T. (eds.), *Experiments and Antitrust Policy* (Cambridge University Press: Cambridge, England).

Simon, H. A., 1955, 'A behavioral model of rational choice,' *Quarterly Journal of Economics*, 69, pp. 99–118.

Snyder, C. M., 1996, 'A Dynamic Theory of Countervailing Power,' *RAND Journal of Economics*, 27, pp. 747–769.

Snyder, C. M., 1998, 'Why Do Large Buyers Pay Lower Prices? Intense Supplier Competition,' *Economics Letters*, 58, pp. 205–209.

Snyder, C. M., 2008, 'Countervailing Power,' in Durlauf, S. N. and Blume, L. E. (eds.) *The New Palgrave Dictionary of Economics* 2nd ed. (Palgrave Macmillan, Houndmills, Basingstoke, Hampshire, England).

Stole, L. A. and Zwiebel, J., 1996, 'Organizational Design and Technology Choice under Intrafirm Bargaining,' *American Economic Review*, 86, pp. 88–102.

Walton, S. and Huey, J., 1993, *Sam Walton: Made in America* (Bantam Books, New York, New York, U.S.A.).

APPENDIX

INSTRUCTIONS FOR *MEGA-RETAILER* TREATMENT

Participant Instructions—Part 1

Introduction. This is an experiment in the economics of market decision-making. Funds for this experiment have been provided by various research foundations. Take the time to read carefully the instructions. A good understanding of the instructions and well thought out decisions during the experiment can earn you a considerable amount of money. All earnings from the experiment will be paid to you in cash at the end of the experiment.

In this experiment, we are going to create numerous markets. All participants in these markets act as sellers. In each market, there will be 3 sellers of a fictitious commodity. You will participate as a seller along with 2 other sellers in your market. There are two parts to this experiment. In the first part of the experiment, each market will be repeated with the same group of sellers for 15 periods.

Random Buyer Size and Seller Pricing. In each period, you and the other 2 sellers will face one computerized buyer of a randomly determined size. The buyer's size is drawn from a uniform distribution from the values 14, 6, 2 and 1. So, there is an equal chance (1/4) that the buyer will be of size 14, of size 6, of size 2 and of size 1. The buyer's size indicates the number of units the buyer is interested in purchasing. Thus, for example, a buyer of size 14 wishes to purchase 14 units, while a buyer of size 1 seeks to purchase a single unit.

As a seller, you and the other 2 sellers will simultaneously choose a price each period. Each seller privately selects one of the following 3 prices: 10, 8, 2. If you all choose the same price, then you share equally the number of units demanded by the buyer. Your profit in this case is determined by the number (or fraction of the number) of units supplied, multiplied by the price. If you choose different prices, then the seller (or sellers) who selected the lowest price supplies the number of units demanded by the buyer. The profit of the lowest-priced seller (or sellers) is again given by the number (or fraction of the number) of units supplied multiplied by the chosen price. Note that sellers do not incur any cost for the units they supply: their costs of production are zero.

Example. To illustrate this game, consider the following example with a buyer of size 6. This means the buyer seeks to purchase 6 units and will pay whatever price you charge. If you and the other 2 sellers all choose the same price, then you each supply 6/3 = 2 units at the price selected. Thus, you each earn 2 units times the price selected. If all of the prices are not the same, then the seller with the lowest price supplies all 6 units at the price he selected. The other higher-priced sellers each earn 0. If it is the case that more than one seller chose this same lowest price, then these sellers split the 6 units equally between them at the chosen price.

Each Period. You will observe the outcome of your price decision at the end of each period. That is to say, you will observe the number of units you supplied to the buyer, your profit for the period and the other 2 sellers' chosen prices.

In each of the 15 periods, you and the same 2 other sellers will face the identical set of 3 prices from which to choose. The only change between periods is the randomly determined size of the buyer you face together: in each period, the number of units the computerized buyer seeks to purchase is chosen with equal probability from the sizes 14, 6, 2, 1.

Payments. At the completion of the second part of experiment, you will be paid your cumulative earnings from all 15 rounds. Your earnings will be paid in cash at a rate of 1 shekel for every 15 units of profit you accumulate.

Before Beginning the Experiment. Before beginning the experiment, the experimenter will read aloud the instructions. Next, you will also be asked to answer a number of questions on the computer to confirm that you have understood the rules of the experiment. Finally, you will participate in 5 practice rounds to allow you to familiarize yourself with the software and improve further your understanding of the experiment. The profit from these practice rounds will not count toward your cash payment from the actual experiment. Otherwise, the rules from the 5 rounds are identical to those of the experiment. At the end of the 5 practice rounds, you will be randomly rematched with 2 different sellers with whom you will participate in the 15 rounds of play.

If at any stage you have any questions about the instructions, please raise your hand and a monitor will come to assist you.

Participant Instructions—Part 2

This part of the experiment is identical to Part 1, except that instead of playing one game consisting of 15 rounds, you will play a number of games one after another, each game of randomly determined length.

To repeat the basic rules, once again you will be matched with 2 other sellers. You will face a buyer of randomly determined size in each round. The buyer's size (meaning the number of units the buyer wishes to purchase) is either 14, 6, 2 or 1. Each of these four buyer sizes occurs with an equal probability of 1/4. At the beginning of each round, you will be shown the buyer's size and need to choose a price 10, 8 or 2. Your profit in a round is determined by your price, its relation to the other 2 sellers' prices and the size of the buyer. The seller with the lowest price supplies all of the buyer's demand at the price selected. If more than one seller chooses the same price, then they share equally the demand at the price chosen. Any seller that chooses a higher price earns 0 for that round.

Number of Periods. In the first game, you will again be randomly matched with 2 sellers. The number of periods in which you will participate with this same pair of sellers will be randomly determined. At the end of each period, a computer-generated random number between 1 and 100 will determine whether you proceed to the next period with the same set of sellers. If the random number is less than or equal to 80, you and the other 2 sellers will play another period together. You will continue to play together for additional periods as long as the randomly generated number at the end of the period is less than or equal to 80.

If, at the end of a given period, the random number is greater than 80, you will cease to participate in this market and, time permitting, begin to participate in an identical market with the sole change being the other 2 seller participants. The rules in this new market are the same as the previous one (i.e., same number of sellers, randomly determined buyer size each period, same set of prices to choose from); the only difference is that you will be matched in this new market with 2 sellers randomly selected from the group of participants with whom you have not participated in a previous market. You will continue to interact with this same new cohort of 2 sellers for as many periods as a random number less than or equal to 80 is drawn. If a number greater than 80 is drawn at the end of a period, then, time permitting, you will again be rematched to play in a third market with 2 new sellers.

In short, in any given round, there is a 0.8 probability that the game will continue to the next round and a 0.2 probability that the game will end at the completion of the current round. If, in fact, the game comes to an end, then you will be rematched with

2 sellers in the next game, time permitting, and continue to play in a market with them until that game ends. You will continue to play in a game until the random draw between 1 and 100 determines that the game ends. You will play as many games as time permits.

Payments. At the completion of the experiment, you will be paid your cumulative earnings from all periods in all markets in which you participated. As in Part 1, each 15 units of profit you earn are worth 1 shekel in cash.

While the monitors are counting your cash payment, you will be asked to complete a questionnaire concerning the experiment that will appear on your computer screen.

If at any stage you have any questions about the instructions, please raise your hand and a monitor will come to assist you. Thank you for your participation.