5-2016

Competition, Cost Innovation, and X-inefficiency in Experimental Markets

Andrew Smyth
Marquette University, andrew.smyth@marquette.edu

Review of Industrial Organization, Vol. 48, No. 3 (May 2016) : 307-331. DOI.

Competition, Cost Innovation, and X-inefficiency in Experimental Markets

Andrew Smyth

Economic Science Institute, Chapman University, One University Drive, Orange, CA

Abstract

This paper examines the relationship between competition, cost innovation, and x-inefficiency in experimental markets. In the lab, oligopolists make closer-to-optimal cost innovation expenditures than do monopolists, which result in lower x-inefficiency in oligopoly than in monopoly. Oligopolies also increase total surplus relative to monopoly, and consumer surplus makes up a larger portion of total surplus in oligopoly than monopoly. The data illustrate how x-inefficiency affects surplus dynamically and suggest price as a mechanism by which competitive pressure increases cost efficiency.

Keywords

X-inefficiency, Cost innovation, Experimental economics

1 Introduction

X-inefficiency theory maintains that firms may not minimize their costs of production, especially in markets where they experience little competitive pressure.[1] To date, there have been no attempts to study x-inefficiency using experimental methods.[2] While there are always questions of external validity with experimental research, the laboratory offers an excellent test bed for x-inefficiency theory because the theory posits a behavioral response to external stimuli—an increase in productive efficiency in response to market competition. In the lab, competitive stimuli can be precisely applied, and any behavioral response that affects efficiency can be readily observed.

In these experimental markets, subjects participated in a repeated, two-stage game. In the first stage, they decided how much to spend on cost innovation. Innovation was stochastic, with their probability of achieving a lower production cost increasing in their expenditure. In the second stage, they made pricing and capacity decisions given their realized production cost. This second, ''market'' stage was a monopoly, or a duopoly, or a quadopoly.

X-inefficiency theory typically assumes that the cost frontier is achievable—if management makes necessary expenditures. In these experiments, as in actual markets, cost innovation was stochastic. Chance dictated that subjects could not always achieve the exogenous cost frontier, so x-inefficiencies are measured relative to optimal innovation paths that are calculated for each market structure.

The data show that subjects in more competitive market structures attempted closer-to-optimal cost innovation. They support the hypothesis that competitive pressure reduces x-inefficiency relative to monopoly. More competitive market structures also increased total surplus and the percentage of total surplus that is consumer surplus, relative to monopoly. Finally, the experiments illustrate how x-inefficiency affects surplus dynamically.

Outside the lab, competitive pressure may limit opportunities for afternoon tee times, opulent office decor, or for using obsolete production methods. In these experiments, differential pricing pressure alone created differences in cost efficiency for monopolies relative to more competitive markets. Thus, the data highlight price as a mechanism through which competitive pressure curbs x-inefficiency.

2 The Experiments

2.1 Motivation

Though no experimental study can incorporate all of the nuance of naturally-occurring markets, the hypothesis that competition reduces x-inefficiency is general, and so if valid, should hold in stylized experimental markets (Plott 1989). Relative to studies that employ naturally-occurring data, the laboratory is a fruitful venue for examining the relationship between competition and cost efficiency. Because demand and costs are known, the primary variables of interest do not have to be estimated. Furthermore, the experimental markets are exogenously switched from monopoly to oligopoly (and vice versa), which allow for clean causal inference of the effect of competition on cost efficiency.

The experimental literature on cost innovation is modest (see for example, Isaac and Reynolds 1992; Darai et al. 2010; Sacco and Schmutzler 2011; Aghion et al. 2014). My design is similar to Isaac and Reynolds (1992), but with several key differences: Isaac and Reynolds employ a repeated, two-stage game to examine competition in both cost innovation and price. They seek an experimental ''existence proof'' of innovation competition akin to that described by Joseph Schumpeter, and report more aggregate expenditure on cost innovation in quadopoly than monopoly.

Isaac and Reynolds's experimental firms have non-constant marginal costs; here, marginal costs are constant. X-inefficiency theory does not depend on this heroic assumption, but the functional form ensures that suboptimal capacity and x-inefficiency are not conflated. There are different capacity constraints across treatments in Isaac and Reynolds (10 units in their monopolies; 5 units in their quadopolies). Here, all firms have a maximum capacity of 8 units regardless of market size, which allows for clean comparisons across treatments. In Isaac and Reynolds, there are 10 market periods with cost innovation; in these experiments, there are 30 such periods.

Finally, and most important, Isaac and Reynolds vary competition between-subjects. Here, competition is varied both between and within subjects. The within-subject variation is motivated by Leibenstein (1966), who gives examples of x-inefficiencies increasing in firms after the departure of management consultants. He implies that successful consultants limit x-inefficiency, but that x-inefficiency rises once they depart. In one set of treatments, subjects are exogenously switched from monopoly to oligopoly to monopoly, to see if the pattern of x-inefficiency tracks the pattern in Leibenstein's anecdote.

2.2 Design

These experiments examine cost efficiency across several market structures. The experimental subjects acted as firm managers. Each period, for a number of periods, they made cost innovation expenditure decisions. Then, given realized costs of production, they selected prices and capacities. Finally, sales of a homogeneous product were allocated according to a rationing rule. The key experimental parameters are presented in Table 1 and are discussed below.

Table 1 Key parameter values

Parameter	Value
Endowment	$5.00
Attempts	[0, 30]
Cost per attempt	$0.10
Prob(Innovation | 1 Attempt)	10 %
Prob(Innovation | *a* Attempts)	$1 - .90^{a}$
Initial production cost	$7.75
Innovation lowered cost by	$0.25
Price	[*c*, $20.00]
Production (sales) capacity	8

All prices and costs are in experimental dollars. Attempts and capacity choices are integers. Endowments were awarded and initial costs reset at the beginning of each block

The experiments lasted 30 periods. Each period was divided into two stages: an Innovation stage and a Market stage. In Innovations stages, subjects decided how many ''innovation processes'' (hereafter, ''attempts'') to purchase. They knew that they would subsequently be sellers in a market, and that if one of their purchased attempts was successful, their costs of producing units in the market would be reduced. But they also knew that attempts were costly, and that success was randomly determined. After deciding how many attempts to buy, subjects learned whether they were successful in lowering their production costs. Then they entered the Market stage and selected a price and a capacity. Finally, they learned their market sales and profit.

Prior to the start of the experiment, subjects were randomly assigned to a market and remained in that market for the entire experiment. The 30 periods were divided into three 10-period blocks. At the beginning of each block, subjects were endowed with $5.00 experimental dollars.[3] In each Innovation stage, subjects could purchase a ($0 \leq a \leq 30$) attempts. Each attempt cost $0.10 and carried a 10 % chance of success. If a subject was successful, her marginal cost of producing units in the current and all future periods in the same block was lowered by $0.25 per unit. Each cost reduction was 3 % of the marginal cost that the subjects faced at the start of a block.

Subjects could receive at most one cost reduction per period and could purchase at most 30 attempts during each Innovation stage (i.e., at the early part of each period).[4] Innovation expenditures were sunk; all *a* attempts were paid for regardless of whether all were necessary to achieve the cost reduction ex post. The probability of achieving the cost reduction was an increasing function of expenditure: $\theta(a) = 1 - 0.90^a$. Innovation success meant lower costs for the innovating seller only (there were no spillovers).

Following the Innovation stage, subjects participated as sellers in a posted-price market (the Market stage). Posted-price markets are analogous to many retail markets in that there is no haggling. All sellers produced a homogeneous product. Prior to period 1 of each block, they shared a constant cost function of \$7.75. In subsequent periods of the block, their cost functions differed (in level) depending on their various successes or failures in the preceding Innovation stages. Between blocks, all sellers' cost functions reset to \$7.75. Costs only applied to units sold: The product was ''made to order.''

The exogenous cost frontier for each period was the same in all treatments. This frontier was achieved if a subject successfully innovated each period—it was \$7.50 in period 1, \$7.25 in period 2, etc. Because it was unlikely that subjects would successfully innovate each and every period, the results in Sect. 2.4 focus on deviations from optimal innovation paths (explained in Sect. 2.3).

In each Market stage, sellers entered a price and a integer production capacity between 0 and 8 units. They could not price discriminate; their posted price applied to all units that they might sell. All sellers, regardless of treatment, shared the exogenous capacity-constraint of 8 units. While subjects could endogenously constrain their output further, it was a weakly dominant strategy always to choose a capacity of 8 units.[5]

The demand side of the market was automated. Subjects were not told the reservation prices of the ''robot'' buyers.[6] Each buyer had a unique reservation price. The highest was \$10.01, with each successive buyer's reservation price \$0.25 lower in value (i.e., \$9.76, \$9.51, \$9.26,...), giving demand a stairstep shape as in Fig. 1. The queue was not random. Buyers purchased in descending order of their reservation price.

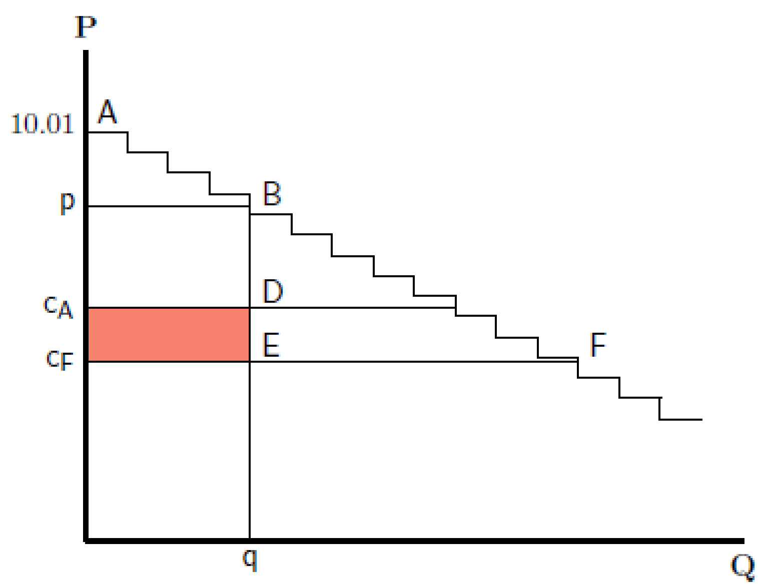

Fig. 1 Experimental market structure

There were monopoly, duopoly, and quadopoly markets. In monopolies, buyers bought one unit from a seller, provided that the seller's posted price was less than the buyer's reservation price. Monopolists were assured of selling 8 units, provided that they posted a price of \$8.26 or less. In duopolies and quadopolies, sellers competed for sales. The seller who posted the lowest price got to make sales first. She took the highest reservation price buyers out of the market. After the lowest price seller made her sales (up to 8), the next lowest price seller had an opportunity to make sales. In the event of a price tie, units were randomly allocated among the price-tied sellers. Thus, in the oligopolies, units of the homogeneous product potentially sold for different prices in the same market.

Table 2 summarizes the design, where n denotes the maximum market size. The ONE treatment was a monopoly baseline. In the TWO and FOUR treatments, subjects were monopolists in Block 1, and thereafter faced competition. In the TWOREV and FOURREV treatments, subjects shared a design history with the TWO and FOUR treatments through the completion of Block 2. Between Blocks 2 and 3, the subjects in these -REV (short for reversion) treatments were exogenously switched back to monopoly.

Table 2 Experimental design

Treatment	Sellers (*n*)	Periods		
		1–10	11–20	21–30
ONE	1	*Monopoly*	*Monopoly*	*Monopoly*
TWO	2	*Monopoly*	*Duopoly*	*Duopoly*
TWOREV	2	*Monopoly*	*Duopoly*	*Monopoly*
FOUR	4	*Monopoly*	*Quadopoly*	*Quadopoly*
FOURREV	4	*Monopoly*	*Quadopoly*	*Monopoly*

The reversion treatments were inspired by Leibenstein (1966). Their market structure pattern across blocks offer the absence (Monopoly), presence (Oligopoly), and absence (Monopoly) of pricing pressure that may cause a pattern of cost reductions that mimic the absence, presence, and eventual absence again of management consultants who, when present, help implement cost reductions.[7]

Subjects were monopolists in all treatments in Block 1. This established an almost constant baseline of experience. I say ''almost constant'' because when subjects in the n>1 treatments were monopolists, they received information after each period about the decisions of the other seller(s) in their market. This feedback was their market companions' innovation attempts, innovation outcomes (success or failure), posted prices, and sales. Subjects in ONE only received feedback about their own decisions and outcomes.[8]

The difference between the monopoly and oligopoly market structures was explained to the subjects in non-technical and context-free language. Subjects in the n>1 treatments knew that they were in a market with another (or three other) subject(s). They knew that market structure was fixed throughout a block, but could vary across blocks. The instructions were explicit that they would not be informed as to which market structure they faced in a given block. Thus, in each of the three blocks, subjects had to learn whether they faced competition for sales (oligopoly) or not (monopoly).[9] See the ''Appendix'' for the instructions.

2.3 Measurement and Optimal Innovation

An advantage of experimental data vis-a´-vis naturally-occurring data is that all of the variables of interest are precisely measurable. Figure 1 illustrates the experimental market structure. Let p and q denote the market price and quantity. Actual costs of production are c_A and the cost frontier is c_F. Maximum total surplus is the area Ac_FF. Actual total surplus is the area Ac_ADB, which can be divided into consumer surplus (ApB) and producer surplus (pc_ADB).

Unfortunately, x-inefficiency has been used casually in the literature to reference both a distance and an area. I follow Leibenstein (1978) and define x-inefficiency as a distance, or, ''the excess of actual over minimum cost for a given output.'' In Fig. 1, x-inefficiency is the distance c_A - c_F. The surplus loss from x-inefficiency is an area; it is a function of both x-inefficiency and output. In Fig. 1, it is the shaded area c_Ac_FED.[10]

In monopoly markets, an optimal pure strategy in price and capacity, {p*(c), q*}, exists. As the notation suggests, this optimal price is a function of the cost level. On the other hand, no pure strategy equilibria in price and capacity exist for the oligopoly markets. Consequently, I use ex post empirical profits to solve for optimal innovation paths instead of ex ante theoretical profits. Because cost reductions carry over from period-to-period, I generate these paths from a finite-horizon, dynamic optimization problem. For period t, the problem is:

$$\max_{a_t} \{\theta(a_t)[\pi_j(s_t - 0.25) + V(s_t - 0.25)] + [1 - \theta(a_t)][\pi_j(s_t) + V(s_t)] - ca_t\}.$$

(1)

The state variable, s_t, is the initial cost level in period t before any innovation attempts are purchased. The choice variable is the number of attempts, a_t. Let c be the cost of 1 attempt; let, $\theta(a) = 1 - 0.90^a$ be the probability that an innovation is successful when a_t attempts are purchased; and let $V(.)$ be the continuation value. Finally, π_j (.) is the observed, average profit for a given block, cost level, and market structure.[11]

The solution to (1) answers the question, ''If I know what average Market stage profit will be, how many attempts should I purchase at the Innovation stage?'' Table 3 lists the optimal innovation paths for each (Block 3) period and market structure—including an optimal path calculated with theoretical monopoly profit. The paths reveal that optimal innovation decreases in time, because earlier cost reductions are more profitable. Monopoly subjects should attempt the most innovation, followed by duopoly and quadopoly subjects. This is because oligopoly price competition can whittle away the gains from innovation that exist under monopoly (and, in the experimental set-up, volume for a seller—whether monopolist or oligopolist—was limited to eight units). Whether monopoly, duopoly, or quadopoly was the most cost-efficient market structure depended on the decisions of the incentivized subjects. The next section analyses their data.

Table 3 Optimal innovation paths

Period	*Monopoly* (Theory)	*Monopoly* (Block 3)	*Duopoly* (Block 3)	*Quadopoly* (Block 3)
1	28.0	28.0	22.0	21.0
2	27.0	27.0	21.5	19.0
3	25.7	26.0	20.0	17.3
4	24.5	24.8	16.5	15.0
5	23.2	23.2	15.0	13.4
6	21.2	21.7	14.0	11.5
7	19.0	19.3	12.1	9.7
8	16.5	16.5	10.1	7.5
9	12.4	13.0	6.4	6.2
10	5.8	5.7	2.4	2.5

The numbers in the table are optimal attempts. The profit measure used in optimization (theoretical or average profit) is in parenthesis

2.4 Data

2.4.1 Preliminaries

The experiments were conducted at a large, public research university in the United States. Subjects were recruited using the ORSEE recruitment system, and the experiment was executed in z-Tree (Greiner 2015; Fischbacher 2007). Subjects received US$10.00 for arriving at the computer lab on time. See Table 4 for mean earnings by treatment. All treatments lasted approximately two hours.

In the figures in this Data section, I pool the monopoly market data across treatments for Blocks 1 and 3.[12] I first present the time paths of market price and quantity to provide a picture of Market stage competition. Then I establish that oligopoly subjects attempted closer-to-optimal cost innovation than did their monopoly counterparts. In fact, oligopolists attempted significantly more innovation than did monopolists in Block 3 and were no ‘‘luckier’’ in achieving cost reductions.

Table 4 Key statistics

	ONE	TWO	TWOREV	FOUR	FOURREV
Subjects	12	16	18	12	16
Markets	12	8	9	3	4
Mean earnings[a]	\$36.01	\$23.11	\$26.99	\$19.77	\$24.23
Actual innovation success rate[b]					
Block 1 (%)	13	11	11	9	9
Block 2 (%)	10	10	10	10	10
Block 3 (%)	10	11	11	10	10

a Includes a US\$10.00 show-up fee. All amounts in this row are for US\$

b The parameterized stochastic success rate was 10 % per try

Using regression analysis, I show that oligopoly markets were significantly less x-inefficient than were monopoly markets. Finally, I report surplus comparisons across the market structures.

2.4.2 Results

Figure 2 graphs both share-weighted market price and average market quantity over time. In a given period, share-weighted market price is

$$\frac{1}{M}\sum_{m=1}^{M}\sum_{i=1}^{n_m} s_{im}\cdot P_{im}.$$

(2)

where m indexes markets and i indexes firms within market m, where $n_m \epsilon \{1, 2, 4\}$. Also, s_im is firm *i*’s market share, and P_{im} is that firm’s price. Average market quantity is calculated as market quantity, averaged across all monopoly, duopoly, and quadopoly markets.

A cost-efficient, profit-maximizing monopolist lowers her price over the course of a block as she achieves cost reductions. Figure 2 shows that actual monopoly subjects posted relatively constant prices. In Block 3, the monopoly price path even trends up over the first several periods. This reflects the fact that -REV subjects realized they were monopolists in Block 3 and raised their prices accordingly.

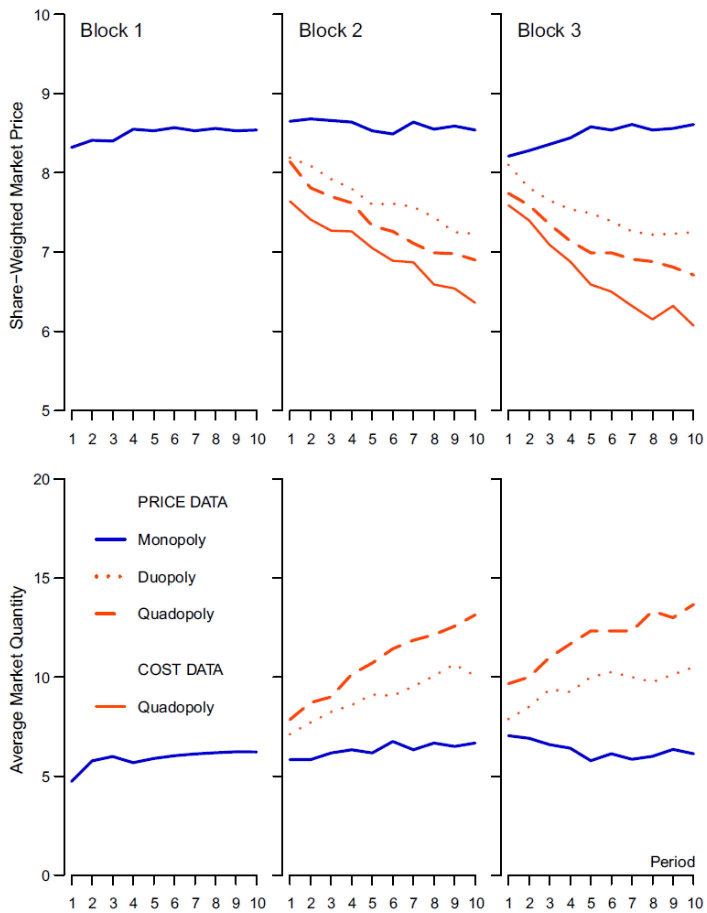

Fig. 2 Price and market quantity over time

Prices in the duopoly and quadopoly markets were markedly lower than prices in the monopolies. Moreover, the quadopolies had lower prices than the duopolies. To provide a clearer picture of pricing in the quadopolies, share-weighted marginal cost is graphed for these markets. A comparison of these cost data to the price path reveals that price roughly tracked cost. The market quantity data are very similar to the price data, only in inverse.

I now turn to cost innovation and the first result.

Result 1 Competition reduced the difference between actual and optimal innovation relative to monopoly.

Support For each block and market structure, optimal innovation paths were calculated as described in Sect. 2.3.[13] Figure 3 shows the deviations in actual attempts from these optimal paths. Subjects under-invested in innovation relative to the optimal path where the series are below 0; where they are above 0, they overinvested.

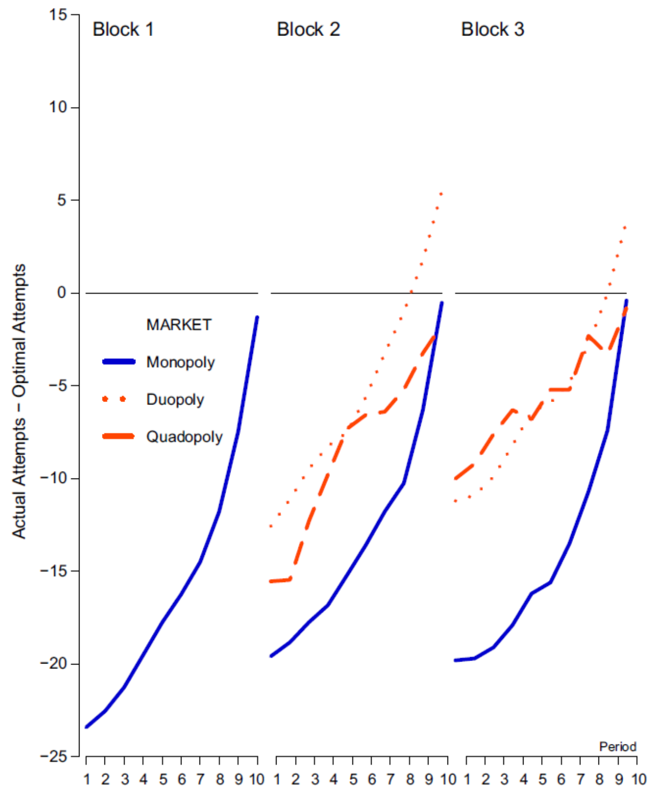

Fig. 3 Optimal innovation path deviations

Wilcoxon rank-sum tests were conducted for the null hypothesis that the deviation paths from different market structures are from populations with the same distribution.[14] In Block 2, the duopoly [quadopoly] deviation path is significantly different from the monopoly path ($p = 0{:}01$ [$p = 0{:}08$]). The deviation paths for the two oligopolies are not significantly different ($p = 0{:}33$). In Block 3, the duopoly [quadopoly] deviation path is also significantly different from the monopoly path ($p = 0{:}01$ [$p = 0{:}01$]). Likewise, the duopoly and quadopoly paths are not significantly different from one another ($p = 0{:}82$).

Subjects in all blocks and all market structures under-invested in cost innovation.[15] This general under-investment does not tarnish the result that subjects in more competitive market structures made closer-to-optimal innovation decisions. The optimal innovation problem here is complex enough to require time, knowledge, and computing power to solve precisely. It is suggested in the Conclusion that in complex environments, competitive prices can better direct firms toward optimal innovation than can monopoly prices.

As the next result shows, oligopolists in Block 3 not only attempted closer-to-optimal innovation than did monopolists, they even attempted more innovation.

Result 2 Competition increased expenditure on cost innovation relative to monopoly.

Support The top panel of Fig. 4 shows average attempts over time. This is calculated as the average number of innovation attempts purchased by a subject in a market, averaged across all monopoly, duopoly, and quadopoly markets. Ordinary least squares (OLS) regression estimates are presented in Table 5. The second and third columns report specifications with attempts as the dependent variable; one estimated with Block 2 data and another with Block 3 data. The variable Period is a linear time trend, and Duopoly and Quadopoly are indicator variables for the market structure. There are fewer than 740 observations (74 subjects x 10 periods) because only non-zero market share data were used in estimation. The coefficient estimates should be interpreted relative to data from the monopoly markets.

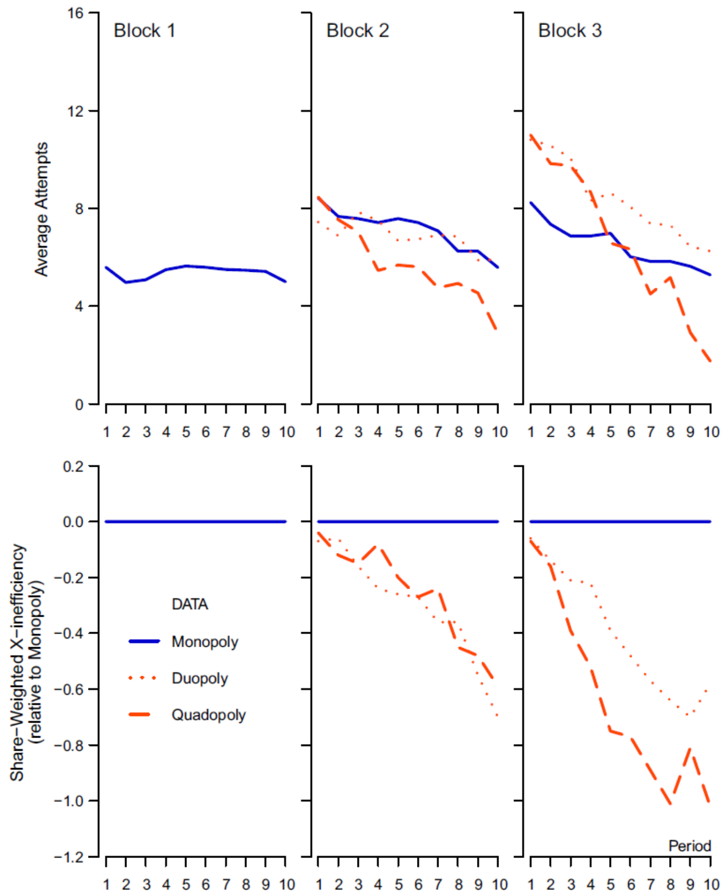

Fig. 4 Innovation attempts and x-inefficiency over time

The regressions indicate no treatment effect in Block 2, as no coefficient estimates are significantly different from zero (save for that on the constant). For Block 3, they suggest that competitive pressure initially increased the amount of attempted cost innovation. The intercept estimates for the duopoly and quadopoly markets are both significant and positive in sign. Note that the slope estimate (x Period) is not significant for the duopoly markets, but is both significant and negative in sign for the quadopoly markets. Thus, in the duopoly markets, the treatment effect persisted over time; in the quadopoly markets, it dissipated. The regressions confirm the visuals in Fig. 4 and indicate that in addition to being closer to the optimal innovation path, oligopolists actually attempted more innovation than did monopolists in Block 3.

It is important to establish that oligopoly subjects did not achieve more innovation success than did monopoly subjects by chance, or vice versa. The empirical innovation success rates are given in Table 4. Recall that the parameterized stochastic success rate was 10 % per try. Note that in Blocks 2 and 3, no treatments are outliers in terms of their actual success rate. Thus, subjects in the oligopoly markets were no ''luckier'' in achieving cost reductions than their monopoly counterparts.

In Block 3, attempted cost innovation in quadopoly started above monopoly, but finished below. For the whole block, there was little difference in average attempts between the quadopoly and monopoly markets (6.65 versus 6.49). However, the aggregate innovation data conceal important firm dynamics. Figure 5 is the Block 3 attempts graph from Fig. 4, with the quadopoly time series disaggregated. Specifically, the quadopoly series is broken up into a time path for quadopolists with a market share above 0.5, and another for those quadopolists with a market share below 0.5.

Figure 5 reveals that quadopoly subjects' innovation expenditures were heterogeneous. Subjects with little or no market share curtailed their innovation expenditures relative to monopolists, and so drove down average attempts in the latter periods of Block 3. On the other hand, large market share subjects invested in more cost innovation than did monopolists until the final two periods of Block 3.

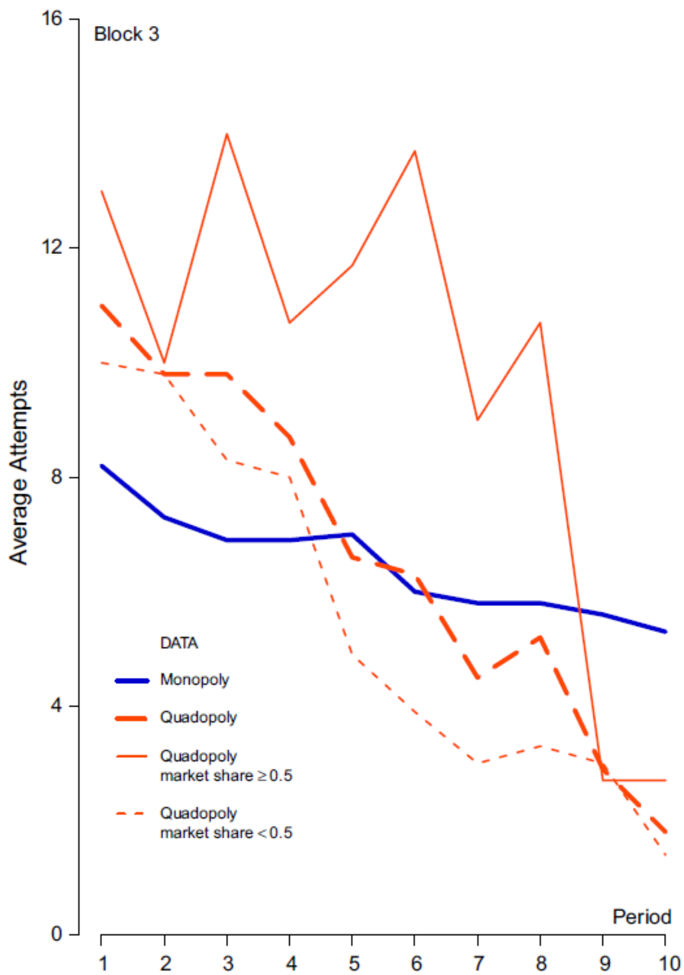

Fig. 5 Innovation attempts in Block 3

The preceding analysis can be summarized: In Blocks 2 and 3, subjects in the more competitive market structures made closer-to-optimal innovation expenditures. In Block 3, duopolists and high-market-share quadopolists actually attempted significantly more innovation than did monopolists. Finally, oligopoly subjects were no ''luckier'' in achieving cost innovations than were their monopolist counterparts.

I now examine whether the significant difference in cost innovation translated into a significant difference in x-inefficiency.

Result 3 Competition reduced x-inefficiency relative to monopoly.

Support X-inefficiencies are calculated relative to the optimal innovation paths, $\{a *_1 ; \ldots ; a *_{10}\}$. For each period and each market structure, the probability of achieving a cost reduction is $\theta(a_t^*) = 1 - .90a_t^*$. The cost frontier in period t relative to the optimal path is calculated according to $CF_t = 7.75 - \sum_t 0.25\,\theta(a_t^*)$.

The bottom panel of Fig. 4 shows share-weighted x-inefficiency over time. The time series are presented relative to x-inefficiency in the monopoly markets, so the monopoly time paths are all horizontal at 0.0. X-inefficiency is strikingly lower in the duopoly and quadopoly markets relative to the monopoly markets.

Table 5 contains regression results for x-inefficiency. The coefficient estimates should be interpreted relative to data from the monopoly markets. The results are consistent with Fig. 4. The estimates of the interaction term coefficients (x Period) are significant and negative in sign for both blocks. In Blocks 2 and 3, x-inefficiency fell by roughly 0.07 each duopoly period relative to monopoly. The estimates also suggest that x-inefficiency fell by 0.05 (0.11) in quadopoly relative to monopoly each Block 2 (Block 3) period.[16]

Table 5 OLS estimates for attempts and x-inefficiency

Dependent variable	Attempts		X-inefficiency	
Block	2	3	2	3
Constant	8.648***	8.130***		
	(1.830)	(0.781)		
Period	-0.266	-0.285***	0.107***	0.113***
	(0.231)	(0.058)	(0.020)	(0.011)
Duopoly	-0.930	3.340*	0.134***	0.039
	(2.063)	(1.821)	(0.026)	(0.054)
Duopoly x period	0.132	-0.227	-0.069***	-0.066***
	(0.246)	(0.163)	(0.024)	(0.018)
Quadopoly	-0.463	5.359**	0.094*	0.016
	(2.189)	(2.293)	(0.051)	(0.073)
Quadopoly x period	-0.029	-0.625**	-0.054**	-0.111***
	(0.265)	(0.283)	(0.023)	(0.016)
Observations	541	650	541	650
R^2	0.02	0.08	0.57	0.64

Clustered standard errors in parentheses (by market). Only non-zero market share data was used in estimation. Significant at the 1 % (***), 5 % (**), and 10 % (*) level

In sum, the analysis suggests that duopoly and quadopoly subjects generated significantly less x-inefficiency than did monopoly subjects. Moreover, quadopolies had less x-inefficiency than did duopolies in Block 3.

Surplus data are considered next.

Result 4 Competition increased total surplus relative to monopoly.

Support The top panel of Fig. 6 shows the time path of (actual) total surplus, averaged across all monopoly, duopoly, and quadopoly markets. Table 6 shows regression results for total surplus. The regressions confirm the visuals in Fig. 6. Oligopoly markets generated statistically significantly more surplus than the did monopolies in Blocks 2 and 3.

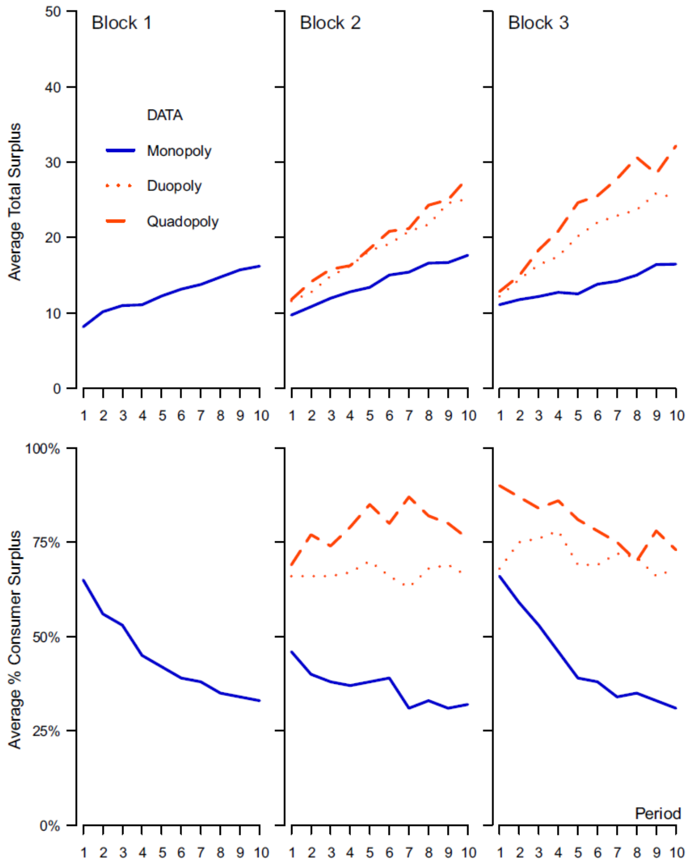

Fig. 6 Surplus and percent consumer surplus over time

The bottom panel of Fig. 6 shows the time path of the percent of total surplus that was consumer surplus. In terms of Fig. 1, it is the ratio of ApB to AcADB. The duopoly and quadopoly markets had a greater percentage of consumer surplus than the did monopolies, though the empirical percentage fell below 100 %. Across all three blocks, consumer surplus averaged 68 % of total surplus in oligopoly compared to 45 % in monopoly.

Table 6 OLS estimates for total surplus

Dependent variable	Total surplus	
Block	2	3
Constant	9.328***	10.366***
	(0.489)	(0.351)
Period	0.872***	0.627***
	(0.197)	(0.116)
Duopoly	0.757	1.499*
	(0.684)	(0.800)
Duopoly x period	0.661**	0.866***
	(0.278)	(0.202)
Quadopoly	0.870	1.541
	(0.938)	(1.015)
Quadopoly x period	0.832***	1.503***
	(0.257)	(0.158)
Observations	541	650

R^2	0.60	0.55

Clustered standard errors in parentheses (by market). Only non-zero market share data was used in estimation. Significant at the 1 % (***), 5 % (**), and 10 % (*) level

3 Conclusion

X-inefficiency exists in these experimental markets, and competition reduces it relative to monopoly. What is it about competitive pressure that spurs cost efficiency? The experiments suggest price as an answer. There has been much written on how the information contained in prices enables the efficient allocation of output. Prices also convey information that motivates the efficient production of output. Price competition grabs sellers' attention and signals firms that they must keep their costs minimized or else lose market share.

The experiments illustrate price competition as a mechanism that promotes cost efficiency. Subjects in the FOUR and FOURREV treatments came from the same student population, received identical instructions, interacted with uniform experimental software, and confronted the same experimental design prior to Block 3. But in that block, FOUR sellers competed with one another for sales while FOURREV sellers were monopolists. The difference in outcomes is stark.

Figure 7 shows the Block 3 time paths for share-weighted price and share-weighted marginal cost. In FOUR, prices fell across time; in FOURREV, they rose. Falling prices in FOUR pushed sellers to reduce their costs or have zero market share. There was no such pricing pressure to spur sellers to cost-innovate in FOURREV. Notice that FOURREV costs were approximately at the level of FOUR prices.

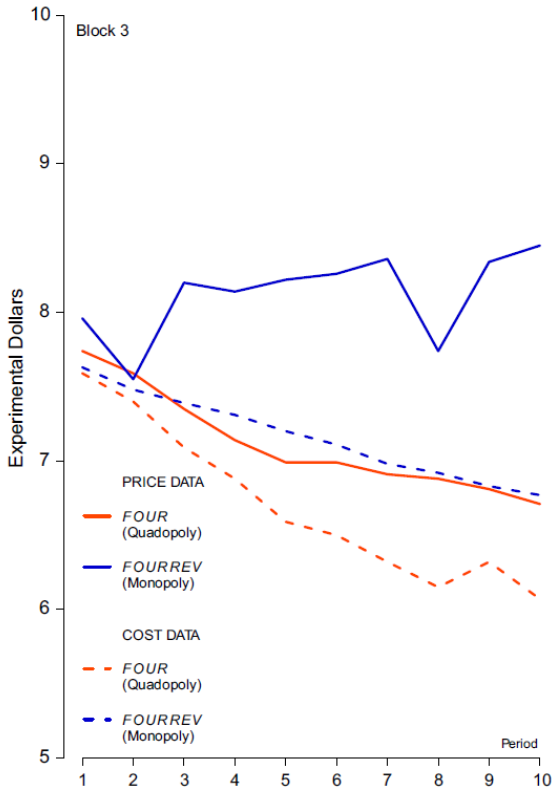

Fig. 7 Comparison of FOUR and FOURREV

Some critics of x-inefficiency theory argue that surplus losses from x-inefficiency are identical to surplus gains from producer leisure.[17] Even if one does not equate surplus losses from x-inefficiency with surplus gains from a ‘‘quiet life,’’ x-inefficiency may not be thought important because it does not diminish consumer surplus; it only ‘‘eats’’ into producer surplus. However, this is only true from a static perspective.

The experiments show how x-inefficiency affects surplus dynamically. Subjects used cost innovation differently in monopoly and oligopoly. In oligopoly, they had to cost-innovate to keep market share—the most x-efficient seller in each quadopoly (duopoly) had the largest market share 70 % (81 %) of the time. As they drove their costs down over time, more buyers were brought into the market, thus increasing total surplus. Because market price tracked the downward trajectory of costs, consumer surplus continued to be a substantial percentage of total surplus in the oligopoly markets. On the other hand, monopolists kept their prices relatively constant and used cost innovation to widen their profit margins. As a result, consumer surplus as a percentage of total surplus fell over time in monopoly.

This research is the first to test experimentally the hypothesis that competitive pressure reduces x-inefficiency. Data generated from the interaction of incentivized human subjects in laboratory markets strongly support the hypothesis. The experiments vividly demonstrate the ills of x-inefficiency, but also the power of competition to spur cost efficiency and lower prices and to promote surplus.

Acknowledgments

I am grateful to Mark Isaac, Taylor Jaworski, Stanley Reynolds, Cortney Rodet, and Bart Wilson for helpful comments and discussion and to Philip Brookins and John Jensenius for programming advice. I also thank my dissertation committee and participants at the xs/fs Experimental Reading Group at Florida State and the 2013 Public Choice Society Conference. Finally, I thank two anonymous referees and especially Lawrence White for comments that have greatly improved this paper. Naturally, any errors are my own.

Appendix

Note: these are the instructions from FOUR and FOURREV. To maintain neutral language, the monopoly market structure was referred to as ‘multiple’ and quadopoly was referred to as ‘single’.

Instructions

This is an experiment on economic decision making.

Please turn off and stow all electronic devices (cell phones, computers, tablets, etc.).

Also, please do not communicate with other subjects from now until the end of today’s experiment.

If you have a question at any point during these instructions or during the experiment, please raise your hand and an experimenter will come to your terminal to address your question privately.

For participating in today’s experiment you will receive a show-up fee of $10 plus the amount you earn during the course of the experiment. During the experiment, your earnings (excluding show-up fee) will be designated in experimental dollars ($). At the conclusion of the experiment, experimental dollars will be converted into U.S. dollars (US$) at an exchange rate of $10.00 to US$1.00.

> Example: If you earn $80.00 during the experiment, you will be paid US$8.00 at the conclusion of the experiment.

You cannot leave today with less than your US$10.00 show-up fee.

This experiment is composed of 3 **blocks**.

Each block is divided into 10 **periods** and each period is further sub-divided into 2 **stages**.

In other words, this experiment consists of 3 blocks, 10 periods and 60 stages total. At any point during the experiment you can determine the block number, period number, and stage by examining the top of your computer screen.

In each period, you will first participate in an **Innovation stage** and then in a **Market stage**. We will discuss the Market stage first.

Market Stage

In each Market stage, you will have the opportunity to sell units of a good.

You will have a production cost per unit, or an amount it costs you to produce a single unit. This production cost per unit is constant regardless of how many units you produce.

Example: Say your production cost per unit is $7.75. Then your first unit costs you $7.75 to produce, your second also costs you $7.75, and so on. If you produce 5 units and your production cost per unit is $7.75, your total production cost is $38.75 (since 5 times 7.75 equals 38.75).

In each Market stage, you will decide the **maximum number of units you wish to sell** in the current Market stage. You may choose to sell 0, 1, 2, 3, 4, 5, 6, 7, or 8 units.

Choosing, say, 7 units as your maximum quantity to sell does not guarantee that you will sell 7 units. The process that determines how many units you actually sell is described later.

If you select 7 as your maximum number to sell but you only sell 3 units your total production cost will be for 3 and not 7 units. Think of units as being ‘‘made to order.’’ You choose the maximum number of units you are willing to sell, but you only produce the number of units you end up selling.

In addition to submitting a maximum number of units to sell, you will also submit a **price**. This price is the price that you are willing to sell all of your units for. You cannot sell different units for different prices. Prices can be any two decimal number from your cost of producing units to $20.00. The computer will not let you sell units at a loss.

> Example: Say you select a maximum number of units to sell of 2 and a price of $8.00. If you end up selling both units you will sell them each for $8.00.

You are in a market with three other sellers. They are currently reading the same instructions you are and will be confronting the same decisions that you will. They will also be selecting a maximum number of units to sell and a price in each Market stage.

What determines how many units you sell? Once you and the other sellers have entered prices, all of your prices will be publicly displayed. You will see the prices they chose and they will see the price you chose. However, you will not know what maximum quantity they selected and they will not know what maximum quantity you selected.

The buyers in this experiment are computerized. Each buyer has some $ value for 1 unit. We refer to this value as their **buyer value**.

If you post a price that is lower than the prices posted by all the other sellers in your market, buyers will ‘‘line up’’ to purchase from you first. They will line up in descending order of their buyer value. In other words, the buyer with the highest buyer value will be ‘‘at the front of the line’’ and the buyer with the lowest buyer value will be at the back of the line.

Remember: You cannot sell more units than the maximum number of units you select!

Example: Say there are 2 buyers: Buyer 1 with a buyer value of $10.00 and Buyer 2 with a buyer value of $5.00. If you post a price of $8.00 and the other sellers post a price of $9.00, the buyers come to you first since your price is the lowest. Buyer 1 values a unit at $10.00 and you are selling units at $8.00, so Buyer 1 is willing to purchase 1 unit from you. Buyer 2 is next in line, but since he only values units at $5.00, he does not buy a unit from you.

Example: Say there are buyers with buyer values of $10.00, $9.00, and $8.00. Say you post a price of $7.80 and a maximum quantity of 2. Suppose that the other sellers in your market post prices of $7.95, $8.00 and $8.10. Since you have the lowest price, the buyers come to you first. All three buyers would like to buy a unit each from you since your price is lower than each of their buyer values. However, only the buyers with buyer values of $10.00 and $9.00 are able to buy from you since you set your maximum quantity to 2.

If you post a higher price than another seller in your market, you must wait until they have sold their maximum quantity. If they have sold their maximum quantity and there are still buyers who wish to buy at your price, these buyers will buy from you.

If you and another seller in your market post identical prices, the computer determines the number of buyers who wish to buy units at your common price. If there are ''extra'' units that cannot be divided evenly among the sellers who submitted the same price, the ''extra'' unit(s) is (are) randomly awarded by the computer.

Example: Say you and another seller both post a price of $8.25 and that 7 buyers are willing to buy units at that price. You and the other seller each sell 3 units and the 7th or ''extra'' unit is randomly assigned to either you or the other seller. You and the other seller each have a 50/50 chance of selling the extra unit.

Summary of How Units Are Sold

If your price is lower:

You sell until you have sold your maximum quantity, or

You sell until your price is greater than any remaining buyer's buyer value

If your price is higher:

You wait until any other sellers (with lower prices) have sold their maximum quantity, then the process is the same as above

If your price is the same:

The computer calculates the number of buyers who wish to buy at your common price. If there are ''extra'' units that cannot be divided evenly among the sellers who submitted the same price, the computer will randomly determine how to allocate the extra units.

You can never sell more units than the maximum number of units you select!

Innovation Stage

You and the other sellers in your market will begin each block with the same production cost per unit: **$7.75 per unit**. However, in the Innovation stage you and the other sellers will each independently make choices which may allow your own production costs to be reduced. Your choices in the Innovation stage can only influence your own production costs and not those of the other sellers, and vice versa.

Your costs will be influenced by a **random process**. This process involves ten numbers: 1, 2, 3, 4, 5, 6, 7, 8, 9 and 10. Each of these numbers is equally likely to be selected by the process. In other words, the probability of any one number, say 6, being selected by the process is 1/10 or 10 %.

The process just described has two outcomes: **success** and **failure**. Success occurs when the process selects the number 10; failure occurs when the process selects any number other than 10 (numbers 1–9). Since the probability that the process will select the number 10 is 1/10 or 10 %, the **probability that the process results in success is 1/10 or 10 %.**

If the process is a success, your costs of producing units in the current and all future periods within a block will be reduced by **$0.25**.

> Example: If you are successful in the Innovation stage of block 1, period 1, your cost of producing units is lowered from $7.75 to $7.50 for period 1 and all the remaining periods (2–10) in block 1.

Note that your profits may be influenced by having lower production costs.

During each Innovation stage, you will be asked to select the number of innovation processes you wish to undertake in the current period. In the first period you may choose any number of processes between 0 and 30. While you can choose up to 30 processes, **you may only receive one $0.25 cost reduction per period**.

> Example: Say you attempt 11 processes. Suppose that the first 7 result in failure, but that the 8th process is a success (the number 10 is randomly selected). Then you will receive a cost reduction for the period. Because you can only have one cost reduction per period, the 9th, 10th and 11th processes are immaterial and will not be displayed.

It is important to note that even if you choose a large number of processes, success is not guaranteed because each process is randomly and independently determined.

The potential benefit of choosing at least one process is the chance of getting a cost reduction. However, **innovation processes are not costless**. Each process costs **$0.10**. So for every $0.10 you agree to spend, you have a 10 % chance of receiving a cost reduction of $0.25.

Note that you and the other sellers in your market make independent process decisions in each Innovation stage.

> Example: Say you chose 6 processes and one of the other sellers chose 16. You spend $0.60 and the other seller spends $1.60, yet random chance may mean that you are successful and that the other seller is not because the computer picks different random outcomes for you both.

After each Innovation stage, you will learn how much the other sellers in your market spent on innovation processes and they will learn how much you spent. You will also learn whether the other sellers were successful or not, and vice versa.

Profit

Your profit for a particular period is calculated as follows:

$$\begin{aligned}\textbf{Profit} = (&\textbf{Your Price} - \textbf{Your Production Cost Per Unit}) * (\textbf{Number of Units Sold}) \\ &- (\textbf{Processes Bought}) * (\$\mathbf{0.10})\end{aligned}$$

Notice that your profit *increases* in your price and in the number of units you sell. Your profit also increases with *decreases* in your production cost. Finally, your profit decreases (by $0.10) with each innovation process you buy.

The computer will calculate your profit for you. After each Market stage you will receive feedback about your profit. You *will not* receive information about the profit of the other sellers in your market, nor will they receive any feedback about your profit.

At the end of each period your profit (calculated according to the above formula) is added to the profit you have previously earned. The result is referred to as your **total profit**.

If your total profit dips below $0.00, the computer will not let you lose additional money. However, if you have $0.00 in total profit you will not be able to purchase any processes in the Innovation stage.

After each block, your total profit will reset. However, *you will be paid your earnings for all 3 blocks* at the conclusion of the experiment!

Additional Instructions

Between blocks, your cost of producing units ‘‘resets’’ to $7.75 per unit.

> Example: Say you had four successes in block 1, so that your cost of producing units was $6.75. At the conclusion of block 1, period 10, but before the start of block 2, period 1, your cost of producing units will reset to $7.75.

There are two types of blocks in this experiment: **single blocks** and **multiple blocks**.

Buyers have the same buyer values regardless of block type and they will only ever buy *at most one* unit from each seller.

In single blocks, each buyer demands one unit total; in multiple blocks, each buyer demands more than one unit total.

Again, even in multiple blocks where buyers demand more than one unit total, they will only buy at most one unit from each seller.

You will not be told whether you are playing in a single block or a multiple block.

Prior to the start of each block, you will be given an **endowment of $5.00**. This endowment is provided to allow you to buy innovation processes in the first period of the block if you so choose. You are under no obligation to buy processes in the first nor in any subsequent period. Your profits may be influenced by having lower production costs. If you wish, you can go the entire experiment without buying an innovation process. If you do so, your endowments will be part of your earnings paid to you at the experiment’s conclusion.

You will remain in the same market throughout the experiment. Thus, the other sellers in your market in block 1 will be the other sellers in your market in subsequent blocks as well.

Review

In Innovation stages you choose a number of processes. Each process has a 1/10 or 10 % chance of success. Processes cost $0.10. Success means a $0.25 per unit cost reduction for the current and all future periods within a block.

In Market stages you choose a price and a maximum quantity. Buyers will buy from you first if your price is less than the other seller’s prices. Whether your price is initially the lowest or not, it must be less than their buyer value for them to purchase.

In single blocks, buyers only demand one unit total; in multiple blocks they demand more than one unit total. Regardless of block type, buyers buy at most one unit from each seller. You will not be told whether the block type is single or multiple.

References

Aghion, P., Bechtold, S., Cassar, L., & Herz, H. (2014). The causal effects of competition on innovation: Experimental evidence. NBER working paper no. 19987.

Comanor, W., & Leibenstein, H. (1969). Allocative efficiency, x-efficiency and the measurement of welfare losses. Economica, 36(143), 304–309.

Darai, D., Sacco, D., & Schmutzler, A. (2010). Competition and innovation: An experimental investigation. Experimental Economics, 13(4), 439–460.

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. Experimental Economics, 10(2), 171–178.

Frantz, R. (1988). X-Efficiency: Theory, evidence and applications. Norwell, MA: Kluwer Academic Publishers.

Frantz, R. (2007). Empirical evidence on x-efficiency, 1967–2004. In R. Frantz (Ed.), Renaissance in behavioral economics: Essays in honor of Harvey Leibenstein (pp. 211–227). London: Routledge.

Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with ORSEE. Journal of the Economic Science Association, 1(1), 114–125.

Isaac, M., & Reynolds, S. (1992). Schumpeterian competition in experimental markets. Journal of Economic Behavior and Organization, 17(1), 59–100.

Leibenstein, H. (1966). Allocative efficiency vs. x-efficiency. The American Economic Review, 56(3), 392–415.

Leibenstein, H. (1978). On the basic proposition of x-efficiency theory. The American Economic Review, 68(2), 328–332.

Martin, S. (2010). Industrial organization in context. Oxford: Oxford University Press.

Parish, R., & Ng, Y.-K. (1972). Monopoly, x-efficiency and the measurement of welfare loss. Economica, 39(155), 301–308.

Plott, C. (1989). An updated review of industrial organization: Applications of experimental methods. In R. Schmalensee & R. Willig (Eds.), Handbook of industrial organization (Vol. 2, pp. 1109–1176). Amsterdam: North Holland.

Sacco, D., & Schmutzler, A. (2011). Is there a u-shaped relation between competition and investment? International Journal of Industrial Organization, 29(1), 65–73.

Stigler, G. (1976). The xistence of x-efficiency. The American Economic Review, 66(1), 213–216.

Waldman, D., & Jensen, E. (2007). Industrial organization: Theory and practice (3rd ed.). Upper Saddle River, NJ: Prentice Hall.

Notes

1 X-inefficiency is defined as the difference between actual and minimum cost of production for a given output (Leibenstein 1978).

2 To the best of my knowledge. For surveys of the existing theoretical and (non-experimental) empirical literature, see Frantz (1988, 2007).

3 The exchange rate between experimental and U.S. currency was $10.00 experimental dollars to US$1.00.

4 If, over the course of several periods, a subject used his endowment on innovation and did not offset this expenditure with earnings from the Market stage, his cumulative profit fell. If a subject's cumulative profit dipped below $3.00 during a block, the experimental software only let him purchase as many attempts as he had funds for. His cumulative profit could reach $0.00, but could not go negative. Across all treatments there were only six cases where a subject's cumulative profit reached $0.00. In each, the

subject in question subsequently earned profits in a Market stage and ended the block with a cumulative profit above $0.00.

5 Subjects set their capacity to 8 units 77.5 % of the time, so while the Market stage was a price-capacity game, competition *de facto* focused on price. Subjects very rarely made capacity choices that, given their marginal cost and chosen price, restricted their sales (and thus their profit). Out of 2220 capacity choices in the data, 104 (4.7 %) were ‘‘mistakes’’ in this sense. 59.6 % of these mistakes occurred in either periods 1 or 2 of Block 1. In Blocks 2 and 3, subjects made capacity mistakes just 2.2 % of the time.

6 As in naturally-occurring markets, sellers had incomplete information about demand. Though this added complexity to their decision task, optimal price was a discovery process for subjects in *all* treatments.

7 An alternative motivation is the presence of a barrier to competitive entry (*Monopoly*), its removal (*Oligopoly*), and eventual reinstatement (*Monopoly*).

8 There were two reasons for this differential feedback. First, subjects in the $n > 1$ treatments were given more feedback in monopoly so that their feedback would be constant in both monopoly and oligopoly. Second, it was possible that with more feedback, imitation would induce subjects in the $n > 1$ treatments to attempt more innovation in Block 1 than did ONE subjects. Unreported regressions (available from the author) indicate that this did not occur: There was no significant difference in innovation expenditure across treatments.

9 Though this feature of the design is perhaps unrealistic, subjects were in the dark about competition only during the Innovation stage of period 1. Thereafter they knew exactly how many competitors they faced.

10 This is the standard measure with linear demand and constant marginal cost. For example, see Figure 1 in Comanor and Leibenstein (1969), Figure 13 in Frantz (1988), Figure 3.10 in Waldman and Jensen (2007), and Figure 2.9 in Martin (2010). See Parish and Ng (1972) for an alternative measure.

11 Formally, for market structure $j \in \{monopoly;\ duopoly;\ quadopoly\}$ and a given block, average profit for cost level $\bar{s}$ is:

$$\pi_j(\bar{s}) = \frac{\sum_{t=1}^{10} \sum_{i=1}^{N_j} \mathbb{1}\{s_{it} = \bar{s}\}(p_{it} - s_t)q_{it}}{\sum_{i=1}^{N_j} \mathbb{1}\{s_{it} = \bar{s}\}}$$

where i indexes subjects in market structure j, N_j is the number of subjects in market structure j in the block, prices are denoted by p, sales by q, and $\mathbb{1}(\cdot)$ is the indicator function.

12 Unreported regressions (available from the author) indicate no significant differences across treatments in Block 1. In Block 3, price was initially lower in TWOREV and FOURREV than in ONE, because subjects in the former treatments were not sure if they were monopolists or oligopolists. They quickly figured out that they were monopolists and raised their prices.

13 6 of the 77 average profit figures that were used to calculate these optimal paths were imputed from linear or polynomial regressions because a few cost levels were never reached in certain market structures.

14 These results are subject to the caveat that independence across periods in the deviation paths is not strictly satisfied.

15 Isaac and Reynolds (1992) also report under-investment in cost innovation in their similar environment.

16 When x-inefficiency is calculated relative to the exogenous cost frontier (optimality is a cost reduction in each period), or relative to the subject with the lowest cost in each treatment (the ‘‘best practice firm’’), x-inefficiency is significantly lower in oligopoly than monopoly in Block 3.

17 For a notable critique of x-inefficiency theory, see Stigler (1976).