

Quantity discounts on a virtual good: The results of a massive pricing experiment at King Digital Entertainment

Steven D. Levitt^{a,b,1}, John A. List^{a,b}, Susanne Neckermann^c, and David Nelson^d

^aDepartment of Economics, University of Chicago, Chicago, IL 60637; ^bNational Bureau of Economic Research, Cambridge, MA 02138; ^cDepartment of Economics, Erasmus University Rotterdam, 3062 PA Rotterdam, The Netherlands; and ^dKing Digital Entertainment, 11134 Stockholm, Sweden

Edited by Richard M. Shiffrin, Indiana University, Bloomington, IN, and approved May 26, 2016 (received for review July 10, 2015)

We report on a natural field experiment on quantity discounts involving more than 14 million consumers. Implementing price reductions ranging from 9–70% for large purchases, we found remarkably little impact on revenue, either positively or negatively. There was virtually no increase in the quantity of customers making a purchase; all the observed changes occurred for customers who already were buyers. We found evidence that infrequent purchasers are more responsive to discounts than frequent purchasers. There was some evidence of habit formation when prices returned to pre-experiment levels. There also was some evidence that consumers contemplating small purchases are discouraged by the presence of extreme quantity discounts for large purchases.

price discrimination | natural experiment | quantity discount | virtual | King

The “Big Data” revolution offers enormous opportunities not only for firms (1, 2) but also for scientific advances (3). (Big data are often defined as data that are generated and available in real time and that are more granular, less structured, and at a larger scale than static datasets of the past.) As noted by Levitt and List (4), the data that are increasingly generated by firms represent a rich and largely untapped source for exploring theories that otherwise are difficult to test empirically. The value of such data is magnified when coupled with randomization in the form of natural field experiments (5) allowing clear, causal inference. Although still rare, academic–firm research joint ventures therefore are ideal to allow both researchers and firms to profit from the “data explosion.”

This paper reports one such big data experiment carried out as a partnership between King Digital Entertainment (hereafter King), one of the world’s leading gaming companies, and academics at the University of Chicago and Erasmus University Rotterdam. Together, we designed and implemented a randomized pricing experiment involving more than 14 million customers and aimed at understanding the effects of quantity discounts on revenue and game play.

King offers its games for free but makes profit from consumers buying in-game content, such as gold bars that allow users to move up the levels of the game faster. Most customers, however, never buy such content, and a large percent of the people buying buy only small quantities. Hence, the majority of the revenue for the company stems from a small percentage of users.

Historically, King has used very simple pricing strategies. Prices are the same for all customers, and quantity discounts have been minimal. The per-unit price for nine gold bars (the smallest bundle of in-game currency offered) is only 9% more than the per-unit price for a purchase of 1,000 gold bars (the largest bundle available). Given the enormous heterogeneity in demand across customers, one might expect such a simple pricing scheme to be far from optimal. There is a large theory literature in economics exploring volume discounts or what economists call “second-degree price discrimination” (6–10). These authors highlight the potential gains associated with a variety of nonlinear pricing strategies such as two-part tariffs (in which consumers pay both a fixed fee that is

independent of the quantity consumed and a price per unit) and quantity discounts.

In contrast to the theory of this topic, which by now is well understood, the empirical exploration of these issues is far less well developed. Borenstein (11) and Shephard (12) provide two early analyses of price discrimination on the quality dimension in the retail gasoline market. More recent investigations include McManus (13), Leslie (14), Busse and Rysman (15), Cohen (16), and Olken and Baron (17). [There is also an extensive literature on quantity discounts in the operations research literature. See, for instance, the survey done by Munson and Rosenblatt (18), which concludes that firms generally have not implemented the pricing models suggested by the academic literature.] In all these papers, the authors analyze differences in prices that arise in a particular setting and try to reconcile the observed differences with economic theory. Each of these studies faces a fundamental limitation: They were all based on observational data.

Our paper differs fundamentally from the existing empirical literature because we actually had the power to change prices and did so in a randomized field experiment in which quantity discounts were varied over an extremely broad range. We observed customer behavior before the randomization, during a 3-mo period in which customers faced very different price schedules, and for 2 mo after the experiment ended when all consumers once again experienced minimal quantity discounts. In contrast to almost all previous studies, we were able to observe not only market outcomes but also the individual actions of more than 14 million customers. Another benefit of our setting is that the marginal cost to King of providing the in-game currency is zero, and for regulatory reasons all customers are offered the same menu of prices, except when the firm carries out explicitly defined experiments such as the one reported in this paper. Thus, we were able to focus on pure price discrimination, without considering price differences arising from delivery costs or differential bargaining power on the part of consumers.

In designing the experimental price menus, we faced three main constraints. First, a two-part tariff in which customers paid a fixed fee to have the right to make in-game purchases was ruled out because King’s software is not designed to handle such a scheme. Second, because consumers have the opportunity to make repeat purchases and because of perceived fairness concerns, the per-unit cost had to be flat or decreasing in the

This paper results from the Arthur M. Sackler Colloquium of the National Academy of Sciences, “Drawing Causal Inference from Big Data,” held March 26–27, 2015, at the National Academies of Sciences in Washington, DC. The complete program and video recordings of most presentations are available on the NAS website at www.nasonline.org/Big-data.

Author contributions: S.D.L., J.A.L., S.N., and D.N. designed research, performed research, analyzed data, and wrote the paper.

Conflict of interest statement: D.N. is an employee of King Digital Entertainment, which potentially stands to benefit from the publication of this paper.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. Email: slevitt@midway.uchicago.edu.

quantity purchased. Otherwise, consumers could simply make multiple smaller purchases at the lowest per-unit cost. Finally, we were not able to change the prices on bundles of gold bars ranging in number from 9 to 59. Thus, our experimental design problem boiled down to constructing price menus with differing degrees of convexity. The existing theory provided little guidance regarding the optimal amount of convexity, so we experimented across an extremely wide range, with discounts for large purchases varying from 9% (King's historical offering) to up to more than 70% in the most extreme intervention. The four experimental price menus are shown in Fig. 1.

The experiment addresses two further questions of interest. The first is the issue of heterogeneity in consumer response. There is every reason to expect consumers to respond differently to quantity discounts. The great majority of King consumers have never made a purchase of gold bars, another group (whom we call "medium-value" customers) makes occasional and typically small purchases, and a small number of "high-value" players (who account for a large share of total revenue) make frequent/large purchases. The second is that immediate expenditure is not the only, or even the primary, objective of King. They are at least equally concerned with the quantity of game play, which, at least correlationally, is associated both with future play and spending both in this particular game and on other King game offerings. (Wall Street analysts also exhibit interest in game-play metrics not immediately tied to revenues. To address such interests, King provides daily and monthly active and unique user numbers in its quarterly earnings reports.) Thus, a critical issue in this experiment was how different levels of quantity discounts would affect long-term levels of both expenditure and game play. Theoretically, habit-formation models (19, 20) imply that big discounts in the present will be associated with greater play both in the present and later. More standard economic models with decreasing marginal utility might suggest that temporary quantity discounts might lead to more play in the present and less play in the future because consumers opt to play when play is "cheap."

A number of insights emerge from the experiment. First, over a wide range of quantity discounts, revenue and profit were essentially unchanged. In economic terms, this result implies that the price elasticity of demand (the percent of change in the

quantity demanded induced by a 1% increase in price) is close to one. Second, heterogeneity in response across consumers is evident. Medium-value customers spend more when presented with radical quantity discounts, but high-value customers spend less. These two effects are offsetting. Third, we find quantity discounts have virtually no impact on the share of consumers making a purchase; i.e., in economics terms, there is no impact on the "extensive margin," only changes in behavior among those who are already purchasing. Fourth, we find some limited evidence of habit formation: Medium-value consumers who temporarily faced low prices for large quantities continued to consume more after these large discounts were removed. However, the ultimate conclusion of our experiment is that the potential gains to King from the forms of price discrimination we explored are remarkably small, in contrast to what one might have expected based on the prior theoretical and empirical literature. From a profit perspective, King is "unlucky" that the relative responsiveness of medium-value and high-value customers is such that quantity discounts turn out not to be revenue enhancing.

The remainder of this paper is structured as follows. The next section provides background information about King and describes the experimental design and its implementation in greater detail. We then report the results of the experiment during the 13 wk in which consumers saw different prices across treatments. The fourth section investigates whether there were lasting effects on behavior in the 6 mo following the experiment, when all consumers faced the same prices. The fifth section presents our conclusions.

Background, Experimental Design, and Implementation

King is one of the world's most successful makers of online games, most notably the blockbuster game Candy Crush, which has been downloaded onto more than 500 million devices. King is widely recognized as being among an elite set of firms that have parlayed big data acumen and a culture of experimentation into business success (21, 22). (As of this writing, King has a market capitalization of nearly \$5 billion, annual revenues of more than \$2 billion, and net profits of \$575 million.)

King games are provided to users free of charge. The great majority of King's revenues are generated through in-game purchases.

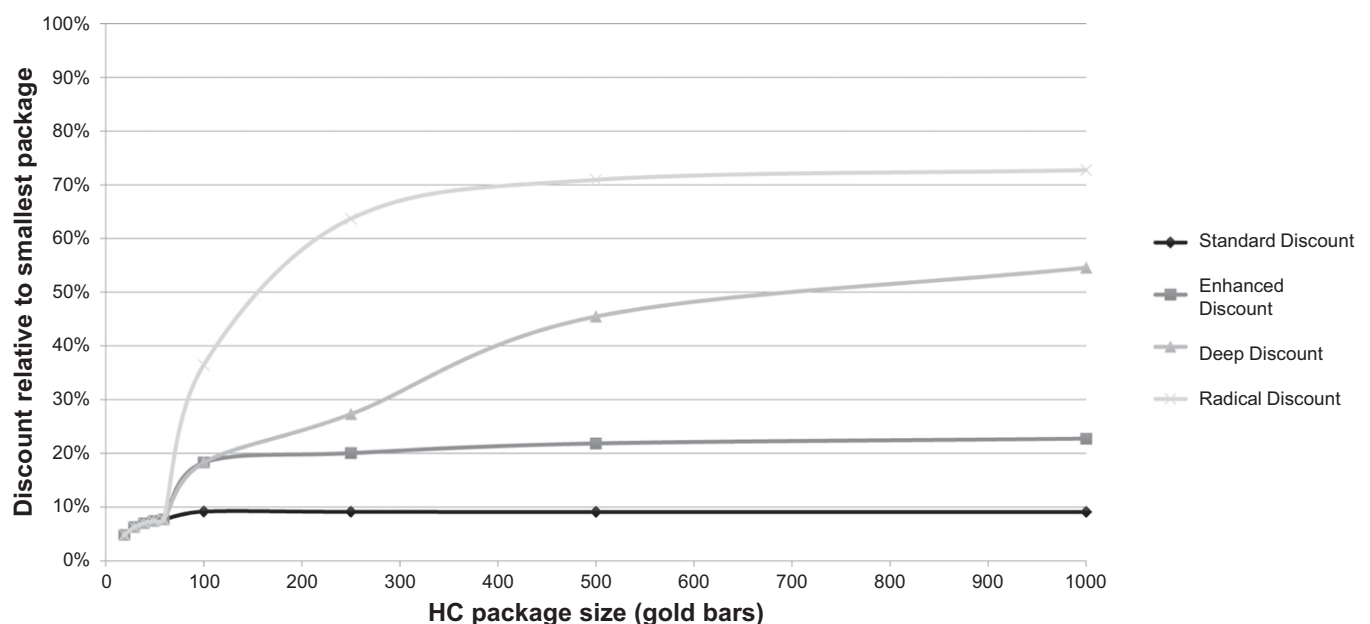


Fig. 1. Quantity discounts presented to customers across treatment arms.

In the game we study, players are allotted a fixed set of moves to complete a task. If the player is successful, he/she moves on to the next level; if unsuccessful, the player repeats the current level until success is achieved. During the course of play, consumers have the opportunity to purchase a virtual currency (gold bars) which can be redeemed for game features that facilitate completion of the level. In most cases, nine gold bars are needed to purchase extra moves or boosters, and this is the smallest quantity of gold bars sold. The consumer pays 99 cents for nine gold bars, or 11 cents per bar. Historically King has offered extremely modest discounts: The price per bar is reduced by less than 10% for a purchase of 1,000 gold bars, the largest quantity offered. This practice stands in stark contrast to many of King's direct competitors and to many other producers of consumer goods. (For instance, in the popular game Clash of Clans, Supercell offers quantity discounts up to 28% percent for in-game gems purchases. Likewise, at the McDonalds restaurant in Hyde Park, IL, a 10-piece Chicken McNugget order sells for \$5.12, and 20 pieces sell for \$5.70.) Note that although, in general, price discrimination is carried out by firms to maximize profit, it also may benefit consumers. In this particular experiment, for instance, all our treatment interventions involved reducing prices relative to King's status quo.

Our experimental intervention took the form of four treatment arms that offered different degrees of quantity discounts. Fig. 1 shows the four different price schedules. Prices were held fixed for small purchases (9–59 gold bars) across the four treatment arms. Prices varied only for purchases of 100 or more gold bars, allowing us to isolate the impact of quantity discounts from the impact of lower prices per se. Historically, about 10% of purchases involved at least 100 gold bars, and these purchases account for 45% of revenue. The bottom line in Fig. 1, which we denote the “standard discount,” mirrors historical pricing by King and offers quantity discounts of less than 10%, even for purchases of 1,000 gold bars. The second treatment “enhanced discount,” mirrors the shape of the historical pricing pattern but roughly doubles the quantity discount offered. In a third treatment, which we denote “deep discount,” the per-unit discount rises monotonically to almost 60% for the largest purchases. In the most extreme intervention, which we call the “radical discount” treatment, even intermediate-sized purchases were offered per-unit price discounts of more than 60%; the largest purchases were rewarded with discounts of more than 70%.

In total, more than 14 million consumers were included in the experiment. A given consumer saw the same price schedule for the 3-mo duration of the experiment, after which all prices reverted to the historical price discounts offered by King. Consumers were not informed that prices were being experimentally varied. (We can find no evidence that our experimental prices were discussed on forums or chat groups related to the game.) Other than the differences in prices, the purchase screen was identical across treatments and was similar to how it always had appeared. No mention was made during the experiment as to whether the prices presented were temporary or permanent. Subjects were randomized into one of the four possible price schedules, with a 20% probability of assignment to the standard discount or the radical discount and a 30% probability of being assigned to the enhanced discount or the deep discount. The researchers received only aggregated data. The data used in this analysis are available to other scholars for replication purposes. Because no individual data were provided to us, and because the experiment was carried on as part of normal business operations at King, our project was deemed exempt from human subjects regulations by the University of Chicago Institutional Review Board.

Table 1 examines the extent to which the four arms of the study are balanced on observable pretreatment characteristics.

Each row of the table represents a different pretreatment characteristic. The four columns of the table correspond to the four treatment arms. In each cell of the table, we report the mean within the treatment arm along with SEs where relevant. At the request of King, we normalized many of the outcome variables to be equal to 100 for the standard discount group in the preperiod. Despite the massive samples, revenues per user differ by up to 3% across treatments because there is a long right tail in revenues, with a small number of high-value customers. The other measures are well balanced across treatments: levels attempted, levels completed, and the share of customers falling into each of three customer segments, i.e., players who have never made purchases, medium-value customers, and high-value customers.

Short-Term Results of the Experiment

Table 2 summarizes the key results of the experiment for the 12 wk in which it was running. The four columns of the table again correspond to the four treatment arms. Each row represents a different outcome measure. Entries in the table are means across all subjects in a treatment. The top row reports revenues per customer. Note that all these values are higher than in the top row of Table 1 because the pre-experimental window was 1 mo, but the experimental period was 12 wk. The results are remarkably similar across the standard discount, enhanced discount, and deep discount treatments, with revenues varying less than 1% across those treatments and with no statistical differences between the treatments. This finding implies that, over an extremely wide range of quantity discounts, the short-run price elasticity of demand is very close to one, i.e., that a 1% decrease in price leads to a 1% increase in quantity. Only for the radical discount is this pattern broken, with revenues falling ~10%.

The small differences in revenue across treatments disguise significant impacts on quantities purchased. Relative to the standard discount treatment, the progressively deeper discounts drive quantity increases of 6.7%, 11.2%, and 44.9%, respectively. (Consumers who respond to the discounts by buying larger quantities use the items they purchase quickly rather than saving them, even though the items are storable.) These differences are all highly statistically significant. The average price paid falls by similar proportions, however, leaving revenue essentially unchanged.

Despite the increases in average quantity, the fourth row of the table shows that there is no difference across treatments in the share of consumers making at least one purchase during the treatment period: In all four treatment arms this share is 0.026. The number of consumers making a purchase is affected by two potential forces in this experiment: (i) bigger quantity discounts may induce more purchases along the extensive margin, and (ii) large quantity discounts may discourage consumers who are interested only in making a small purchase but who are made to feel that small purchases are a “bad deal” in the face of steep quantity discounts. We believe that both these forces are at work, offsetting one another, but that both are relatively weak. For instance, among those who have never made a purchase before, the radical discount treatment induces a 16% increase in purchases (although this increase is from an extremely low base rate). Among customers who have made a positive number of small purchases in the past, the number making a purchase in radical discount scheme actually falls by 2% relative to the standard discount; this result is consistent with the second force highlighted above. (Indeed, roughly 4 wk into the experiment, the share of players making at least one purchase fell monotonically with the size of the discounts offered, exactly the opposite of the predictions of standard economic theory, which ignores the second force highlighted above. These effects were undone in the latter part of the experiment.) The data suggest that large numbers of low-priced boosters have little allure for consumers who historically have made little use of boosters. Whether such

Table 1. Testing for balance across predetermined customer characteristics, by treatment arm

Variable	Standard discount: maximum discount 9% (SE)	Enhanced discount: maximum discount 23% (SE)	Deep discount: maximum discount 55% (SE)	Radical discount: maximum discount 73% (SE)
Revenue per user, prior 30 d	100.00 (0.993)	100.74 (0.843)	100.62 (0.848)	97.77 (0.937)
Game rounds per user, prior 30 d	100.00 (0.113)	99.96 (0.092)	99.98 (0.092)	99.79 (0.113)
Levels completed per user, prior 30 d	100.00 (0.105)	99.89 (0.085)	99.95 (0.085)	99.74 (0.104)
Share of players who have never paid	0.982	0.982	0.982	0.982
Share of medium-value players	0.048	0.048	0.048	0.048
Share of high-value players	0.009	0.009	0.009	0.009
Number of players assigned	4,024,826	6,036,239	6,035,011	4,021,156

This table presents a comparison across treatment arms of observable customer characteristics before the treatment intervention. The variables revenue per user, game rounds per user, and levels completed per user are indexed to be equal to 100 on the pretreatment standard discount group.

consumers could be transformed through price discounts on smaller bundles remains an open question and one that is of central interest to King.

The quantity discounts had only a minor impact on game play, as demonstrated in the final two rows of Table 2. The total number of rounds played (i.e., roughly 194 rounds) was virtually unaffected across treatments. The number of levels successfully completed per user rose slightly with the magnitude of the discounts, but the increase was only a fraction of a percent. Given the results on purchases, this outcome is not surprising. Although the quantity purchased rose substantially in percentage terms, purchases were quite rare overall, so these changes did not have a major impact on game play.

The aggregate results reported above potentially hide substantial heterogeneity across customers. Unrelated to and before our experiment, players had been segmented into categories based on past purchasing behavior. Table 3 reports the results when all customers who had previously made a purchase in the game are divided into two mutually exclusive and exhaustive categories based on these predefined segments, medium-value and high-value players. (There also was a large group of players who had never made a purchase before the experiment, but these players accounted for a trivial share of revenue in the experimental period.) Once again, treatments are shown in four columns. Outcome variables are presented first for the medium-value segment and then for the high-spending segment.

It is interesting that the radical discount, which performed so poorly overall, actually generated the most revenue from the medium-value group (2–3% more than the other treatments, statistically significant at the 0.05 level) but fared very badly for the high-value users, with a highly statistically significant decline in revenue of almost 15%. For medium-value customers, the price discounts induced sufficient substitution from small to large purchases to be revenue enhancing. Among high-value customers,

the quantity response was not sufficient to offset the price declines. Put another way, the high-value customers are less price elastic than the medium-value customers. Absent other constraints, it would be profit maximizing to charge the high-value players higher prices. For legal reasons and because of fairness concerns, King would never do so. (One potential pricing scheme King could consider would be a price per unit of 4 for the first 20 units, a price per unit of 2 for the next 10 units, and a price per unit of 3 thereafter. The average price never increases with quantity in this scheme, so consumers do not have an incentive to divide their purchases into smaller chunks, but it does provide a lower marginal price for purchases of intermediate size. We thank a referee for suggesting this mechanism.) Had the high-value customers been more price elastic than other players, then steeper quantity discounts likely would have been strongly profit enhancing. The enhanced discount yields the best results for high-value customers but generates only a statistically insignificant 1% increase in revenues for that group.

The other important difference that emerges between medium-value customers and high-value customers is the pattern of levels successfully completed per customer. For medium-value customers, we see little impact; among high-value customers that number jumps by 12% between the standard discount and the radical discount treatments. Note also that in the radical discount treatment the increase in successful completions is accomplished using slightly fewer game rounds: The players use their gold bars more frequently to complete the task.

Impacts of the Pricing Experiment on Behavior After It Ends

It is impossible to determine from the results presented above which of the pricing schedules yields the greatest value to King. The short-term impact on revenue is so small that any persistent effects of the experiment may swamp these differences.

Table 4 presents results, broken down into medium-value consumers and high-value consumers, for the month following

Table 2. The impact of quantity discounts on customer behavior during the experimental period

Outcome measure	Standard discount: maximum discount 9% (SE)	Enhanced discount: maximum discount 23% (SE)	Deep discount: maximum discount 55% (SE)	Radical discount: maximum discount 73% (SE)
Revenues per user	222.84 (2.108)	224.50 (1.720)	222.44 (1.669)	198.72 (1.650)
Quantity purchased per user	206.34 (1.974)	220.11 (1.742)	229.39 (1.845)	298.95 (2.935)
Average price paid per gold bar	108.00 (1.453)	101.99 (1.123)	96.97 (1.067)	66.47 (0.855)
Share of players making at least one purchase	0.026	0.026	0.026	0.026
Game rounds played per user	214.08 (0.259)	214.39 (0.212)	214.11 (0.212)	213.89 (0.259)
Levels completed per user	182.56 (0.226)	182.86 (0.184)	182.79 (0.183)	183.17 (0.223)
Number of players observed in period	2,933,225	4,397,622	4,399,290	2,930,341

This table presents treatment effects during the experimental period. The variables revenue per user, quantity purchased per user, average price paid per gold bar, game rounds per user, and levels completed per user are indexed to 100 on the pretreatment standard discount group. Per user calculations are based on the sample of customers who played the game at least once in the experimental period.

Table 3. The impact of quantity discounts across medium-value and high-value players

Outcome measure	Standard discount: maximum discount 9% (SE)	Enhanced discount: maximum discount 23% (SE)	Deep discount: maximum discount 55% (SE)	Radical discount: maximum discount 73% (SE)
Medium-value players				
Revenues per user	793.04 (5.789)	789.25 (4.689)	793.61 (4.729)	815.42 (5.870)
Quantity purchased per user	724.76 (5.353)	746.51 (4.581)	768.86 (4.893)	1,081.66 (9.986)
Average price paid per gold bar	109.42 (1.136)	105.72 (0.903)	103.22 (0.900)	75.39 (0.883)
Share of players making at least one purchase	0.291	0.292	0.291	0.292
Game rounds played per user	365.30 (1.411)	363.82 (1.149)	364.47 (1.154)	362.16 (1.399)
Levels completed per user	267.75 (1.112)	266.11 (0.885)	266.76 (0.909)	267.98 (1.092)
Number of players	171,775	257,533	256,798	172,284
High-value players				
Revenues per user	15,189.27 (159.245)	15,311.06 (129.297)	15,190.16 (125.157)	12,995.71 (119.943)
Quantity purchased per user	14,112.89 (149.530)	15,149.61 (131.927)	15,914.06 (140.692)	20,284.65 (222.299)
Average price paid per gold bar	107.627 (1.604)	101.066 (1.226)	95.451 (1.154)	64.067 (0.918)
Share of players making at least one purchase	0.734	0.738	0.737	0.734
Game rounds played per user	515.55 (3.793)	513.89 (3.028)	517.18 (3.088)	508.27 (3.649)
Levels completed per user	437.89 (3.213)	449.25 (2.783)	457.31 (2.828)	491.24 (3.696)
Number of players	34,058	51,194	50,992	33,988

This table presents results during the experimental period across medium-value and high-value players. The variables revenue per user, quantity purchased per user, average price paid per gold bar, game rounds per user, and levels completed per user are indexed to 100 on the pretreatment, standard discount group. Per user calculations are based on the sample of customers who played the game at least once in the experimental period.

the end of the experiment when all consumers saw the same price schedule, i.e., the standard discount. For medium-value players, those who had been exposed to the radical discount continued to spend a statistically insignificant additional 3–4%, as is consistent with habit formation. For high-value customers, there are no clear patterns of behavioral spillovers into the postexperimental period.

Conclusions

In this paper we report on a massive field experiment investigating the impact of quantity discounts in a virtual environment. A number of findings are surprising, at least to the authors. First, varying quantity discounts across an extremely wide range had almost no profit impact in the short term. Second, almost all of the impact of the price changes was among those already making a purchase; radical price reductions induced almost no new customers to buy. Third, there was heterogeneity in response, especially to the radical discount treatment, which led to increased revenue from medium-value customers but a sharp reduction in revenue from high-value customers. Finally, we observe

few differences in behavior in the postexperimental period, although there is some evidence of habit formation among medium-value consumers who bought only small amounts of gold bars and did so sporadically before the price cuts.

It is difficult to know to what extent the results of this experiment will generalize to other settings. The product we examined has many unusual aspects: The goods being purchased are virtual; consumers can play the game without these goods, which only enhance the experience; and consumers had substantial experience with the game before this pricing experiment and potentially had already formed habits in advance of the experiment. Given how little is known empirically about consumer response to quantity discounts, every additional data point represents a material increase in knowledge.

From a corporate perspective, this experiment was somewhat a failure. It did not reveal a pricing strategy that led to a large immediate increase in profitability. Had this experiment yielded different results, however, it easily could have generated hundreds of millions of dollars in profit for King. Even though the experiment did not directly generate profits, it strongly suggests

Table 4. Postexperimental impacts on medium-value and high-value players

Outcome measure	Standard discount: maximum discount 9% (SE)	Enhanced discount: maximum discount 23% (SE)	Deep discount: maximum discount 55% (SE)	Radical discount: maximum discount 73% (SE)
Medium-value players				
Revenue per user	430.93 (7.915)	419.93 (5.547)	437.93 (6.685)	446.10 (8.794)
Game rounds per user	152.95 (0.710)	152.25 (0.575)	152.13 (0.576)	150.56 (0.698)
Level completed per user	99.21 (0.557)	98.16 (0.421)	98.24 (0.435)	97.87 (0.518)
Number of players	119,887	179,696	179,787	120,119
High-value players				
Revenue per user	5,306.79 (71.571)	5,357.60 (59.485)	5,368.36 (59.817)	5,241.66 (70.106)
Game rounds per user	186.21 (1.617)	183.61 (1.284)	183.04 (1.283)	177.79 (1.527)
Levels completed per user	140.76 (1.337)	140.20 (1.103)	140.19 (1.058)	138.79 (1.343)
Number of players	26,546	39,992	39,652	26,588

This table presents sample averages by treatment arm in the time period after the experiment had ended, when all consumers saw the same prices. Results are shown by player segment, based on the pre-experimental classifications. The variables revenue per user, game rounds per user, and levels completed per user are indexed to 100 on the pretreatment standard discount group. Per user calculations are based on the sample of customers who played the game at least once in the postexperimental period.

future experiments that could answer important theoretical questions and/or generate increased profit. For instance, the limited evidence of movement along the extensive margin makes it more likely than otherwise that price increases for small quantity purchases would be profit enhancing, a possibility that the company has not previously tested.

From an academic perspective, however, the experiment was at least partially successful in that it raises many challenges to the conventional thinking. Based on both the theory literature and previous empirical research, one might have suspected that increased profits would result from second-degree price discrimination. In practice, medium-value and high-value customers reacted in offsetting ways, diluting the gains from the quantity discounts. Understanding why King's customers exhibit this pattern would be valuable. Another potentially interesting and surprising feature observed in the data is that this experiment suggests that some consumers who would have made small purchases were discouraged from doing so when faced with large-quantity discounts. That phenomenon would not occur under standard economic models. Not having anticipated that particular result, we did not design the experiment to isolate it cleanly, but it suggests a fruitful avenue for future research. Our results imply that prices change people's perception of the value of goods—perhaps more so in virtual settings—but little is known about this topic.

There can be little doubt that partnering with firms offers opportunities for scientific advances that otherwise would be out of reach (23, 24). Without collaboration with a firm such as King, it is hard to imagine a project like this one ever being carried out as academic research. Moreover, there was essentially no cost to the experiment. King had to offer some menu of prices to these 14 million customers. Because King has invested in an infrastructure for and a mindset of experimentation, only a few person days of effort were needed to implement and analyze the basic findings. The cost per subject of running this experiment was perhaps 1/100th of one cent. Although there are many notable examples of firms sharing data for academic analysis, to date there are surprisingly few cases in which academics and firms have collaborated to run randomized experiments and have made the findings available to the scientific community. In the quest to test economic theory in real-world settings, firm-based field experiments represent a unique opportunity previously out of reach for academic economists. This path appears to be a promising one for future scientific advances.

ACKNOWLEDGMENTS. We thank Chris Goldammer, Michelle Kim, Annie Mail, Daniel Neiberg, Aditya Tata, Richard Thompson, Mattie Toma, Alexandra Vo, the editor, and two anonymous referees for their input and ideas on this project.

- Chen H, Chiang RH, Storey V (2012) Business intelligence and analytics: From big data to big impact. *Manage Inf Syst Q* 36(4):1165–1188.
- McAfee A, Brynjolfsson E (2012) Big data: The management revolution. *Harv Bus Rev* 90(10):60–66, 68, 128.
- Einav L, Levin J (2013) The data revolution and economic analysis. NBER Working Paper 19035 (National Bureau of Economic Research, Cambridge, MA).
- Levitt S, List J (2009) Field experiments in economics: The past, the present, and the future. *Eur Econ Rev* 53(1):1–18.
- Harrison G, List J (2004) Field experiments. *J Econ Lit* 42(4):1009–1055.
- Mussa M, Rosen S (1978) Monopoly and product quality. *J Econ Theory* 18(2):301–317.
- Katz ML (1984) Price discrimination and monopolistic competition. *Econometrica* 52(6):1453–1471.
- Maskin E, Riley J (1984) Monopoly with incomplete information. *Rand J Econ* 15(2):171–196.
- Varian HR (1989) Price discrimination. *Handbook of Industrial Organization*, eds Schmalensee R, Willig R (Elsevier Science, Cambridge, MA), pp 597–654.
- Rochet JC, Stole L (2002) Nonlinear pricing with random participation. *Rev Econ Stud* 69(1):277–311.
- Borenstein S (1991) Selling costs and switching costs: Explaining retail gasoline margins. *Rand J Econ* 22(3):354–369.
- Shepard A (1991) Price discrimination and retail configuration. *J Polit Econ* 99(1):30–53.
- McManus B (2007) Nonlinear pricing in an oligopoly market: The case of specialty coffee. *Rand J Econ* 38(2):512–532.
- Leslie P (2004) Price discrimination in Broadway theater. *Rand J Econ* 35(3):520–541.
- Busse M, Rysman M (2005) Competition and price discrimination in Yellow Pages advertising. *Rand J Econ* 36(2):378–390.
- Cohen A (2008) Package size and price discrimination in the paper towel market. *Int J Ind Organ* 26(2):502–516.
- Olken B, Barron P (2009) The simple economics of extortion: Evidence from trucking in Aceh. *J Polit Econ* 117(3):417–452.
- Munson C, Rosenblatt M (1998) Theories and realities of quantity discounts: An exploratory study. *Prod Oper Manag* 7(4):352–369.
- Skinner B (1972) *Beyond Freedom and Dignity* (Batman Books, New York).
- Becker G, Murphy K (1988) A theory of rational addiction. *J Polit Econ* 96(4):675–700.
- Toumpouris L (2015) King Digital Entertainment Creates and Captures Value Using Data. *Harvard Business School Open Forum*, April 14, 2015.
- Baldwin H (2015) For big data, it was the best of times, it was the worst of times. Available at www.forbes.com/sites/howardbaldwin/2015/04/06/for-big-data-it-was-the-best-of-times-it-was-the-worst-of-times/#165962c62e20. Accessed June 11, 2015.
- Ostrovsky M, Schwarz M (2011) Reserve prices in internet advertising auctions: A field experiment. *Proceedings of the 12th ACM Conference on Electronic Commerce*: 59–60.
- Lewis R, Reiley D (2014) Online ads and offline sales: Measuring the effect of retail advertising via a controlled experiment on Yahoo! *Quant Mark Econ* 12(3):235–266.