

WELFARE PROPERTIES OF PROFIT MAXIMIZING RECOMMENDER SYSTEMS: THEORY AND RESULTS FROM A RANDOMIZED EXPERIMENT¹

Xiaochen Zhang

Uber, 1455 Market Street # 400,
San Francisco, CA 94103 U.S.A. {xiaochen.zhang@gmail.com}

Pedro Ferreira

Heinz College, Carnegie Mellon University, 5000 Forbes Avenue,
Pittsburgh PA 15213 U.S.A. {pedro@cmu.edu}

Miguel Godinho de Matos

Universidade Católica Portuguesa, Católica Lisbon School of Business and Economics,
Palma de Cima, 1649-023 Lisboa PORTUGAL {miguel.godinhomatos@ucp.pt}

Rodrigo Belo

Rotterdam School of Management, Erasmus University, Burgemeester Oudlaan 50,
3062 PA Rotterdam, THE NETHERLANDS {rbelo@rsm.nl}

Recommender systems have been introduced to help consumers navigate large sets of alternatives. They usually lead to more sales, which may increase consumer surplus and firm profit. In this paper, we ask whether the firms' choice of recommender system might hurt consumers. We use data from a large-scale field experiment in video-on-demand to measure the price elasticity of demand for movies placed in salient and non-salient slots on the TV screen. During this experiment, the firm randomized the prices and slots in which movies were recommended to consumers. This setting readily allows for identifying the effects of price and slot on demand, and thus computes consumer surplus. We find that consumers are less price-elastic toward movies placed in salient slots. Using the outcomes of this experiment, we simulate how consumer surplus and welfare change when the firm implements several types of recommender systems, namely one that maximizes profit. We show that this system hurts consumer surplus and welfare relative to a system that maximizes welfare. We also show that, in our setting, the system that maximizes profit does not generate less consumer surplus than several recommender systems frequently used in practice. Yet, the amount of extra rent the firm can extract from strategically placing movies in salient slots is still a function of the popularity and quality of movies used to do so. Ultimately, our results question whether recommender systems embed mechanisms that extract excessive surplus from consumers, which may call for better scrutiny.

Keywords: Recommender systems, welfare properties, randomized experiment

¹Bin Gu was the accepting senior editor for this paper. Ofir Turel served as the associate editor.

Introduction

Digitization significantly increased the number of products and services available to consumers. In this context, recommender systems became a pervasive tool to help navigate large catalogs. The very first implementations of these systems were aimed at maximizing matching using heuristics. For example, content-based recommender systems would suggest items similar to the ones that the consumer previously bought and collaborative filters would suggest items bought by consumers with similar purchasing history. Several papers showed that the introduction of these systems allowed consumers to more easily find niche products, extending and fattening the distribution of sales. This effect has been empirically observed across industries, in particular in the case of experience goods (Brynjolfsson et al. 2006; Dewan and Ramaprasad 2012; Elberse and Oberholzer-Gee 2007; Oestreicher-Singer and Sundararajan 2012, Tucker and Zhang 2011) and there is evidence that it increases the surplus of consumers (Brynjolfsson et al. 2003, 2010). Other works, using modeling and simulation, show that the effect of recommender systems on the distribution of sales depends essentially on how they manipulate product saliency (Fleder and Hosanagar 2009) and that the effect of these systems on consumer surplus and on welfare may be ambiguous (Choudhary and Zhang 2016; Fleder and Hosanagar 2009). Recommender systems may also affect consumer preferences through anchoring, which may hamper maximizing their surplus (Adomavicius et al. 2013).

The question of how recommender systems affect consumers becomes even more important as these systems mature and start allowing firms to introduce elements of profit maximization into the original implementations (Jannach and Adomavicius 2017). Today, firms use the suggestions from content-based and collaborative filtering recommender systems to predetermine sets of items that are highly relevant to consumers but then order them in ways that may maximize profit (Das et al. 2009) and ensure some level of consumer trust leading to repeated purchases. However, little is known about the welfare properties of these profit-aware recommender systems. This is not surprising because instantiating such a system requires firms to know the demand curve for their products. A firm that does not know this curve is likely to simply order suggestions to consumers in decreasing order of similarity, which is a rule of thumb often used in practice that is believed to work well. A firm that knows the demand curve may instead use similarity measures to identify products that are likely to align with the preferences of consumers, and thus help maintain trust, and then use the demand curve to reorder suggestions to maximize profit. Our paper starts by introducing an analytical model showing conditions under which recommending items to consumers to maximize profits hurts

the surplus of the latter. We then show how a firm can use a randomized control trial to econometrically identify the demand curve and, finally, we show what happens to the profit of the firm and to the surplus of consumers when the firm uses this knowledge strategically.

We estimate the demand curve for movies, which is an experience good, in the Video-on-Demand (VoD) system of a large telecommunications provider and simulate the behavior of consumers when they face suggestions from different types of recommender systems. We are primarily interested in comparing the consumer surplus obtained by the recommender system that maximizes the profit of the firm and that obtained by the recommender system that maximizes consumer surplus in order to learn whether they are different and, if so, by how much. We design, implement, and analyze outcomes from a randomized field experiment with more than 600,000 households over a period of half a year. We identify the price elasticity of demand in this setting, which is the parameter of interest that shapes the demand curve allowing us later to compute the appropriate counterfactuals for welfare analysis. Recommender systems operate by manipulating the saliency across products. They highlight some products, making others relatively more hidden. Therefore, in our context, we model the demand for a movie as a function of price and saliency. Specifically, saliency is associated with the slot in which movies are shown to consumers on their TV screen. In the country we study, people read from left to right and, therefore, placing movies toward the left of the TV screen primes consumers to pay more attention to them. We randomize independently both the slot and the price of each and every movie suggested to consumers during our experiment, which allows us to identify the price elasticity of demand for movies placed in salient and non-salient slots. We show empirical evidence that, in our setting, consumers exhibit a lower price elasticity of demand for movies placed in the slots toward the left of the TV screen (irrespective of the movies placed in these slots). These movies are highlighted to consumers and, as a consequence, consumers are more tolerant when their prices increase. Different recommender systems—operationalized by manipulating differently the slots where movies are shown—yield different levels of sales, profits, and consumer surplus that we compare using simulations.

The difference in the price elasticity of demand between movies placed in slots toward the left and toward the right of the TV screen allows the firm to extract surplus from consumers. We show that under certain conditions this difference leads the firm to suggest movies to consumers to maximize profit in a way that is different from the way that consumers would like to receive recommendations in order to maximize consumer surplus. In our setting, and on average across our simulations, the consumer surplus generated by the recom-

mender system that maximizes profit is 8.44% lower than the consumer surplus that would be attained by the recommender system that would maximize consumer surplus. Total welfare may also reduce when the firm maximizes profit. Across our simulations, and on average, when the firm maximizes profit, welfare reduces 2.07% relative to the welfare that would be attained by the recommender system that would maximize welfare. We also show that for certain subsets of movies used during our experiment, these losses can be as high as 28.7% and 12.2%, for consumer surplus and welfare respectively. We also study what may happen when consumers interact repeatedly with the recommender system. We discuss that the degree to which recommender systems hurt consumer surplus is, in part, determined by whether or not consumers realize that sometimes they are suggested “second best” products. How exactly this occurs in real world settings is still an open question. For example, only by incurring additional search costs can consumers know whether there are better products for them. However, not only do those that search further become a self-selected sample, but also, by doing so, their surplus is already reduced due to the additional search costs. Firms may also use consumer feedback, such as number of likes and reviews, to learn more about consumer enjoyment, but again such signals are noisy, scant, and come from self-selected samples of consumers. Finally, we also provide empirical evidence that, in our setting, consumers exhibit a lower price elasticity of demand toward the movies placed in salient slots immediately after the experiment started and that this effect reduces over time. Therefore, in our setting, consumers seem to be able to learn how to discern good movies from bad movies that are placed in salient slots by our exogenous manipulations.

We expect that our measurements of the welfare inefficiency associated with profit-aware recommender systems will trigger additional research on this subject, and in particular research aimed at understanding the relative magnitude of these losses *vis-à-vis* those introduced by other mechanisms that firms typically use to extract surplus from consumers, such as price discrimination. This type of research is also likely to better inform us about whether recommender systems need scrutiny to protect consumers.

Literature Review

Digitization changed many industries by lowering production, distribution, and consumption costs (Bakos 1997). This reduction in costs allowed for an unprecedented increase in product availability and variety, which is likely to attract more consumers (Anderson 2006; Brynjolfsson et al. 2003; Clemons et al. 2006). However, large sets of alternatives

might increase search costs disproportionately, which, under some conditions, may even lead to less choice (Kuksov and Villas-Boas 2010). Digital recommender systems have been introduced to help consumers navigate these large sets of alternatives (Resnick and Varian 1997). At their core, these systems change product saliency. They highlight some products to consumers at the expense of other products that become relatively more hidden. The early work on recommender systems has shown that, in general, these systems allow consumers to more easily find niche products online, reducing the concentration of sales (Brynjolfsson et al. 2011).

However, and ultimately, what happens to the distribution of sales depends heavily on how recommender systems affect product saliency. For example, content-based recommender systems, which increase the saliency of products similar to the products that consumers have already purchased in the past, tend to reduce the concentration of sales. Several authors find empirical evidence of this fact, for example at Amazon’s bookstore (Oestreicher-Singer and Sundararajan 2012), in the context of wedding service vendors (Tucker and Zhang 2011), in the demand for niche home video products (Elberse and Oberholzer-Gee 2007), and in the sales of music (Chellappa et al. 2007; Dewan and Ramaprasad 2012). On the other hand, collaborative-filtering recommender systems, which increase the saliency of the products bought by similar consumers, may instead increase the concentration of sales because they disproportionately recommend the most popular products (Fleder and Hosanagar 2009). Self-reinforcing mechanisms that increase the popularity of already popular artists leading to a super-star effect have been both proposed theoretically (Frank and Cook 1995; Rosen 1981) and observed empirically (Salganik and Watts 2008). For in-depth reviews of the types of recommender systems currently available and how they work, refer to Adomavicius and Tuzhilin (2005) and Breese et al. (1998).

Knowing how the introduction of recommender systems changes the distribution of sales is useful from a managerial point of view, in particular for inventory management and contract negotiation. However, this tells us very little about changes in the profit of firms or in the welfare of consumers. A couple of market-level empirical studies provide some evidence that the long-tail effect at Amazon’s bookstore increased consumer surplus (Brynjolfsson et al. 2003, 2010). However, theoretically, recommender systems with low accuracy may also decrease consumer surplus (Choudhary and Zhang 2016). In addition, simulations have also been used to show that the effect of recommender systems on welfare is hard to anticipate (Fleder and Hosanagar 2009). On the one hand, recommender systems help consumers find products that may better match their preferences, which were hard to find before. This is likely to increase consumer surplus. On

the other hand, recommender systems may also lead consumers to purchase products that otherwise they would not want through, for example, anchoring (Adomavicius et al. 2013). In sum, the prior literature lacks individual-level empirical studies looking at the effect of different recommender systems on firm profit and consumer welfare.

The prior literature has also focused on whether saliency changes the likelihood of purchase. A couple of seminal works show that popularity is self-reinforcing for middle quality songs (Salganik et al. 2006) and that consumers are able to correct manipulations to the ratings of very bad and very good songs (Salganik and Watts, 2008). These studies look at how saliency shapes consumer behavior, but no prices were involved. The subjects could download songs for free, which hampers the generalization of these results to real-world settings. More recently, priming effects have been shown to be asymmetric in the case of a news aggregator website (Muchnik et al. 2013). Again, this is an example in which the researchers manipulated saliency, but no prices were involved, which limits the ability to learn how priming can actually be used in a profit maximizing context. In a setting closer to ours, Godinho de Matos et al. (2016) partnered with a large telecommunications company to analyze the impact of peer ratings in a real-world VoD market. After experimentally changing the number of likes displayed to consumers, the authors found that likes influence consumer behavior independently from the underlying product quality. However, they also found little evidence of long-term bias as a result of priming effects, at least in their setting. Specifically, when movies were artificially promoted or demoted in peer rating lists, subsequent likes caused them to return to their true quality position relatively quickly. Again, this is an example in which the authors manipulate saliency but leave prices untouched, and thus they could not appreciate the tradeoff between saliency and prices, which eventually prevented them from discussing how firms can strategically manipulate saliency to increase profits and how such behavior may hurt consumers.

The prior literature has also established that humans are not perfect at maximizing their welfare. They make mistakes when making decisions and are primed to make certain types of decisions that might not always enhance their welfare. In particular, the way messages are framed and how they are delivered to individuals makes a difference. As a consequence, firms may use messages to strategically take advantage of the inefficiencies associated with consumer choice (Thaler and Sunstein 2003, 2008). This line of research is concerned with the design of specific nudges that steer consumer behavior. Recommender systems are a specific type of nudge. They affect the search cost of the recommended

products (Pathak et al. 2010) and affect individual *ex ante* expectations about the *ex post* enjoyment from consuming the recommended products. The gap that we identify in the prior literature is the absence of a clear, direct link between saliency and price elasticity of demand, which would provide a clear measure of how priming consumers changes their willingness to pay, and, from there, how it affects their well-being.

Our work is also related to prior literature in marketing by looking at how both price and non-price oriented advertising affect the sensitivity of consumers. Results from a meta-analysis conclude that retail-oriented price advertising tends to increase price sensitivity whereas non-price advertising tends to decrease price sensitivity (Kaul and Wittink 1995). Yet, such an aggregate meta-analysis may hide a number of complex effects that may move price elasticity in different directions. For example, effects may be heterogeneous with respect to the brands' baseline market shares (Kalra and Goodstein 1998). In addition, there is also debate about the effect of advertising on the price elasticity of demand. Some researchers argue that advertising decreases price elasticity because it enhances market power, while other researchers argue the opposite because advertising provides more information about substitutes (Mitra and Lynch 1995). A framework to reconcile these findings has been proposed where advertising affects price elasticity through two mechanisms: the size of the consideration set and the relative strength of consumer preferences (Mitra and Lynch 1995). On the one hand, advertising enlarges the consideration set leading consumers to consider more substitutes, which increases price elasticity. On the other, advertising provides additional information about the products, allowing consumers to reduce uncertainty and to decisively prefer specific products for which they may be willing to pay more, thus reducing the price elasticity of demand. Recommender systems are a specific tool to advertise some products more than others and often are used to convey information to consumers about substitutes.

Theoretical Model

Setup

Consider a representative consumer, two products indexed by j , called H and L , and two slots, a salient slot and a non-salient slot, where these products will be placed by a monopolistic firm. Let d_j represent the demand for product j and let s_j represent whether product j is placed in the salient slot. We

have $(s_H, s_L) \in \{(0, 1); (1, 0)\}$. Let p_j represent the price of product j and define

$$d_j(p_j; s_j) = p_j^{\beta_0 + \beta_1 s_j}$$

This standard log-linear functional form is commonly used in the literature to specify the demand for digital goods, such as in the seminal work by Brynjolfsson et al. (2003). The elasticity of demand for the product placed in the non-salient slot is β_0 and the elasticity of demand for the product placed in the salient slot is $\beta_0 + \beta_1$. Without loss of generality, let $\beta_1 > 0$. This entails that consumers are less price elastic with respect to products placed in the salient slot. Let c_H and c_L represent the marginal cost of products H and L , respectively. Without loss of generality, let $0 < c_L < c_H$. Let $m_j = p_j(s_j) - c_j$ represent the margin associated to product j . In this setup, the profit of the firm and the consumer surplus are given by

$$\begin{aligned} \Pi(p_H, p_L; s_H, s_L) &= \sum_{j=H,L} d_j(p_j; s_j)(p_j(s_j) - c_j) \\ &= \sum_{j=H,L} m_j p_j^{\beta_0 + \beta_1 s_j} \end{aligned}$$

$$\begin{aligned} CS(p_H, p_L; s_H, s_L) &= \sum_{j=H,L} \int_{p_j}^{+\infty} d_j(p; s_j) dp \\ &= \sum_{j=H,L} \frac{p_j^{1 + \beta_0 + \beta_1 s_j}}{1 + \beta_0 + \beta_1 s_j} \end{aligned}$$

Profit Maximization

For sake of simplicity assume price elasticities of demand less than -1 , that is, $1 + \beta_0 < 0$ and $1 + \beta_0 + \beta_1 < 0$. Given an allocation of products to slots the firm sets prices to maximize total profit, which yields

$$p_j^*(s_j) = c_j \frac{\beta_0 + \beta_1 s_j}{1 + \beta_0 + \beta_1 s_j}, m_j^*(s_j) = -c_j \frac{1}{1 + \beta_0 + \beta_1 s_j}$$

Note that in this setup, conditional on slot, the product with a higher cost yields a higher margin. The following profit arises for each allocation of products to slots:

$$\begin{aligned} \Pi(p_H^*, p_L^*; s_H = 1, s_L = 0) &= -\frac{(\beta_0 + \beta_1)^{\beta_0 + \beta_1}}{(1 + \beta_0 + \beta_1)^{1 + \beta_0 + \beta_1}} c_H^{1 + \beta_0 + \beta_1} \\ &\quad - \frac{\beta_0^{\beta_0}}{(1 + \beta_0)^{1 + \beta_0}} c_L^{1 + \beta_0} \end{aligned}$$

$$\begin{aligned} \Pi(p_H^*, p_L^*; s_H = 0, s_L = 1) &= -\frac{(\beta_0 + \beta_1)^{\beta_0 + \beta_1}}{(1 + \beta_0 + \beta_1)^{1 + \beta_0 + \beta_1}} c_L^{1 + \beta_0 + \beta_1} \\ &\quad - \frac{\beta_0^{\beta_0}}{(1 + \beta_0)^{1 + \beta_0}} c_H^{1 + \beta_0} \end{aligned}$$

RESULT 1: The firm places the product with the highest cost in the salient slot if both costs are lower than a threshold F , i.e., if $c_H < F$, with

$$F = \frac{1 + \beta_0 + \beta_1}{\beta_0 + \beta_1} \left[\frac{\beta_0 (1 + \beta_0 + \beta_1)}{(1 + \beta_0)(\beta_0 + \beta_1)} \right]^{\frac{\beta_0}{\beta_1}}$$

PROOF: Let

$$\begin{aligned} &\Pi(p_H^*, p_L^*; s_H = 1, s_L = 0) - \\ &\Pi(p_H^*, p_L^*; s_H = 0, s_L = 1) = f(c_H) - f(c_L) \end{aligned}$$

in which $f(\cdot)$ is defined as

$$f(c) = -\frac{(\beta_0 + \beta_1)^{\beta_0 + \beta_1}}{(1 + \beta_0 + \beta_1)^{1 + \beta_0 + \beta_1}} c^{1 + \beta_0 + \beta_1} + \frac{\beta_0^{\beta_0}}{(1 + \beta_0)^{1 + \beta_0}} c^{1 + \beta_0}$$

and thus

$$f'(c) = -\left(\frac{\beta_0 + \beta_1}{1 + \beta_0 + \beta_1} \right)^{\beta_0 + \beta_1} c^{\beta_0} \left[c^{\beta_1} - \left(\frac{1 + \beta_0 + \beta_1}{\beta_0 + \beta_1} \right)^{\beta_0 + \beta_1} \left(\frac{\beta_0}{1 + \beta_0} \right)^{\beta_0} \right]$$

Therefore, $f(\cdot)$ is increasing when $c < F$ and decreasing otherwise. Consequently, $f(c_H) > f(c_L)$ when $v_{c_H} < F$. Thus, in this case, $\Pi(p_H^*, p_L^*; s_H = 1, s_L = 0) > \Pi(p_H^*, p_L^*; s_H = 0, s_L = 1)$ and the firm places the product with the highest cost in the salient slot. Similarly, the firm places the product with lower cost in the salient slot when $c_L > F$.

The intuition behind this result is as follows: $f(c)$ can be interpreted as the extra profit originated from promoting a product with cost c from the non-salient slot to the salient slot, ceteris paribus. In other words, $f(c)$ is the difference between the product-level profit curve in the salient slot and the product-level profit curve in the non-salient slot for the same product. This function has a maximum at F , meaning that a (hypothetical) product with cost F would be the best product to promote to the salient slot. This is the product for which the difference in profits between being in the salient slot and in the non-salient slot is maximized, ceteris paribus. Note that whenever a product is promoted to the salient slot another product needs to be demoted. In any case, promoting (demoting) any other product with a cost different than F to (from) the salient slot has a lower impact on profit. Therefore, the best strategy for the firm is to place the product with

cost closest to F in the salient slot: when the cost of both products is lower than F , then the product with highest cost (closest to F) should be placed in the salient slot; when the cost of both products is higher than F , the product with lowest cost (closest to F) should be placed in the salient slot. More information would be required to determine the best configuration when F is higher than the cost of one product but lower than the cost of the other product. Result 2 relates the gain in profit from profit maximization with β_1 :

RESULT 2: The gain in profit when profit is maximized increases with β_1 .

PROOF: See proofs for Results 2–8 in Appendix A.

Effect on Consumer Surplus

The following consumer surplus arises when the firm maximizes profit:

$$CS(p_H^*, p_L^*; s_H = 1, s_L = 0) = \frac{(\beta_0 + \beta_1)^{1+\beta_0+\beta_1}}{(1 + \beta_0 + \beta_1)^{2+\beta_0+\beta_1}} c_H^{1+\beta_0+\beta_1} - \frac{\beta_0^{1+\beta_0}}{(1 + \beta_0)^{2+\beta_0}} c_L^{1+\beta_0}$$

$$CS(p_H^*, p_L^*; s_H = 0, s_L = 1) = \frac{(\beta_0 + \beta_1)^{1+\beta_0+\beta_1}}{(1 + \beta_0 + \beta_1)^{2+\beta_0+\beta_1}} c_L^{1+\beta_0+\beta_1} - \frac{\beta_0^{1+\beta_0}}{(1 + \beta_0)^{2+\beta_0}} c_H^{1+\beta_0}$$

RESULT 3: The consumer surplus is maximized when the firm places the product with the highest cost in the non-salient slot if $c_L > G$ with

$$G = \frac{1 + \beta_0 + \beta_1}{\beta_0 + \beta_1} \left[\frac{\beta_0 (1 + \beta_0 + \beta_1)}{(1 + \beta_0)(\beta_0 + \beta_1)} \right]^{\frac{1+\beta_0}{\beta_1}}$$

The intuition behind this result is analogous to that behind Result 1: $g(c)$, defined in Appendix A analogously to $f(c)$, can be interpreted as the additional consumer surplus originated from promoting a product with cost c from the non-salient slot to the salient slot, *ceteris paribus*. Like in the case of Result 1, G corresponds to the cost of the product that maximizes the difference in consumer surplus generated by placing that product in the salient slot and placing it in the non-salient slot. Thus, the strategy that maximizes consumer surplus places the

product with cost closest to G in the salient slot. Combining Result 1 and Result 3 leads to the following result:

RESULT 4: The allocation of products to slots that maximizes profit does not maximize consumer surplus when $G < c_L < c_H < F$.

Result 4 shows that when the cost of the products falls in the (G, F) region a conflict arises between the way that the firm recommends products to maximize profits and the way in which consumers would like to be recommended products to maximize consumer surplus. Figure 1 illustrates this conflict. Within the shaded area maximizing profit requires placing the product with the highest cost in the salient slot, while maximizing consumer surplus requires placing the product with the lowest cost in this slot. Outside this area there is no conflict; placing the product with the highest cost in the salient slot maximizes both profit and consumer surplus when costs are lower than G , and placing the product with the lowest cost in this slot maximizes both profit and consumer surplus when costs are higher than G . It is clear that no conflict arises when $\beta_1 = 0$ because in this case both slots are similar. The results that follow relate the size of the conflict region and the magnitude of the loss in consumer surplus with the difference in the price elasticity of demand between slots.

RESULT 5: The size of the conflict region for consumer surplus increases with β_1 .

RESULT 6: The loss in consumer surplus when profit is maximized increases with β_1 .

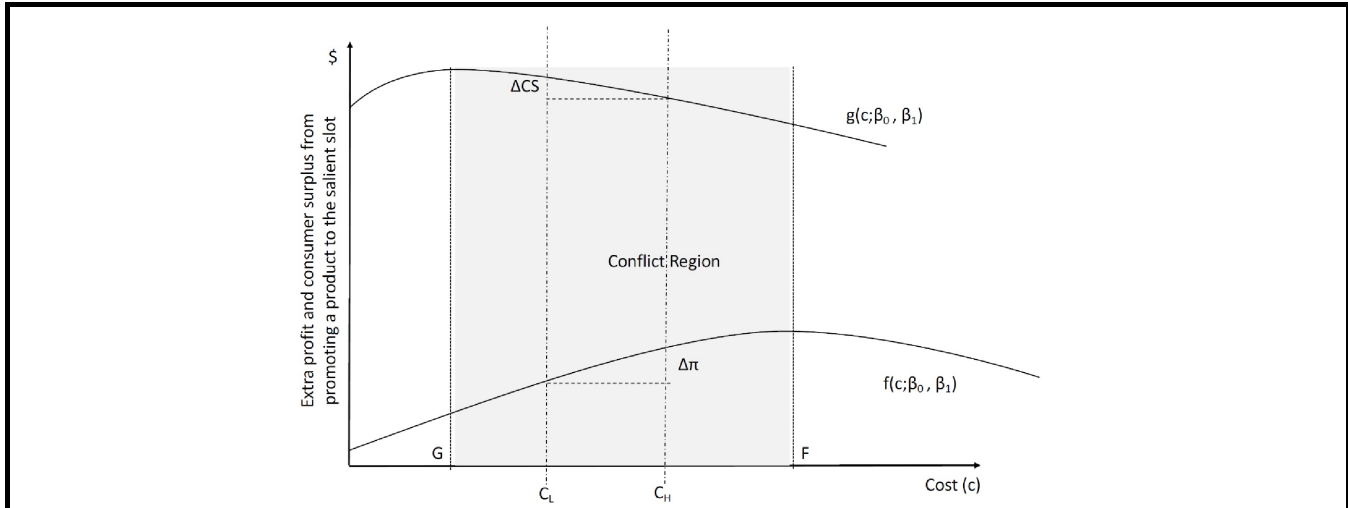
The previous results show the nuclear role that β_1 plays in our setting. β_1 measures the difference in the price elasticity of demand across slots in the recommender system and it is the reason why conflicts arise between the firm and consumers with respect to how recommendations should be issued. Furthermore, the larger this difference the more likely conflicts are to arise and the larger they are in magnitude. In sum, β_1 embodies the mechanism by which conflicts arise and thus our empirical application later in this paper will be used to identify this parameter.

Effect on Total Welfare

Define total welfare as the sum of consumer surplus and profits $W(p_H^*, p_L^*; s_H, s_L) = \Pi(p_H^*, p_L^*; s_H, s_L) + CS(p_H^*, p_L^*; s_H, s_L)$.

The following results follow:

RESULT 7: Total welfare is maximized when the firm places the product with the highest cost in the non-salient slot when $c_L > H$.



The extra profit earned by the firm from maximizing profit, compared to the recommender system that maximizes consumer surplus, is given by $f(c_H) - f(c_L)$. The corresponding loss in customer surplus is given by $g(c_H) - g(c_L)$.

Figure 1. Illustration of the Conflict Between the Maximization of Profit and the Maximization of Consumer Surplus

RESULT 8: The allocation of products to slots that maximizes profit does not maximize total welfare when $F > c_H > c_L > H$.

Result 7 shows that when the marginal cost of the products falls in the (H, F) region, a conflict arises between the order in which the firm recommends products to maximize profits and the order in which these products should be recommended to consumers to maximize total welfare. Therefore, the selfish behavior of the firm with respect to the order in which products are recommended to consumers may even hurt welfare besides hurting consumer surplus. We also note that this region where welfare decreases is a subset of the region previously identified where consumer surplus decreases. This arises because $G < H$, which follows immediately from the fact that

$$H = G \left[\frac{(1 + 2\beta_0)(\beta_0 + \beta_1)}{\beta_0(1 + 2\beta_0 + 2\beta_1)} \right]^{\frac{1}{\beta_1}},$$


$$\beta_1 > 0, 1 + \beta_0 < 0, 1 + \beta_0 + \beta_1 < 0$$

This section shows that under certain conditions profit maximizing firms recommend products to consumers in ways that hurt consumer surplus. In particular, this result arises when consumers exhibit different price elasticities of demand with respect to the different slots that the firm uses to recommend products. Appendix C shows numerical examples that illustrate our results. The first example considers only one repre-

sentative consumer. The second example introduces a taste parameter allowing for two types of consumers. Both cases show how loss in consumer surplus arises. The theoretical model developed in this section considers only one representative consumer, which is enough to show that a profit maximizing firm recommends products to consumers in ways that may hurt consumer surplus. This model can be easily extended to multiple consumers, which would essentially represent additional independent markets, each of which could generate loss in consumer surplus. Finally, we note that there are significant economic instances in which consumers use the recommender system only once, or only once from time to time with significantly large inter-visit times. This happens, for example, when purchasing expensive products, which consumers are more likely to buy only occasionally, such as cars and houses. Therefore, the results presented so far are informative in several occasions. In any case, in the next subsection we analyze what happens when consumers interact repeatedly with the recommender system.

Dynamics with Repeated Interactions

Consider now two consecutive time periods in which the representative consumer interacts with the recommender system. In this case, consider $\beta_{0,t}$ and $\beta_{1,t}$ for period t , $\beta_{0,t+1}$ and $\beta_{1,t+1}$ for period $t + 1$ and assume that β_1 reduces from t to $t + 1$ when there is loss in consumer surplus at time t , that is, consumers become more flexible toward the product placed

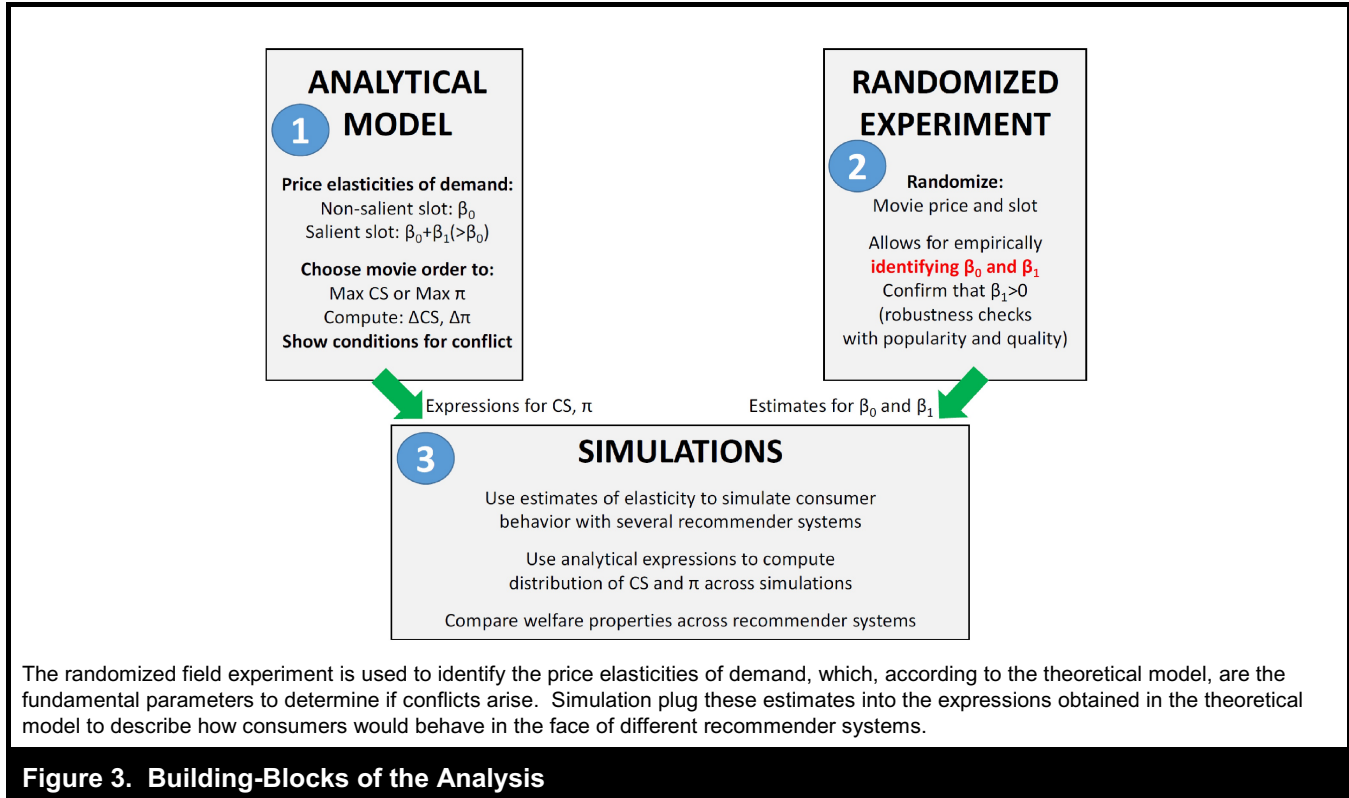


in the salient slot when they were less so in the past and lost surplus. Figure 2 shows how this reduction in β_1 changes the conflict region from period t to period $t + 1$. Result 5 implies that this region shrinks. Result 6 implies that the loss in consumer surplus in period $t + 1$ will be smaller than that experienced in period t . Likewise, Result 2 implies that the gain in profit in period $t = 1$ will be smaller than that enjoyed by the firm in period t .

Generalizing the rationale above, β_1 reduces over time yielding lower profits for the firm and lower loss in consumer surplus with each interaction. Furthermore, this dynamic process embodies loss of surplus for consumers at each and every interaction. This result comes in line with the findings by Godinho de Matos et al. (2016), where the authors experimented with a video-on-demand system and found that movies that have been promoted (or demoted) to a better (worse) slot return to their natural slot after a dynamic process during which consumers issue less (more) likes to promoted (demoted) movies. Also in this case, the process by which movies converge back to their natural slots is achieved at the expense of consumers who find that bad movies have been promoted to good slots and that good movies have been demoted to bad slots. Whether β_1 converges to zero cannot be unequivocally determined in our model. If β_1 reaches zero, then the conflict region becomes inexistent. In this case, consumers enjoy the maximum surplus possible and the firm collects the profit associated to that equilibrium. If β_1 does not reach zero then the conflict region is sustained by other

mechanisms such as, for example, search costs. If the non-salient slot embodies higher search costs irrespective of the quality of the products placed in both the salient and in the non-salient slot then consumers may prefer to purchase from the salient slot, even if the product placed there is of lower quality and they know about it, because they save the search cost. This would result in a positive β_1 that would likely to sustained over time.

In practice, whether consumers perceive that they may have been shown only a “second best” product in the salient slot is far from trivial because firms are unlikely to implement recommender systems that maximize consumer surplus and thus it is hard for consumers to directly observe the ideal world for them. In addition, there might be several roughly equally good products to the consumer. The firm will choose the one with the highest margin for the salient slot, but the consumer will have a hard time to tell roughly equally good products apart. This may allow the firm, in practice, to sustain extra rents over time. Finally, note that each new consumer needs to go through the adjustment process described before whereby β_1 keeps reducing. Therefore, and in a real-world setting, the extra profit that the firm can collect from strategically allocating products to slots is, in part, shaped by the mix of consumers that use the recommender system in terms of new consumers (who still need to adjust and thus still exhibit a higher β_1) versus repeated consumers (who have already adjusted and thus exhibit a smaller β_1).



Our Work and Our Experimental Context

Research Approach

Figure 3 illustrates the research approach that we follow to study the welfare properties of recommender systems. The first step in this approach was introduced in the previous section and includes the theoretical model showing conditions under which there is a conflict between the order in which the firm wants to suggest products to consumers and the order in which consumers would like to receive these recommendations. This model shows that conflicts arise due to differences in price elasticity of demand between the salient and the non-salient slot. Therefore, the second step in our work is to run a large-scale randomized field experiment to measure these elasticities and test whether they are different. Below, we introduce our industry partner and describe in detail the experiment that we run to accomplish this goal. Finally, in the third step of our work, we plug in the estimates for the price elasticities of demand identified using our field experiment into the expressions for profit, consumer surplus and welfare derived in our theoretical model, and provide simulation results showing how different recommender systems compare in terms of how much surplus they generate.

Our simulations compute how consumers would behave in face of different recommender systems, which allow us to show how consumer surplus is affected when the firm maximizes profit *vis-à-vis*, for example, when the recommender system is setup to maximize the former. As discussed before, it is unlikely that, in practice, firms arrange recommendations to only maximize short-term profits because they need to maintain consumer trust across repeated interactions. Also, they are unlikely to arrange recommendations in ways that fully maximize consumer surplus. Therefore, when we compare these recommender systems we are, in fact, providing bounds for the loss in consumer surplus that consumers may experience in practice. This exercise is similar to that performed in economic analysis when one compares the consumer surplus associated to a monopoly (that chooses price to maximize profit) and that associated to the competitive market (in which price is set to maximize consumer surplus). Also, as in such cases, our measure of consumer surplus reflects everything that may affect the consumers' willingness to pay for the product. For example, some firms are, at times, interested in improving the "user experience" as a way to ensure that consumers keep coming back. User experience is captured by consumer surplus to the extent that consumers are willing to pay more for a better user experience. In other words, using consumer surplus as a measure of consumer

well-being captures everything about the latter except traits that do not affect the consumer's willingness to pay.

Our Industry Partner

Our industry partner, hereinafter called TELCO, is a major multinational telecommunications provider that offers TV, Internet, telephony and Video-on-Demand (VoD). It serves more than 1 million households, 69% of which purchase triple play bundles. Households subscribe to either standard or premium service. The latter is more expensive. It includes more TV channels and better Internet connection speeds. We had access to TELCO's VoD database between June 2013 and March 2014. Roughly half of the active households subscribe to VoD, of which a fifth paid for VoD content at least once during this 9-month period. During this period, we observe roughly 800 thousand movie leases. We have the anonymized identifier of the household requesting each transaction through the anonymized identifier of the MAC address of the corresponding Set-Top Box (STB). For each transaction, we have the identifier of the movie leased, the price of the movie, and a timestamp for when the lease took place. We also have a number of covariates for each movie in TELCO's database such as title, director, studio, play length, synopses, cast, and genre.

Figure 4 shows the look and feel of the TV screen in this VoD system, which consumers can access using a hot-key in their STB remote control. This screen, called the Highlights Section, contains a set of menus filled with movies chosen by an editorial team. Movies are organized into menus such as promotions, suggestions, dramas, and comedies. Each menu has a header with a name that clearly identifies the type of movies it contains. Menus are horizontal lines stacked vertically on the TV screen. Two menus fit the TV screen at any time and users can scroll up and down across menus. Upon scrolling to a new menu, the cursor highlights the movie farthest to the left. Users can scroll left and right across movies under the same menu. They can also scroll past the last movie cover on the right of the TV screen to unveil hidden movies. TELCO displays 15 movies per menu and 11 menus in the highlights section. The title of the movie highlighted by the cursor is shown toward the bottom of the screen. The number of likes of the movie is shown close to the title for premium consumers. Clicking on the cover of a movie leads to a new screen with the year of release, play length, cast and synopsis. A number of actions are available from this new screen such as lease the movie, use a promotional coupon to lease the movie, or watch the movie trailer (if one is available). Leased movies can be watched only for a period of 48 hours. Premium consumers can also issue likes for movies.

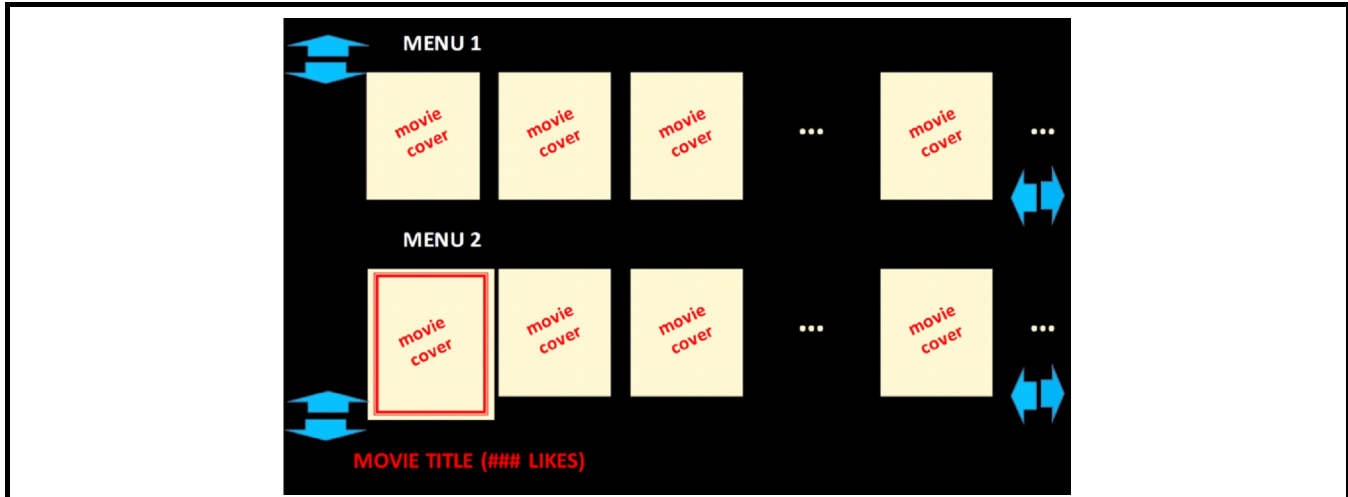
Figure 4 also illustrates the fact that, under each menu, slots toward the left of the TV screen are more salient than slots toward the right of the screen. This is a known fact in countries where consumers read from left to right (Godinho de Matos et al. 2016). These movies are more prominent to the eyes of consumers and, for example, consumers may be more tolerant when their prices increase. Finally, users can leave the highlights section of the VoD interface and search for movies in the complete catalog. The catalog is hierarchically organized into content categories such as movies, documentaries, and TV shows. Within each of these categories, screens are organized as described above with menus for genres. All households see the same number of movies both in the highlights section as well as in the full catalog. TELCO's full VoD catalog offers access to more than 2,000 titles at any point in time.

Randomized Experiment

The experimental design followed in our study is illustrated in Figure 5. All households at TELCO with VoD service were included in the experiment. A new menu, called "Good Deals" was introduced into the highlights section of the VoD system. This menu was accessible to both standard and premium households with three clicks down on the remote control. At the beginning of the experiment, households were randomly named "original" or "discount." During the experiment, the former households got movies under the new menu at their original price, which was decided by TELCO. The latter households got movies under the new menu at discounted prices. The discounted price was the marginal cost of the movie for TELCO, which was negotiated between TELCO and content owners. The average discount was 25% across the movies used during this experiment.

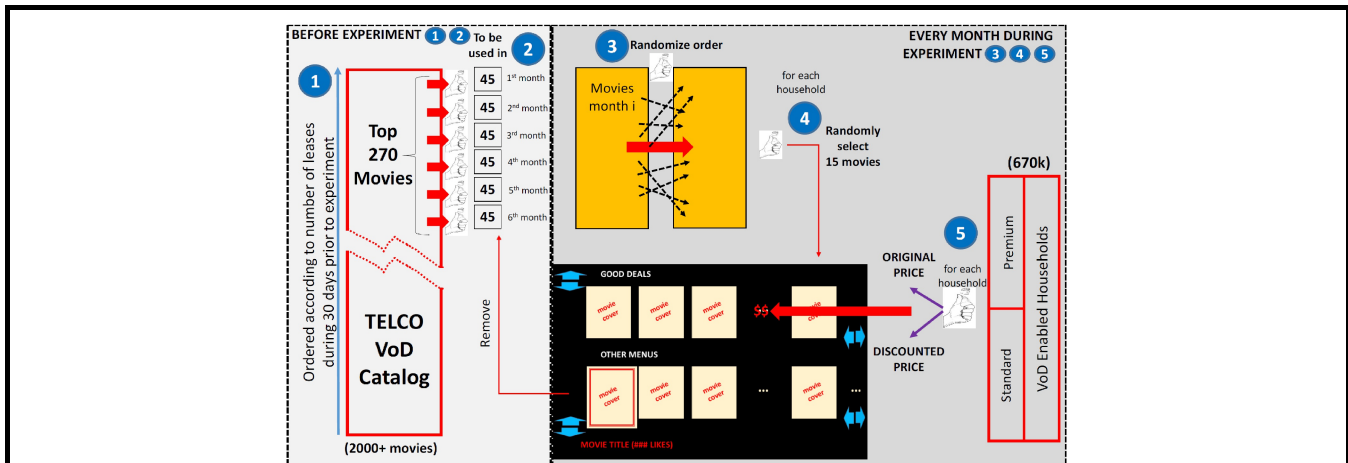
Before the experiment started, all movies in TELCO's VoD catalog were ordered according to sales during the previous 30 days and the 270 movies at the top of this list were used in the experiment. This is represented by step 1 in Figure 5. This set of 270 movies was randomly subsetted into 6 non-overlapping sets of 45 movies each of them to be used during each month of the experiment. This is represented by step 2 in Figure 5. Therefore, a different subset of the most sold movies in TELCO's VoD catalog during the 30 days before our experiment started was used each month in our experiment. Let M_t represent the subset of movies used in month t .

Then, a number of operations were repeated on a monthly basis. For month t the movies in M_t that were not shown in other VoD menus (to avoid double counting) were randomly ordered. This is represented by step 3 in Figure 5. Let OM_t represent this randomly ordered list of movies. Then, for each household separately, a random subset of 15 movies in OM_t



Menus are stacked vertically on the screen and movies are shown horizontally under each menu. The title of the movie highlighted by the cursor shows on the bottom line of the screen (alongside the number of likes for premium consumers).

Figure 4. Illustration of the TV Screen for the Video-on-Demand System at TELCO.



All VoD-enabled households at TELCO were included in the experiment as well as the 270 movies with more sales during the 30 days before the experiment started. Households were exposed to movies in random orders on the TV screen under the new menu and were randomly assigned the original price or discounted price.

Figure 5. Illustration of the Experimental Design

were shown under the “Good Deals” menu on the TV screen. This is represented by step 4 in Figure 5. Therefore, each household in our experiment was shown a different (randomized) subset of movies in a randomized order. The set of 15 movies allocated to each household remained in the “Good Deals” menu for the entire month until this procedure was repeated in the next month with a different set of 45 movies. That is, the recommender system that we study in this paper does not shift nor replaces movies once consumers buy them. If a consumer buys a movie from the “Good

Deals” menu this movie remains in the menu until the end of the month. This is the common practice in VoD systems. However, a recommender system that would dynamically adjust items to refresh those shown to consumers in salient slots would likely register even more sales for the movies placed in these slots, yielding an even larger difference in the price elasticity of demand between salient and non-salient slots. Finally, the movies under the “Good Deals” menu were shown, for each household, either at their original price or at discounted prices (step 4 in Figure 5).

Empirical Strategy and Identification ■

We want to estimate the demand curve for movies as a function of their price and the slot where they are placed on the TV screen for both standard and premium households following the functional form used in the equation below. Let i represent a movie, let $m \in \{s, p\}$ represent the type of household (s for standard and p for premium), let $l \in \{d, o\}$ represent the type of price for the movies offered under the “Good Deals” menu (d for discounted and o for original) and let t represent a week. Let d_{imit} represent the demand for movie i from households of type m with prices of type l in week t , let p_{imit} represent the price of movie i in week t for households of type m with prices of type l , and let s_{imit} represent the average slot on the TV screen in which movie i is placed in week t across households of type m with prices of type l . We estimate the following reduced form equation:

$$\begin{aligned} \log(d_{imit}) = & \beta + \beta_{0p} \log(p_{imit}) + \beta_{1p} \log(p_{imit}) \mathbf{1}(s_{imit} < R) \\ & + \beta_{0s} \log(p_{imit}) \mathbf{1}(m = s) + \beta_{1s} \log(p_{imit}) \mathbf{1}(s_{imit} < R) \mathbf{1}(m = s) \\ & + i + m + t + \tau_{im} + \varepsilon_{imit} \end{aligned}$$

where i , m , and t represent dummies for movies, household types and weeks, respectively, τ_{im} are dummies to control for the different sampling probabilities of movies across different types of households and ε_{imit} is the idiosyncratic error term. There are standard and premium households in our setting. Therefore, we control for household type. We use 270 movies during our experiment and thus the need for movie fixed effects. These fixed effects capture the effect of the different characteristics that these movies have (e.g., different genre, year of release). Given the layered structure in our setting (i.e., movies are shown to different household types), we also include an interaction fixed effect (τ in our model) to control for the fact that there is a different number of standard and premium households in our experiment from where different movies may end up having different exposure. Finally, and given that our experiment developed over several weeks, we also include week dummies to control for common time shocks.

For an illustration of our setup, refer to Figure 6. Slots before R , which are farther to the left of the TV screen, represent, on aggregate, the salient slot in our theoretical model (that is, when s_j is one in equation 1). In other words, slot k on the TV screen is salient if $k < R$ and non-salient otherwise (in this definition, slot 1 is the one farther to the left on the TV screen). Therefore, in the reduced form above, β_{0p} and $\beta_{0p} + \beta_{0s}$ represent the price elasticity of demand of premium and standard consumers with respect to movies placed in the non-salient slots. Similarly, $\beta_{1p} + \beta_{1p}$ and $\beta_{1p} + \beta_{1p} + \beta_{0s} + \beta_{1s}$ represent the price elasticity of demand of premium and

standard consumers with respect to movies placed in the salient slots. We expect $\beta_{0p} < 0$, $\beta_{0s} + \beta_{0p} < 0$, $\beta_{1p} + \beta_{1p} < 0$, and $\beta_{0p} + \beta_{1p} + \beta_0 + \beta_{0s} < 0$ given that these coefficients measure price elasticities of demand. Our research hypotheses are that $\beta_{1p} > 0$ and $\beta_{1p} + \beta_{1s} > 0$, that is, households are less price elastic with respect to movies placed farther to the left of the TV screen (which, in our setting, are more salient). Finally, we may also expect standard consumers to be more price elastic than premium consumers, that is, $\beta_{0s} < 0$ and $\beta_{0s} + \beta_{1s} < 0$, that they are subscribing to a less expensive service.

We estimate our reduced form equation using a Poisson regression because our dependent variable is a count and thus this regression provides direct and unbiased estimates for the price elasticities of demand that we are interested in. Both the price and the slot of each movie have been randomly assigned during our experiment. Thus, we expect no correlation between our independent covariates and the error term.

Results

Descriptive Statistics and Balance Across TV Slots

We analyze 45 movies per month in our experiment and we aggregate our data on a weekly basis. Our experimental dataset comprises 26 weeks, two types of households, standard and premium, and two levels of prices, discounted and original. Therefore, we analyze, 4,680 observations. We register 2,238 sales from the “Good Deals” menu during our experiment for a total revenue of \$7,663. Table 1 shows that prices, age, IMDb ratings and IMDb votes are similar across movies placed in different slots on the TV screen during our experiment. These results are obtained using ANOVA showing that the slots in which movies are placed do not explain these covariates. This is expected, given the random assignment of movies to TV slots under the “Good Deals” menu during our experiment. This also provides evidence of the good balance that our randomized schedule achieved to identify the price elasticity of demand across TV slots. Appendix E shows Tukey plots comparing all pairs of TV slots. This appendix provides additional evidence that the movies placed in different slots were similar during our experiment. Furthermore, the movies placed in the “Good Deals” menu are similar to the movies placed under other menus of TELCO’s VoD system (results available upon request) and thus the sample of movies that we use in our experiment represents well the set of movies uses under this highlights section. Appendix F provides evidence that the allocation of movies to households was random as well.

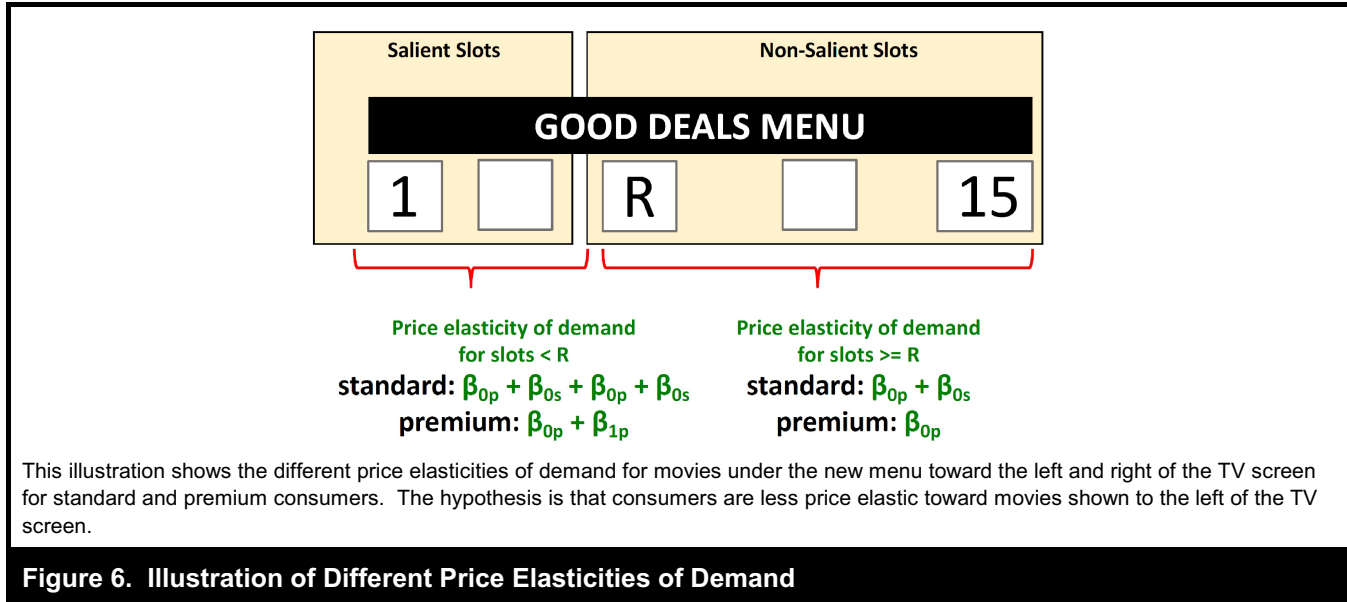


Table 1. Evidence of Balance across Slots for Movie Prices, Ages, IMDb Ratings, and IMDb Votes

Covariate for ANOVA	Df	Sum sq	Mean sq	F-value	p-value
Price					
Slot	7	44722	6389	0.884	0.519
Residuals	413	2986451	7231		
IMDb Rating					
Slot	7	4.18	0.597	0.434	0.88
Residuals	192	263.97	1.375		
IMDb Votes					
Slot	7	1.153e+11	1.647e+10	1.015	0.422
Residuals	192	3.116e+12	1.623e+10		
Movie Age					
Slot	7	251	35.8	0.77	0.613
Residuals	182	8464	46.5		

Table 2 shows the average price per slot for the movies placed in the “Good Deals” menu during our experiment. This table also shows the average price for when the price was low and high, which allows us to confirm the average discount of 25% during our experiment. This table also shows sales per slot, and cumulative sales up to each slot, for the movies placed in this menu during our experiment. We observe that movies with the higher price sell more up to slot 4, which gives us some preliminary evidence that despite the higher price these movies still sell more when placed toward the left of the TV screen. Finally, this table also shows the average IMDb rating, average IMDb votes, and average per slot for these movies. The average for these covariates across all movies placed in the “Good Deals” menu during our experiment are 6.25, 79,869 and 2008.14, respectively.

Empirical Findings

Table 3 shows the results obtained from estimating the reduced form equation presented earlier using the data collected during our randomized experiment. Each column instantiates a different value for *R* as indicated on the top of the corresponding column. The price elasticities of demand for movies in slots *R* and beyond are -0.29 and -0.75 for premium and standard consumers, respectively. In addition, standard consumers are more elastic than premium consumers (*p* – values < 0.0004). More important, we confirm our hypothesis that consumers, both standard and premium, are less price-elastic toward movies placed in slots farther to the left of the TV screen, that is, slots prior to *R*, as shown by the interaction between price elasticity and *slot* < *R* for the cases

Table 2. Prices (In Cents), Sales and Movie Characteristics per Slot for the Movies under the “Good Deals” Menu During Our Experiment

Covariate	Stat	Slot:	1	2	3	4	5	6	7	8
Price	Mean		290.7	302.8	281.5	275.8	285.5	291.0	291.4	302.2
Price price=low	Mean		246.2	252.1	245.3	236.8	237.7	249.6	241.9	253.8
Price price=high	Mean		332.0	340.4	316.3	314.8	325.6	329.2	336.4	344.3
Sales price=low	Sum		89	94	71	76	86	98	101	69
Cumulative sales price=low	Sum		89	183	254	330	416	514	615	684
Sales price=high	Sum		109	86	62	80	72	53	65	62
Cumulative sales price=high	Sum		109	195	257	337	409	462	527	589
IMDb rating	Mean		6.34	6.48	6.49	6.07	6.21	6.03b	6.30	6.37
IMDb votes	Mean		69485	65545	89221	67213	62480	68613	101790	96039

of $R = 2$ and $R = 3$. One could also expect this interaction to become less significant as R increases, that is, as one blurs the definition of left of the TV screen. Empirically, we find that both premium and standard consumers exhibit a smaller price elasticity of demand for the first two slots under the new menu.

Table 4 shows how the estimates for the price elasticity of demand for premium and standard consumers differ for time periods right after the introduction of the “Good Deals” menu compared to later time periods. As in the previous table, each column instantiates a different value of R as indicated at the top of the column. We create a dummy variable that indicates whether observations pertain to the first 4 months of the experiment, represented in the table by $period < P$, or to the last 3 months of the experiment. Therefore, the results in this table show that the difference in the price elasticity of demand between the salient and non-salient slots arises during the time periods immediately after the introduction of this menu, thus providing some evidence that parameter β_1 in the theoretical model introduced earlier attenuates over time in our setting, as could be expected in light of the discussion in that section. These results are obtained controlling for movie age to avoid potential confounding effects as time goes by.

The Role of Movie Popularity, Movie Quality and Movie Age

The previous results show that consumers exhibit a lower price elasticity of demand for movies placed toward the left of the TV screen. This is a result obtained on average across all the movies used during our experiment. Table 5 shows that this result is heterogeneous across movie popularity. We define a dummy variable indicating whether the number of votes on IMDb falls within the first three quartiles of the distribution of IMDb votes across the movies used in our

experiment or else on the top quartile of this distribution. We observe that the difference in the price elasticity of demand for movies placed toward the left of the TV screen is driven by popular movies both for standard and premium consumers. This means that while on average the firm is able to hook up consumers to movies placed toward the left of the TV screen this is not necessarily true for all movies. In fact, this does not seem to be the case for unpopular movies. The firm is only able to hook up consumers to movies that are sufficiently well known, otherwise placing unpopular movies toward the left of the TV screen does not seem to affect the price elasticity of demand. In other words, the firm is only able to profitably increase the price of movies placed toward the left of the TV screen if these movies are sufficiently well known.

Table 6 provides similar heterogeneous results for the case of movie ratings. We define a dummy variable indicating whether IMDb rating falls within the first three quartiles of the distribution of IMDb ratings across the movies used in our experiment or else on the top quartile of this distribution. We observe that the difference in the price elasticity of demand for movies placed toward the left of the TV screen is also driven by better movies for both standard and premium consumers. Similar to the case of movie popularity, placing worse movies toward the left of the TV screen does not seem enough to hook up consumers. The firm is only able to profitably increase the price of the movies placed toward the left of the TV screen if they are good enough.

Finally, Table 7 provides similar heterogeneous results for the case of movie age. In this analysis, movie age is measured in years elapsed since the release of the movie in TELCO’s VoD system and the beginning of our experiment. We observe that the difference in the price elasticity of demand for movies placed toward the left of the TV screen is also driven by more recent movies both for standard and premium consumers. Similar to the cases of movie popularity and movie quality,

Table 3. Estimates for the Price Elasticity of Demand for Premium and Standard Consumers for Movies Placed in Slots Prior to R and in Slot R and Beyond (consumers are less elastic to movies placed in slots toward the left of the TV screen)

	Dependent Variable: Sales							
	R = 2	R = 3	R = 4	R = 5	R = 6	R = 7	R = 8	R = 9
log(price)	-0.298 ^{**}	-0.303 ^{**}	-0.301 ^{**}	-0.301 ^{**}	-0.286 ^{**}	-0.288 ^{**}	-0.293 ^{**}	-0.312 ^{**}
	(0.135)	(0.135)	(0.135)	(0.135)	(0.135)	(0.135)	(0.136)	(0.137)
log(price)*slot < R	0.122 ^{***}	0.074 ^{**}	0.052	0.025	-0.004	0.0004	0.007	0.028
	(0.042)	(0.033)	(0.032)	(0.027)	(0.025)	(0.024)	(0.026)	(0.032)
log(price)*standard	-0.463 ^{***}	-0.461 ^{***}	-0.461 ^{***}	-0.456 ^{***}	-0.460 ^{***}	-0.457 ^{***}	-0.454 ^{***}	-0.444 ^{***}
	(0.130)	(0.131)	(0.131)	(0.131)	(0.131)	(0.131)	(0.131)	(0.131)
log(price)*slot < R*standard	-0.023	0.001	-0.0004	-0.012	-0.002	-0.008	-0.013	-0.019
	(0.030)	(0.022)	(0.020)	(0.019)	(0.017)	(0.017)	(0.017)	(0.017)
Constant	0.093	0.224	0.302	0.256	0.269	0.269	0.272	0.271
	(1.042)	(1.037)	(1.038)	(1.037)	(1.038)	(1.038)	(1.038)	(1.037)
Sampling Probability Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Movie Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Household Type Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,680	4,680	4,680	4,680	4,680	4,680	4,680	4,680
Log Likelihood	-3,417.726	-3,419.441	-3,420.663	-3,421.502	-3,422.024	-3,421.923	-3,421.750	-3,421.186
Akaike Inf. Crit.	7,431.451	7,434.881	7,437.326	7,439.003	7,440.049	7,439.846	7,439.499	7,438.373

Note: *p < 0.1; **p < 0.05; ***p < 0.01

Table 4. Estimates for the Price Elasticity of Demand for Premium and Standard Consumers for Movies Placed in Slots Prior to R and in Slot R and Beyond for Months 1–4 ($period < P$) Versus Months 5–7 (the difference in the price elasticity of demand for slots prior to R arises from time periods right after the introduction of the “Good Deals” menu)

	Dependent Variable: Sales							
	R = 2	R = 3	R = 4	R = 5	R = 6	R = 7	R = 8	R = 9
log(price)	-0.285 ^{**}	-0.298 ^{**}	-0.303 ^{**}	-0.318 ^{**}	-0.330 ^{**}	-0.312 ^{**}	-0.332 ^{**}	-0.340 ^{**}
	(0.136)	(0.136)	(0.136)	(0.137)	(0.137)	(0.137)	(0.138)	(0.139)
log(price)*standard	-0.434 ^{***}	-0.427 ^{***}	-0.416 ^{***}	-0.400 ^{***}	-0.388 ^{***}	-0.372 ^{***}	-0.369 ^{***}	-0.356 ^{***}
	(0.132)	(0.132)	(0.132)	(0.132)	(0.132)	(0.133)	(0.133)	(0.134)
log(price)*slot < R	0.073	-0.040	-0.028	-0.035	-0.087	-0.042	-0.012	0.003
	(0.106)	(0.082)	(0.068)	(0.068)	(0.060)	(0.053)	(0.054)	(0.068)
log(price)*slot < R*standard	-0.154	-0.107 [*]	-0.153 ^{**}	-0.124 ^{***}	-0.115 ^{***}	-0.105 ^{***}	-0.100 ^{***}	-0.114 ^{***}
	(0.099)	(0.063)	(0.060)	(0.046)	(0.042)	(0.036)	(0.035)	(0.034)
log(price)*slot < R*period < P	0.112	0.174 [*]	0.136 [*]	0.141 [*]	0.194 ^{**}	0.086	0.085	0.071
	(0.112)	(0.097)	(0.082)	(0.079)	(0.077)	(0.062)	(0.063)	(0.068)
log(price)*slot < R*period < P*standard	0.146	0.128 [*]	0.177 ^{***}	0.133 ^{***}	0.135 ^{***}	0.117 ^{***}	0.105 ^{***}	0.115 ^{***}
	(0.104)	(0.066)	(0.062)	(0.048)	(0.043)	(0.037)	(0.035)	(0.034)
Constant	0.684	0.729	0.743	0.772	0.661	0.566	0.241	0.325
	(1.489)	(1.487)	(1.487)	(1.486)	(1.490)	(1.510)	(1.526)	(1.530)
Sampling Probability Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Movie Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Household Type Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Movie Age Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,312	4,312	4,312	4,312	4,312	4,312	4,312	4,312
Log Likelihood	-3,320.240	-3,319.344	-3,317.057	-3,317.265	-3,314.865	-3,317.947	-3,318.006	-3,316.294
Akaike Inf. Crit.	7,246.480	7,244.688	7,240.114	7,240.531	7,235.729	7,241.895	7,242.012	7,238.589

Note: *p < 0.1; **p < 0.05; ***p < 0.01

Table 5. The Role of Popularity on the Estimates for the Price Elasticity of Demand for Premium and Standard Consumers for Movies Placed in Slots Prior to R and in Slot R and Beyond (consumers are less elastic to popular movies placed in slots toward the left of the TV screen)

	<i>Dependent Variable: Sales</i>							
	R = 2	R = 3	R = 4	R = 5	R = 6	R = 7	R = 8	R = 9
log(price)	0.078 (0.238)	0.068 (0.239)	0.052 (0.240)	0.058 (0.240)	0.072 (0.239)	0.078 (0.239)	0.074 (0.240)	0.052 (0.240)
log(price)*bottom quartiles	-0.527** (0.258)	-0.514** (0.258)	-0.494* (0.258)	-0.496* (0.259)	-0.498* (0.258)	-0.507** (0.258)	-0.505* (0.259)	-0.504* (0.259)
log(price)*slot < R	0.249*** (0.073)	0.182*** (0.063)	0.185*** (0.062)	0.142*** (0.055)	0.105** (0.051)	0.078 (0.048)	0.078 (0.048)	0.090 (0.055)
log(price)*slot < R*bottom quartiles	-0.443*** (0.133)	-0.454*** (0.133)	-0.456*** (0.133)	-0.450*** (0.132)	-0.450*** (0.132)	-0.444*** (0.133)	-0.437*** (0.133)	-0.429*** (0.135)
log(price)*standard	-0.175** (0.083)	-0.146** (0.070)	-0.181*** (0.070)	-0.152** (0.061)	-0.139** (0.056)	-0.100* (0.055)	-0.094* (0.055)	-0.077 (0.058)
log(price)*bottom quartiles*standard	-0.003 (0.020)	-0.001 (0.022)	-0.001 (0.022)	0.001 (0.023)	-0.007 (0.024)	-0.008 (0.025)	-0.011 (0.027)	-0.006 (0.029)
log(price)*slot < R*standard	-0.100 (0.089)	0.0003 (0.049)	0.004 (0.043)	-0.006 (0.041)	-0.017 (0.038)	-0.026 (0.037)	-0.030 (0.035)	-0.026 (0.035)
log(price)*slot < R*bottom quartiles*standard	0.087 (0.095)	0.011 (0.056)	0.004 (0.048)	-0.002 (0.046)	0.023 (0.043)	0.026 (0.042)	0.025 (0.040)	0.012 (0.040)
Constant	0.816 (1.095)	0.864 (1.093)	0.856 (1.094)	0.716 (1.096)	0.758 (1.096)	0.826 (1.096)	0.845 (1.097)	0.835 (1.100)
Sampling Probability Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Movie Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Household Type Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,592	4,592	4,592	4,592	4,592	4,592	4,592	4,592
Log Likelihood	-3,352.648	-3,354.390	-3,354.443	-3,355.648	-3,356.218	-3,357.516	-3,357.563	-3,357.579
Akaike Inf. Crit.	7,295.296	7,298.779	7,298.885	7,301.295	7,302.435	7,305.033	7,305.126	7,305.158

Note: *p < 0.1; **p < 0.05; ***p < 0.01

Table 6. The Role of Quality on the Estimates for the Price Elasticity of Demand for Premium and Standard Consumers for Movies Placed in Slots Prior to R and in Slot R and Beyond (consumers are less elastic to better movies placed in slots toward the left of the TV screen)

	Dependent Variable: Sales							
	R = 2	R = 3	R = 4	R = 5	R = 6	R = 7	R = 8	R = 9
log(price)	-0.154 (0.235)	-0.158 (0.235)	-0.165 (0.236)	-0.164 (0.238)	-0.147 (0.237)	-0.149 (0.237)	-0.144 (0.238)	-0.187 (0.240)
log(price)*bottom quartiles	-0.219 (0.257)	-0.218 (0.257)	-0.209 (0.257)	-0.207 (0.259)	-0.211 (0.258)	-0.209 (0.258)	-0.223 (0.258)	-0.192 (0.260)
log(price)*slot < R	0.244*** (0.072)	0.137** (0.058)	0.138** (0.057)	0.053 (0.045)	0.038 (0.042)	0.036 (0.039)	0.027 (0.041)	0.074 (0.048)
log(price)*slot < R*bottom quartiles	-0.472*** (0.134)	-0.476*** (0.135)	-0.475*** (0.134)	-0.463*** (0.134)	-0.462*** (0.134)	-0.456*** (0.135)	-0.437*** (0.136)	-0.400*** (0.137)
log(price)*standard	-0.175** (0.083)	-0.095 (0.067)	-0.128* (0.068)	-0.052 (0.059)	-0.066 (0.051)	-0.059 (0.048)	-0.033 (0.049)	-0.067 (0.053)
log(price)*bottom quartiles*standard	0.006 (0.022)	0.016 (0.023)	0.017 (0.024)	0.009 (0.025)	-0.002 (0.026)	-0.001 (0.027)	-0.014 (0.030)	-0.041 (0.032)
log(price)*slot < R*standard	-0.053 (0.053)	0.014 (0.040)	0.015 (0.039)	-0.014 (0.038)	-0.028 (0.036)	-0.031 (0.036)	-0.049 (0.036)	-0.083** (0.037)
log(price)*slot < R*bottom quartiles*standard	0.048 (0.066)	-0.004 (0.049)	-0.009 (0.046)	0.010 (0.043)	0.039 (0.042)	0.033 (0.041)	0.050 (0.041)	0.085** (0.042)
Constant	0.415 (1.101)	0.531 (1.098)	0.539 (1.100)	0.531 (1.107)	0.522 (1.107)	0.541 (1.105)	0.637 (1.107)	0.587 (1.104)
Sampling Probability Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Movie Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Household Type Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,592	4,592	4,592	4,592	4,592	4,592	4,592	4,592
Log Likelihood	-3,354.225	-3,357.022	-3,357.399	-3,360.004	-3,359.639	-3,359.785	-3,359.743	-3,357.664
Akaike Inf. Crit.	7,298.450	7,304.044	7,304.797	7,310.008	7,309.277	7,309.570	7,309.485	7,305.329

Note: *p < 0.1; **p < 0.05; ***p < 0.01

Table 7. The Role of Movie Age on the Estimates for the Price Elasticity of Demand for Premium and Standard Consumers for Movies Placed in Slots Prior to R and in Slot R and Beyond (consumers are less elastic to more recent movies placed in slots toward the left of the TV screen)

	Dependent Variable: Sales							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log(price)	-0.747*** (0.235)	-0.721*** (0.235)	-0.676*** (0.238)	-0.750*** (0.236)	-0.709*** (0.237)	-0.754*** (0.235)	-0.766*** (0.236)	-0.794*** (0.237)
log(price)*movieage	0.059** (0.026)	0.056** (0.026)	0.048* (0.027)	0.061** (0.027)	0.054** (0.027)	0.061** (0.026)	0.063** (0.027)	0.065** (0.027)
log(price)*slot < R	0.072 (0.093)	0.005 (0.071)	-0.071 (0.078)	0.019 (0.069)	-0.062 (0.062)	-0.004 (0.060)	0.015 (0.061)	0.052 (0.070)
log(price)*slot < R*movieage	-0.413*** (0.143)	-0.438*** (0.144)	-0.433*** (0.144)	-0.433*** (0.143)	-0.419*** (0.143)	-0.414*** (0.142)	-0.413*** (0.142)	-0.408*** (0.142)
log(price)*standard	0.009 (0.014)	0.012 (0.012)	0.023* (0.014)	0.001 (0.013)	0.011 (0.011)	0.001 (0.009)	-0.002 (0.009)	-0.004 (0.010)
log(price)*standard*movieage	0.002 (0.002)	0.0003 (0.003)	0.001 (0.003)	0.0002 (0.003)	0.002 (0.003)	0.002 (0.003)	0.002 (0.003)	0.003 (0.003)
log(price)*slot < R*standard	-0.014 (0.039)	0.001 (0.031)	0.003 (0.029)	-0.014 (0.028)	0.010 (0.027)	0.006 (0.027)	0.002 (0.027)	0.0002 (0.028)
log(price)*slot < R*standard*movieage	-0.002 (0.003)	0.0003 (0.003)	-0.0001 (0.003)	0.001 (0.003)	-0.002 (0.004)	-0.002 (0.004)	-0.002 (0.004)	-0.003 (0.004)
Constant	-0.386 (1.079)	-0.263 (1.076)	-0.083 (1.081)	-0.282 (1.087)	-0.108 (1.087)	-0.232 (1.076)	-0.254 (1.080)	-0.290 (1.081)
Sampling Probability Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Movie Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Household Type Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,312	4,312	4,312	4,312	4,312	4,312	4,312	4,312
Log Likelihood	-3,330.722	-3,332.283	-3,332.615	-3,334.990	-3,334.821	-3,335.089	-3,334.906	-3,334.326
Akaike Inf. Crit.	7,219.444	7,222.565	7,223.230	7,227.980	7,227.642	7,228.178	7,227.811	7,226.651

Note: *p < 0.1; **p < 0.05; ***p < 0.01

Table 8. Absence of Effect of Slot on the Number of Likes Received, Indicating That During Our Experiment Households Did Not Seem to Change Their Preferences for Media

	Dependent Variable: Number of Likes							
	R = 2	R = 3	R = 4	R = 5	R = 6	R = 7	R = 8	R = 9
log(price)	-0.327 (0.287)	-0.325 (0.287)	-0.325 (0.287)	-0.321 (0.287)	-0.317 (0.287)	-0.322 (0.287)	-0.317 (0.287)	-0.318 (0.287)
slot < R	0.507 (0.618)	-0.092 (0.460)	-0.281 (0.417)	-0.232 (0.359)	-0.486 (0.343)	-0.415 (0.316)	-0.596* (0.342)	-0.527 (0.411)
Constant	-17.722 (3,257.668)	-17.547 (3,256.901)	-17.540 (3,256.813)	-17.416 (3,259.311)	-17.328 (3,262.483)	-17.415 (3,259.296)	-17.418 (3,261.095)	-17.271 (3,262.955)
Sampling Probability Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Movie Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,340	2,340	2,340	2,340	2,340	2,340	2,340	2,340
Log Likelihood	-928.518	-928.837	-928.629	-928.647	-927.820	-927.974	-927.270	-928.008
Akaike Inf. Crit.	2,447.037	2,447.675	2,447.258	2,447.294	2,445.640	2,445.947	2,444.541	2,446.016

Note: *p < 0.1; **p < 0.05; ***p < 0.01

placing older movies toward the left of the TV screen does not seem enough to hook up consumers. The firm is only able to profitably increase the price of the movies placed toward the left of the TV screen if these movies are sufficiently new.

Effect on Consumer Tastes

Tables 8 and 9 show the results obtained from estimating the reduced form equation presented earlier only for premium consumers, using the likes issued by households during the experiment. These results are shown only for premium households because standard ones did not see the number of likes on the TV and could not issue likes. The dependent variables in these regressions are the number of likes and the number of likes per lease, respectively (the latter is set to zero when there are no leases). The results in both these tables show that our exogenous manipulations of movies into slots during the experiment did not change the number of likes issued by consumers. This provides some evidence that doing so did not seem to change their preferences, which lessens the concerns that potential losses in consumer surplus could be confounded with the fact that preferences may change as a result of recommendations. Tables 10 and 11 show similar results interacting slot with movie quality, providing additional evidence that consumer preferences did not seem to change during our experiment even when better movies were placed in salient slots.

Effect of Recommender Systems on Consumer Surplus and Welfare

Setup for Simulations

The previous subsections show that, in our setting, both premium and standard consumers are less price elastic for movies placed in the slots toward the left of the TV screen under the new menu. According to the theoretical model presented earlier this is likely to lead TELCO to place movies in these slots in a way that hurts consumer surplus and that may also hurt total welfare. As customary, we resort to policy simulations to understand the potential magnitude of these losses. Our randomized experiment allowed us to empirically identify the relevant parameters of the demand curve. Therefore, we can use them now to determine the appropriate counterfactuals, that is, how would consumers behave if they faced different sets of movies on the TV screen that were suggested to them by different recommender systems. We use the empirical results obtained in Table 3 with $R = 3$ (which yields higher log-likelihood than those obtained with $R = 2$) to run these simulations. The results for $R = 2$ are similar to the ones presented below and available upon request.

Consider a subset of 15 movies out of all the movies used during our experiment. In addition, assume that once the experiment is over TELCO sells these movies to consumers in a new menu in its VoD system at the prices originally negotiated with content providers. TELCO's goal is to place these movies into the 15 slots under the new menu to maximize profit. These movies may be ordered differently from left to right on the TV screen for standard and premium consumers because these consumers have different preferences. Therefore, we compute separately for standard and premium consumers the order of recommendations that maximizes TELCO's profits. For the allocation of movies to slots that maximizes profit we compute total profits, consumer surplus and total welfare. Likewise, for the order of movies that maximizes consumer surplus and welfare. We compute consumer surplus, profit, and welfare for each type of recommender system and compare them. Finally, note that there are more than 8×10^{23} ways to choose a set of 15 movies out of the movies that we used throughout the experiment. We choose 1,000 of these sets at random and average out our results. Our findings do not change if, instead, we choose 10,000 of these sets at random. The results for the latter case are available upon request.

Comparing Different Recommender Systems

Figure 7(a) shows the relationship between consumer surplus when consumer surplus is maximized (y-axis) and consumer surplus when profit is maximized (x-axis). Each dot in this figure is a simulation for a random set of 15 movies. As expected, all dots lie (weakly) above the 45-degree line and there is some variance in terms of their distance to this line, which measures the magnitude of the loss in consumer surplus introduced by the fact that the firm maximizes profit. On average, across our simulations, this loss is 8.44%. Figure 7(b) shows the relationship between profit when profit is maximized (x-axis) and profit when consumer surplus is maximized (y-axis). This figure shows how much profit the firm loses when consumer surplus is maximized. The average loss across our simulations is 7.17%.

Figure 8(a) shows the relationship between total welfare when profit is maximized (x-axis) and total welfare when total welfare is maximized (y-axis). On average across our simulations, this loss is 2.07%. Figure 8(b) shows the relationship between total welfare when consumer surplus is maximized (x-axis) and total welfare when total welfare is maximized (y-axis). On average across our simulations, this loss is 1.62%. Therefore, maximizing consumer surplus does not hurt total welfare as much as maximizing profit (p -value < 0.0001). The figures in Appendix D show the distribution of losses for the four cases presented before. We observe that some sets of

Table 9. Absence of Effect of Slot on the Number of Likes Received per Lease, Indicating That During Our Experiment Households Did Not Seem to Change Their Preferences for Media

	Dependent Variable: Number of Likes/Number of Leases							
	R = 2	R = 3	R = 4	R = 5	R = 6	R = 7	R = 8	R = 9
log(price)	-0.040 (0.034)	-0.040 (0.034)	-0.040 (0.034)	-0.040 (0.034)	-0.040 (0.034)	-0.040 (0.034)	-0.039 (0.034)	-0.039 (0.034)
slot < R	0.045 (0.055)	0.016 (0.041)	0.016 (0.041)	0.014 (0.037)	-0.004 (0.033)	-0.006 (0.030)	-0.034 (0.034)	-0.017 (0.044)
Constant	0.131 (0.210)	0.140 (0.210)	0.143 (0.209)	0.141 (0.210)	0.144 (0.210)	0.145 (0.210)	0.155 (0.210)	0.154 (0.211)
Sampling Probability Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Movie Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,340	2,340	2,340	2,340	2,340	2,340	2,340	2,340
R ²	0.164	0.164	0.164	0.164	0.164	0.164	0.164	0.164
Adjusted R ²	0.044	0.043	0.043	0.043	0.043	0.043	0.044	0.043
Residual Std. Error (df = 2045)	0.278	0.278	0.278	0.278	0.278	0.278	0.278	0.278
F Statistic (df = 294; 2045)	1.364***	1.362***	1.361***	1.361***	1.361***	1.361***	1.365***	1.362***

Note: *p < 0.1; **p < 0.05; ***p < 0.01

Table 10. Absence of Effect of Slot on the Number of Likes Received as a Function of Movie Quality Measured by IMDb Ratings), Indicating That During Our Experiment Households Did Not Seem to Change Their Preferences for Media

	Dependent Variable: Number of Likes							
	R = 2	R = 3	R = 4	R = 5	R = 6	R = 7	R = 8	R = 9
log(price)	-0.374 (0.289)	-0.372 (0.289)	-0.372 (0.289)	-0.364 (0.289)	-0.364 (0.289)	-0.376 (0.289)	-0.364 (0.289)	-0.368 (0.289)
slot < R	0.527 (1.136)	-0.453 (0.741)	-0.457 (0.738)	-0.568 (0.610)	-0.514 (0.544)	-0.138 (0.501)	-0.560 (0.557)	-0.376 (0.659)
slot < R*bottom quantiles	-0.027 (1.328)	0.573 (0.924)	0.257 (0.891)	0.534 (0.770)	0.045 (0.642)	-0.435 (0.621)	-0.055 (0.663)	-0.216 (0.745)
Constant	-17.514 (3,253.859)	-17.351 (3,251.883)	-17.312 (3,252.327)	-17.045 (3,256.255)	-17.102 (3,259.071)	-17.416 (3,254.947)	-17.244 (3,257.351)	-17.149 (3,259.117)
Sampling Probability Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Movie Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,296	2,296	2,296	2,296	2,296	2,296	2,296	2,296
Log Likelihood	-909.506	-909.632	-909.574	-909.396	-908.815	-908.718	-908.255	-908.955
Akaike Inf. Crit.	2,397.013	2,397.264	2,397.148	2,396.793	2,395.629	2,395.435	2,394.511	2,395.909

Note: *p < 0.1; **p < 0.05; ***p < 0.01

Table 11. Absence of Effect of Slot on the Number of Likes Received per Lease as a Function of Movie Quality Measured by IMDb Ratings), Indicating That During Our Experiment Households Did Not Seem to Change Their Preferences for Media.

	Dependent Variable: Number of Likes/Number of Leases							
	R = 2	R = 3	R = 4	R = 5	R = 6	R = 7	R = 8	R = 9
log(price)	-0.044 (0.034)	-0.043 (0.034)	-0.043 (0.034)	-0.043 (0.034)	-0.044 (0.034)	-0.044 (0.034)	-0.043 (0.034)	-0.044 (0.034)
slot < R	0.156 (0.123)	0.010 (0.081)	0.008 (0.081)	0.0004 (0.061)	0.030 (0.055)	0.026 (0.049)	-0.002 (0.052)	0.026 (0.067)
slot < R*bottom quantiles	-0.136 (0.135)	0.008 (0.091)	0.010 (0.094)	0.022 (0.080)	-0.052 (0.068)	-0.051 (0.061)	-0.050 (0.064)	-0.065 (0.074)
Constant	0.147 (0.212)	0.156 (0.212)	0.159 (0.212)	0.160 (0.212)	0.150 (0.212)	0.158 (0.212)	0.166 (0.212)	0.170 (0.213)
Sampling Probability Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Movie Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,296	2,296	2,296	2,296	2,296	2,296	2,296	2,296
R ²	0.165	0.164	0.164	0.164	0.165	0.165	0.165	0.165
Adjusted R ²	0.045	0.044	0.044	0.044	0.045	0.045	0.045	0.045
Residual Std. Error (df = 2007)	0.279	0.279	0.279	0.279	0.279	0.279	0.279	0.279
F Statistic (df = 288; 2007)	1.377***	1.370***	1.370***	1.371***	1.372***	1.373***	1.376***	1.374***

Note: *p < 0.1; **p < 0.05; ***p < 0.01

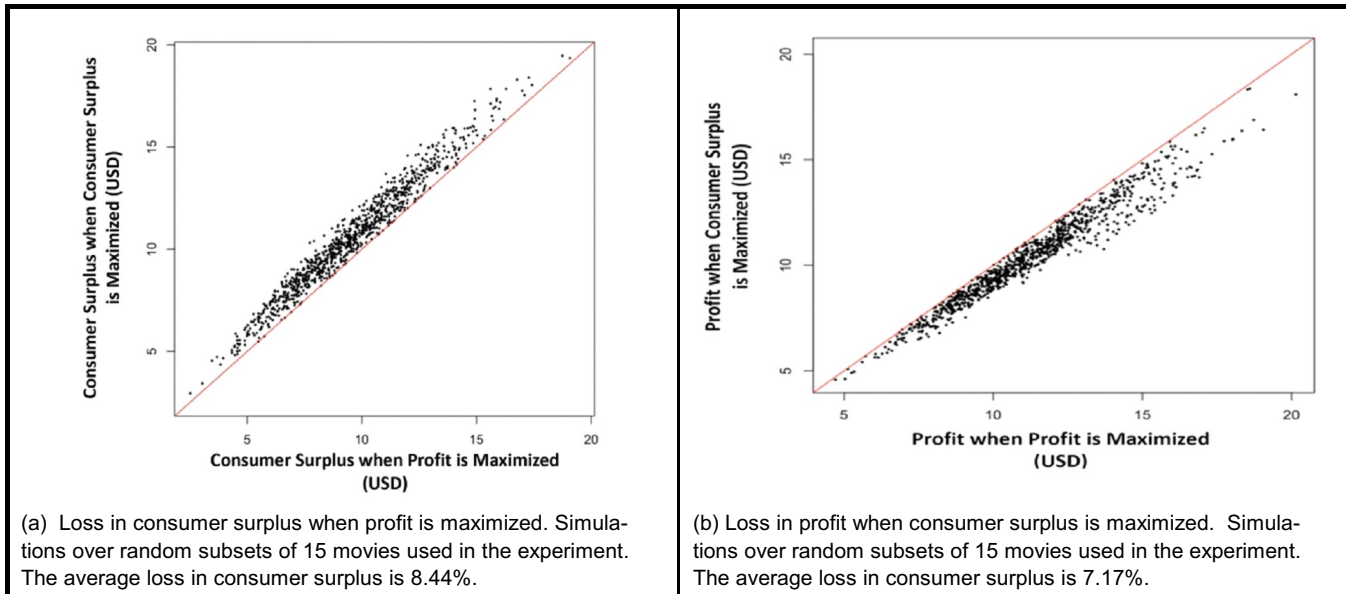
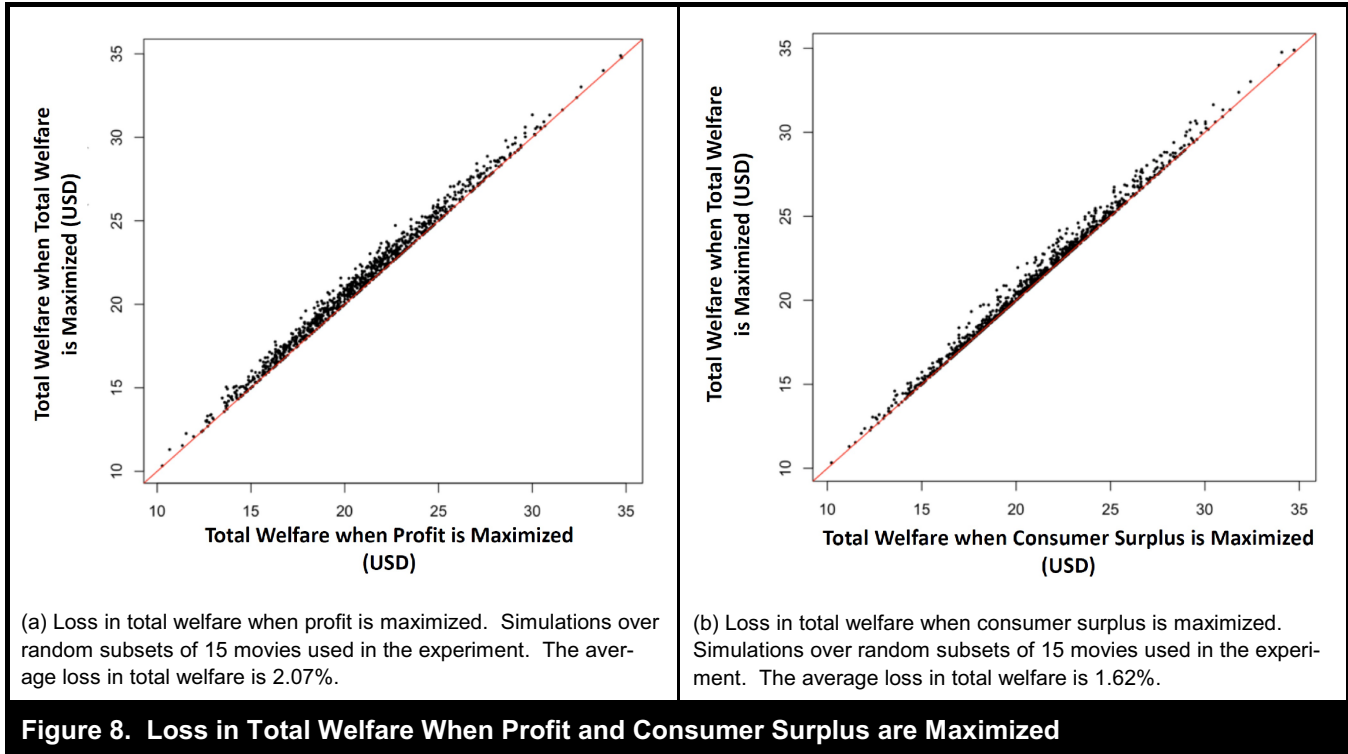


Figure 7. Changes in Consumer Surplus and Profit as a Function of Maximization Goal



movies yield significant losses in consumer surplus and welfare. The largest losses in consumer surplus and welfare across our simulations when profit is maximized are 28.7% and 12.2%, respectively.

Consumers' Trust in Recommender Systems

The previous subsection shows that losses in welfare, and particularly in consumer surplus, can arise when the firm chooses the order in which products are recommended to consumers in order to maximize profits. These losses are likely to arise when consumers exhibit different price elasticities of demand for different slots and the firm takes advantage of this knowledge to maximize profit. Still, the simulations presented in the previous subsection are for a one-shot game, in which consumers buy only once from the recommender system. Our results show that the recommender system that maximizes profits leads consumers to lose surplus. Still, consumers may also lose trust in a recommender system that induces consumer loss every time they use it. This may lead consumers to come more seldom to the market, which may result in less profit for the firm and less total welfare in the long-run.

Figure 9 compares the consumer surplus obtained from ten different recommender systems. As before, each simulation

is for a random subset of 15 movies used during the experiment and we report results across 1,000 simulations. The three recommender systems farthest to the left maximize profit, consumer surplus and welfare, respectively from left to right. The fourth recommender system is a content-based algorithm for movie genre. We identify two sets of TELCO consumers that have different tastes for movie genres as indicated by sales data during the 30 days prior our experiment. One set of consumers prefers drama movies while the other set of consumers prefer action movies. The fourth recommender system shown in this figure places movies of the genre that the consumer prefers in the salient slot (if there are fewer movies of the required genre than salient slots then the algorithm randomizes the movies placed in the extra salient slots; if, on the other hand, there are more movies of the required genre than slots then the algorithm randomizes the movies of the required genre that are shown in the salient slots). The three recommender systems on the far right of this figure place the most sold movies (in TELCO's VoD system during the experiment), the most rated movies (i.e., the most popular movies according to IMDb) and the highest rated movies (i.e., the movies with the highest rank according to IMDb) toward the left of the TV screen. These recommender systems are simple to implement and thus very often used in practice. Examples of firms that use these types of recommender systems include Amazon, Expedia, Netflix, to name a few.

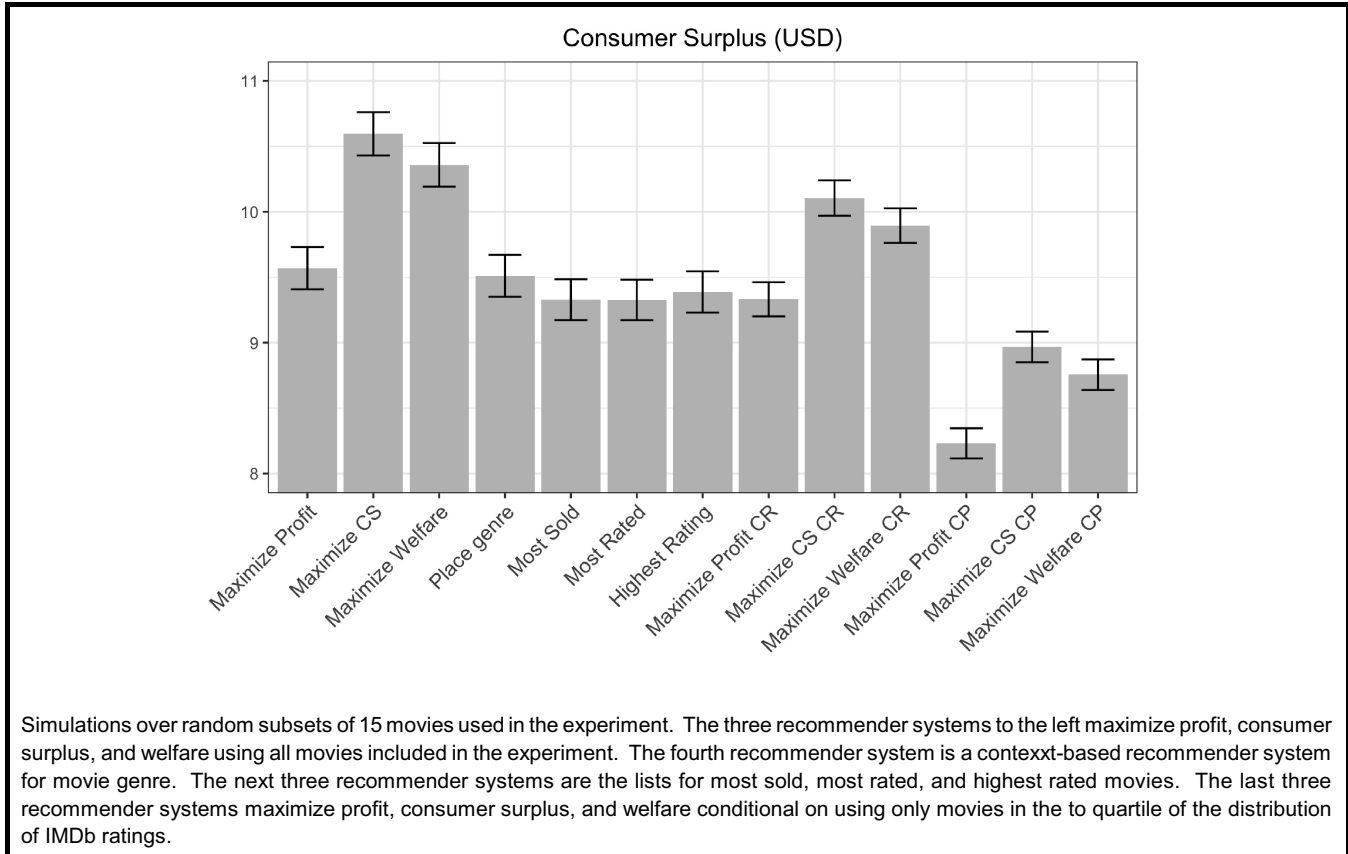
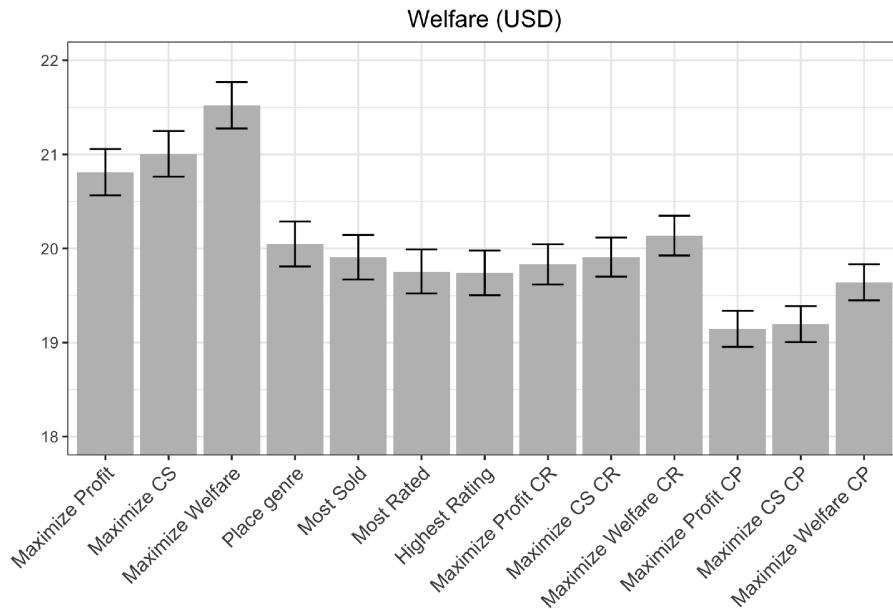


Figure 9. Comparison of the Consumer Surplus Obtained Using Different Recommender Systems

We observe that the recommender system that maximizes profit yields as much consumer surplus as the content-based recommender system. The latter also generates as much consumer surplus as the recommender systems that highlights the most sold, the most rated and the highest rated movies to consumers. Appendix G shows this result is not driven by the fact that TELCO charges consumers in the inelastic portion of the demand curve. Figure 10 presents the same type of information for the case of total welfare. In this case, the recommender system that maximizes profit yields more total welfare than the other recommender systems in the figure. Therefore, our analysis provides some empirical evidence that, at least in our context, consumers are unlikely to lose trust in a recommender system that maximizes profit every time they use it. Such a recommender system generates lower consumer surplus than a recommender system that maximizes consumer surplus, but the latter is unlikely available to consumers and the former seems to yield as much consumer surplus (as well as much total welfare) as that generated by the other recommender systems that we refer above. Therefore, consumers are unlikely to notice, or at least unlikely to complain, if the company moves from those

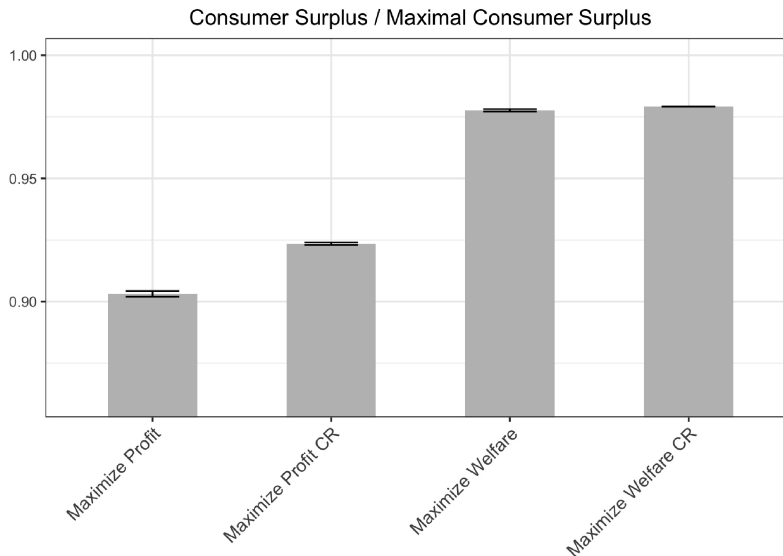
systems to a system that recommends products to maximize profit.

Yet, it may still be the case that firms may want to maintain consumer trust by recommending only good movies that would speak for the reputation of their VoD system. We model this behavior by considering recommender systems that maximize profit, consumer surplus, or welfare that can only use movies included in our experiment that belong in the top quartile of the distribution of IMDb ratings. These three recommender systems are represented toward the far right of Figures 9 and 10. As expected, such a system generates less consumer surplus and less welfare given that it operates only over a constrained set of movies. More important, maximizing profit conditional on showing only these movies to consumers reduces the loss in consumer surplus relative to the maximum consumer surplus that could be generated. Columns (1) and (2) in Figure 11 show this result. Column (1) shows that when all movies are available the recommender system that maximizes profit generates 90.3% of the maximum consumer surplus. Column (2) shows that when only movies in the top quartile of the distribution of IMDb ratings



Simulations over random subsets of 15 movies used in the experiment. The three recommender systems to the left maximize profit, consumer surplus, and welfare using all movies included in the experiment. The fourth recommender system is a content-based recommender system for movie genre. The next three recommender systems are the lists for most sold, most rated, and highest rated movies. The last three recommender systems maximize profit, consumer surplus, and welfare conditional on using only movies in the top quartile of the distribution of IMDb ratings.

Figure 10. Comparison of the Welfare Obtained Using Different Recommender Systems



Comparison in the loss of consumer surplus when the firm maximizes profit using all movies used in the experiment versus using only the ones in the to quartile of the distribution of IMDb ratings (CR models) (simulations over random subsets of 15 movies). The latter yield relatively lower loss in consumer surplus.

Figure 11. Comparison of the Loss in Consumer Surplus

are available the recommender system that maximizes profit generates 92.3% of the maximal consumer surplus and these statistics are different (p -value < 0.001). Columns (3) and (4) in this figure provide similar statistics for when welfare is maximized and show that in this case limiting the set of movies that can be shown to consumers does not make a significant difference. Overall, these results show that the loss in consumer surplus reduces when the firm recommends only good movies in an attempt to maintain its reputation and earn consumers' trust.

Conclusions

Resorting to recommendations from friends and from the crowd to learn more about products is not new (Dellarocas 2003). However, the complexity of recommender systems increased significantly in recent times with the advent of big data. Furthermore, recommender systems are now ubiquitous and used to find restaurants (e.g., Yelp), shoes (e.g., Zappos), cars (e.g., AutoTrader), households (e.g., Zillow), doctors (e.g., Healthgrades), love partners (e.g., Match.com). It is, therefore, important to study the welfare properties of these systems as they evolve and become increasingly entrenched in our daily life activities. Recommender systems were introduced to help consumers navigate large catalogs of products to find the ones that better match their preferences. But can firms also use recommender systems to extract surplus from consumers? Recommender systems generate more sales compared to markets without them, and this may increase both firm and consumer surplus, but do firms hurt consumers when they choose which type of recommender system to use? In this paper, we compare the welfare properties of different recommender systems. We compare the recommender system that maximizes profit with the recommender system that maximizes consumer surplus and with the recommender system that maximizes total welfare. The former is different from the latter two because under certain conditions firms have incentives to suggest products to consumers in ways that are different from how consumers would like to receive suggestions from the firm. We provide an analytical model showing when this happens and discuss its dynamic properties when consumers interact with the recommender system repeatedly.

We confirm our theoretical results empirically using outcomes from a randomized experiment in the Video-on-Demand system of a large telecommunications provider. For six months, this provider randomized the order in which movies were suggested to consumers on the TV screen and whether they were sold at the original prices negotiated with the content providers or at discounted prices. Using data from this experiment we estimate the demand for movies in this VoD

system as a function of both price and slot on the TV screen. Then, we simulate how consumers would have behaved when presented with movie recommendations from different recommender systems. These systems place different movies in different slots on the TV screen, thus manipulating their saliency, which yields different levels of sales, profit, consumer surplus and welfare. We show that when the firm recommends movies to consumers in an order that maximizes profit, consumers lose surplus relative to the maximum surplus that they could enjoy: 8.44% in our setting, on average across our simulations. Welfare also reduces when the firm maximizes profit: 2.07%, on average across our simulations. Our simulations also show that, for certain subsets of the movies used during our experiment, losses can be significantly larger than these means. We also compare the welfare properties of the recommender systems that maximize profit, consumer surplus and welfare to those of simpler recommender systems often used in practice today, such as lists of the most sold, the most rated and the highest rated movies. We find that the former three recommender systems generate more consumer surplus and more welfare than the latter three. Furthermore, we simulate a content-based recommender system based on movie genre and we show that this system generates as much consumer surplus as the recommender system that maximizes profit. Firms use this type of system to match consumers with products that they are likely to prefer, which may increase their welfare. However, a firm that does so is also likely to be able to charge more for these products, which may then reduce consumer surplus. Therefore, how these systems change consumer well-being is difficult to anticipate and most likely a subject of empirical analysis on a case-by-case basis.

We also study what may happen when consumers interact repeatedly with the recommender system. Analytically, we show that over time the firm is unlikely to maintain extra rents if consumers perceive that they were suggested only "second best" movies in the salient slots. Empirically, we provide evidence that in our setting consumers exhibit a lower price elasticity of demand toward movies placed in the salient slots immediately after the experiment started and that this effect dampens over time. We also show that the preferences of consumers in our experiment did not seem to change over time as a function of the movies suggested to them, lessening the concern that recommendations were only allowing consumers to change their preferences instead of lowering their surplus. We also model recommender systems that maximize profit, consumer surplus and welfare conditional on showing only good movies to consumers. Such recommender systems could be used by the firm to maintain reputation. We show that these systems yield smaller losses in consumer surplus when compared to the ones that use all the movies included in our experiment.

The contribution of our paper is threefold. First, we develop a model to show how firms can use recommender systems to extract surplus from consumers. Using this model, we identify a mechanism underlying recommender systems that firms can exploit to do so, namely that consumers are less price elastic toward products under highlight. Second, we show empirical evidence that the difference in the price elasticity of demand between products placed in salient slots and in non-salient slots, arises in real world settings. In particular, we devise a randomized field experiment to identify this difference empirically. The fact that we use data from a randomized experiment to measure demand is a step forward vis-à-vis the observational work done to date to study how recommender systems affect the distribution of sales. Also, our experiment is large scale with organic consumers, which obviates the usual lack of statistical power and potential biases in the results that arise from small lab-based experiments. Third, we use simulation to compute the appropriate counterfactuals to show how recommender systems compare in terms of welfare properties. To the best of our knowledge, our paper is the first to effectively measure the consumer surplus generated by these systems.

The contribution of our paper to managers and practitioners is significant. With our paper, practitioners are exposed to the appropriate experimental protocol that they should implement to measure how much they can increase profit in their setting relative to their current business practice. Specifically, and in terms of explicit managerial implications, we show conditions under which placing items with larger margins in salient slots is likely to increase profits. In addition, we also show conditions under which doing so hurts consumer surplus. Therefore, our paper also challenges the current state of regulation in online industries that resort to recommender systems to sell products. A significant body of legislation prevents firms from extracting excessive surplus from consumers. For example, mechanisms such as price discrimination are highly scrutinized. Our paper shows that recommender systems may embed a different way of extracting surplus from consumers, one not tied to pricing but rather to how products are presented to consumers in the first place. We provide estimates for the loss in consumer surplus associated to such practice in the case of recommendations for VoD movies. Our findings may generalize to other goods sold in similar online settings given that the recommendation technologies that we study are fairly general and used without much change across many industries. Furthermore, the theoretical model that we explore in this paper is also fairly general and agnostic to the type of product under investigation. Yet, additional empirical studies may still be needed to study how much manipulating recommendations affects consumers in other settings. Our hope is that our paper propels other researchers to develop additional

research on the topic to better inform us of how recommender systems are likely to shape welfare in online commerce.

Finally, we note that our paper does not come without limitations. For example, in our empirical work we aggregate our data across slots on the TV screen, namely for the slots prior to slot and for slots and beyond. A larger and longer experiment could perhaps provide enough statistical power to identify a different price elasticity of demand for each slot separately, which in practice would allow the firm to further increase its profits. Whether this increased level of personalization results in more loss of consumer surplus is still an open question. Our work also suggests that future research could compare the loss in consumer surplus that arises from maximizing profit to that generated by other profit maximizing strategies that firms typically use, such as price discrimination schedules. In addition, one may also compare the welfare properties of recommender systems that maximize profit or that maximize consumer surplus to those of pure content-based recommender systems and of pure collaborative filtering recommender systems. The latter are employed by firms that do not know the demand curve and thus use similarity across items and across users as heuristics to issue reasonable recommendations. Understanding better the relationship between these systems and those that directly maximize welfare properties may be of interest in order to know more about how these heuristics correlate to profit and consumer surplus. Also, our theoretical approach models a pure profit maximizer; that is, a recommender system that maximizes profit irrespective of how similar the items suggested are to the ones previously bought. However, in practice, firms may be more likely to use two-stage profit-aware recommender systems, where in the first stage they determine a set of items that are reasonably similar to the ones bought before and then in the second stage they order such the set of items in a way that maximizes profit. Our approach in this paper models such a system with a loose threshold for similarity, which does not allow us to characterize the trade-offs between profit and similarity that managers can affect in practice by how they set the similarity threshold. Future research may also study in more detail how consumers adjust their demand over time when they realize that they might have been suggested only “second best” products. This question is far from trivial because consumers are unlikely to be shown recommendations in a way that maximize consumer surplus and thus they have no direct observation of their ideal world unless they invest the additional search costs, which then may offset the benefit from easily finding their ideal product. However, whether and how consumers realize this is fundamental to understand how firms need to manage their recommender systems to maintain reputation and consumer trust in settings with repeated interactions. Finally, future research may also

focus on formalizing the potential effects of content-based and collaborative filtering recommender systems on profit, consumer surplus and welfare. This line of research could focus on studying in detail the tension between showing products to consumers that they are more likely to buy, which should increase their surplus, and charging more for them, given that these products match the preferences of consumers better, which should then decrease their surplus.

Acknowledgments

This work has been partially supported by the Portuguese Foundation Science and Technology through fellowships SFRH/BPD/94212/2013, UID/GES/00407/2013, and PTDC/EGE-OGE/27968/2017, and by the iLab and the Living Analytics Research Center at the Heinz College. We also thank our Industry Partner for their support and interest in running such a large-scale randomized experiment.

References

- Adomavicius, G., Bockstedt, J. C., Curley, S. P., and Zhang, J. 2013. "Do Recommender Systems Manipulate Consumer Preferences? A Study of Anchoring Effects," *Information Systems Research* (24:4), pp. 956-975.
- Adomavicius, G., and Tuzhilin, A. 2005. "Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions," *IEEE Transactions on Knowledge and Data Engineering* (17:6), pp. 734-749.
- Anderson, C. 2006. *The Long Tail: Why the Future of Business Is Selling Less of More*, New York: Hyperion.
- Bakos, J. Y. 1997. "Reducing Buyer Search Costs: Implications for Electronic Marketplaces," *Management Science* (43:12), pp. 1676-1692.
- Breese, J. S., Heckerman, D., and Kadie, C. 1998. "Empirical Analysis of Predictive Algorithms for Collaborative Filtering," in *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence*, San Francisco: Morgan Kaufmann Publishers Inc., pp. 43-52.
- Brynjolfsson, E., Hu, Y. J., and Simester, D. 2011. "Goodbye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales," *Management Science* (57:8), pp. 1373-1386.
- Brynjolfsson, E., Hu, Y. J., and Smith, M. D. 2003. "Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers," *Management Science* (49:11), pp. 1580-1596.
- Brynjolfsson, E., Hu, Y. J., and Smith, M. D. 2006. "From Niches to Riches: Anatomy of the Long Tail," *Sloan Management Review* (47:4), pp. 67-71.
- Brynjolfsson, E., Hu, Y. J., and Smith, M. D. 2010. "The Longer Tail: The Changing Shape of Amazon's Sales Distribution Curve," Social Science Research Network.
- Chellappa, R., Konsynski, B., Sambamurthy, V., and Shivendu, S. 2007. "An Empirical Study of the Myths and Facts of Digitization in the Music Industry," paper presented at the 2007 Workshop Information Systems Economics (WISE), Montreal.
- Choudhary, V., and Zhang, Z. 2016. "Recommender Systems and Consumer Product Search," paper presented at the Theory in Economics of Information Systems Workshop, Sonoma, CA.
- Clemons, E. K., Gao, G. G., and Hitt, L. M. 2006. "When Online Reviews Meet Hyperdifferentiation: A Study of the Craft Beer Industry," *Journal of Management Information Systems* (23:2), pp. 149-171.
- Das, A., Mathieu, C., and Ricketts, D. 2009. "Maximizing Profit Using Recommender Systems," arXiv: 0908.3633v1 [CS.cy].
- Dellarocas, C. 2003. "The Digitization of Word of Mouth: Promise and Challenges of Online Feedback Mechanisms," *Management Science* (49:10), pp. 1407-1424.
- Dewan, S., and Ramaprasad, J. 2012. "Research Note—Music Blogging, Online Sampling, and the Long Tail," *Information Systems Research* (23:3, Part 2), pp. 1056-1067.
- Elberse, A., and Oberholzer-Gee, F. 2007. "Superstars and Underdogs: An Examination of the Long Tail Phenomenon in Video Sales," Working Paper, Harvard Business School, Division of Research.
- Fleder, D. M., and Hosanagar, K. 2009. "Blockbuster Culture's Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity," *Management Science* (55:5), pp. 697-712.
- Frank, R. H., and Cook, P. J. 1995. *The Winner-Take-All Society: How More and More Americans Compete for Fewer and Bigger Prizes, Encouraging Economic Waste, Income Inequality, and an Impoverished Cultural Life*, New York: Free Press.
- Godinho de Matos, M., Ferreira, P., Smith, M. D., and Telang, R. 2016. "Culling the Herd: Using Real-World Randomized Experiments to Measure Social Bias with Known Costly Goods," *Management Science* (62:9), 2563-2580.
- Jannach, D., and Adomavicius, G. 2017. "Price and Profit Awareness in Recommender Systems," paper presented at the 2017 Workshop on Value-Aware and Multi-Stakeholder Recommendations.
- Kalra, A., and Goodstein, R. 1998. "The Impact of Advertising Positioning Strategies on Consumer Price Sensitivity," *Journal of Marketing Research* (35:2), pp. 210-224.
- Kaul, A., and Wittink, D. R. 1995. "Empirical Generalizations About the Impact of Advertising on Price Sensitivity and Price," *Marketing Science* (14:3, Supplement), pp. G151-G160.
- Kuksov, D., and Villas-Boas, J. M. 2010. "When More Alternatives Lead to less Choice," *Marketing Science* (29:3), pp. 507-524.
- Mitra, A., and Lynch Jr., J. G. 1995. "Toward a Reconciliation of Market Power and Information Theories of Advertising Effects on Price Elasticity," *Journal of Consumer Research* (21:4), pp. 644-659.
- Muchnik, L., Aral, S., and Taylor, S. J. 2013. "Social Influence Bias: A Randomized Experiment," *Science* (341:6146), pp. 647-651.
- Oestreicher-Singer, G., and Sundararajan, A. 2012. "Recommendation Networks and the Long Tail of Electronic Commerce," *MIS Quarterly* (36:4), pp. 65-83.

- Pathak, B., Garfinkel, R., Gopal, R. D., Venkatesan, R., and Yin, F. 2010. "Empirical Analysis of the Impact of Recommender Systems on Sales," *Journal of Management Information Systems* (27:2), pp. 159-188.
- Resnick, P., and Varian, H. R. 1997. "Recommender Systems," *Communications of the ACM* (40:3), pp. 56-58.
- Rosen, S. 1981. "The Economics of Superstars," *The American Economic Review* (71:5), pp. 845-858.
- Salganik, M. J., Dodds, P. S., and Watts, D. J. 2006. "Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market," *Science* (311:5762), pp. 854-856.
- Salganik, M. J., and Watts, D. J. 2008. "Leading the Herd Astray: An Experimental Study of Self-Fulfilling Prophecies in an Artificial Cultural Market," *Social Psychology Quarterly* (71:4), pp. 338-355.
- Stigler, G. J. 1961. "The Economics of Information," *The Journal of Political Economy* (69:3), pp. 213-225.
- Thaler, R. H., and Sunstein, C. R. 2003. "Libertarian Paternalism," *American Economic Review* (93:2), pp. 175-179.
- Thaler, R. H., and Sunstein, C. R. 2008. *Nudge: Improving Decisions about Health, Wealth, and Happiness*, New Haven, CT: Yale University Press.
- Tucker, C., and Zhang, J. 2007. "How Does Popularity Information Affect Choices? A Field Experiment," *Management Science* (57:5), pp. 828-842.

About the Authors

Xiaochen Zhang is a senior data scientist at Uber. Her Ph.D. work focuses on the welfare impact of Internet technologies, such as recommender systems. She used randomized experiments to estimate the elasticities of demand, allowing her to obtain welfare results using simulations. The insights learned from her empirical work helped shape the analytic and economic model that she developed to describe the players' behavior and predict welfare outcomes.

Pedro Ferreira is an associate professor at the Heinz College, Carnegie Mellon University. His work focuses on how people use technology to consume experience goods and influence others to do

so. These are inextricably linked to how firms behave and how public policies affect market structures. Pedro's work focuses on the application of robust empirical identification methods to analyze large datasets obtained from organic *in vivo* large-scale network-centric randomized experiments. His research spans two interrelated applied areas, the impact of information and communication technologies on education and peer-influence and consumption in the media industry. The bulk of his work is on mediallytics, using big data analytics to understand the future of the media industry. In addition, Pedro has been studying competition, consumer churn, and switching costs in telecommunications.

Miguel Godinho de Matos is an associate professor of Information Systems and Management at Católica Lisbon School of Business and Economics. He is also a visiting research scientist at the Heinz College from Carnegie Mellon University. Miguel's research interests focus on analyzing the impact of digitization, social networks, and peer influence on consumer behavior. Miguel has published his work in top journals and top peer-reviewed research conferences such as *Marketing Science*, *Management Science*, *MIS Quarterly*, *Information Systems Research*, and the Economics of Digitization Seminar Series of the National Bureau of Economic Research.

Rodrigo Belo is an associate professor at the Rotterdam School of Management, Erasmus University. His research focuses on the effects of information systems on organizations and on the impacts of social network structures and peer influence on consumer behavior. His work has been published in top journal and peer-reviewed conferences such as *Management Science*, *Marketing Science*, and the International Conference of Information Systems. Rodrigo has lead and collaborated in multiple projects with established firms and startups in the online and telecommunications sectors. His engagements include the design and deployment of large-scale real world randomized experiments to assess the effectiveness of marketing campaigns and to optimize online user engagement. Rodrigo holds a Ph.D. in Technological Change and Entrepreneurship from Carnegie Mellon University, an M.Sc. in Engineering and Public Policy from Carnegie Mellon University, and a B.Sc. in Computer Science and Engineering from Instituto Superior Técnico, University of Lisbon.

Appendix A

Proofs of All Results in the Paper

RESULT 1: The firm places the product with the highest cost in the salient slot if both costs are lower than a threshold F , that is, if $c_H < F$, with

$$F = \frac{1 + \beta_0 + \beta_1}{\beta_0 + \beta_1} \left[\frac{\beta_0 (1 + \beta_0 + \beta_1)}{(1 + \beta_0)(\beta_0 + \beta_1)} \right]^{\frac{\beta_0}{\beta_1}}$$

PROOF:

$$\prod(p_H^*, p_L^*; s_H = 1, s_L = 0) - \prod(p_H^*, p_L^*; s_H = 0, s_L = 1) = f(c_H) - f(c_L)$$

$$f(c) = -\frac{(\beta_0 + \beta_1)^{\beta_0 + \beta_1}}{(1 + \beta_0 + \beta_1)^{1 + \beta_0 + \beta_1}} c^{1 + \beta_0 + \beta_1} + \frac{\beta_0^{\beta_0}}{(1 + \beta_0)^{1 + \beta_0}} c^{1 + \beta_0}$$

$$f'(c) = -\left(\frac{\beta_0 + \beta_1}{1 + \beta_0 + \beta_1} \right)^{\beta_0 + \beta_1} c^{\beta_0} \left[c^{\beta_1} - \left(\frac{1 + \beta_0 + \beta_1}{\beta_0 + \beta_1} \right)^{\beta_0 + \beta_1} \left(\frac{\beta_0}{1 + \beta_0} \right)^{\beta_0} \right]$$

Therefore, $f(\cdot)$ is increasing when $c < F$ and decreasing otherwise. Consequently, $f(c_H) > f(c_L)$ when $c_H < F$. Thus, in this case,

$$\prod(p_H^*, p_L^*; s_H = 1, s_L = 0) > \prod(p_H^*, p_L^*; s_H = 0, s_L = 1)$$

and the firm places the product with the highest cost in the salient slot. Similarly, the firm places the product with lower cost in the salient slot when $c_L > F$.

RESULT 2: The gain in profit when profit is maximized increases with β_1 .

PROOF:

$$\begin{aligned} \Delta\Pi &= -\left[\frac{(\beta_0 + \beta_1)^{\beta_0 + \beta_1}}{(1 + \beta_0 + \beta_1)^{1 + \beta_0 + \beta_1}} c_H^{1 + \beta_0 + \beta_1} + \frac{\beta_0^{\beta_0}}{(1 + \beta_0)^{1 + \beta_0}} c_L^{1 + \beta_0} \right] \\ &\quad + \left[\frac{(\beta_0 + \beta_1)^{\beta_0 + \beta_1}}{(1 + \beta_0 + \beta_1)^{1 + \beta_0 + \beta_1}} c_L^{1 + \beta_0 + \beta_1} + \frac{\beta_0^{\beta_0}}{(1 + \beta_0)^{1 + \beta_0}} c_H^{1 + \beta_0} \right] \\ \frac{\partial\Delta\Pi}{\partial\beta_1} &= -\frac{(\beta_0 + \beta_1)^{\beta_0 + \beta_1}}{(1 + \beta_0 + \beta_1)^{1 + \beta_0 + \beta_1}} [r(c_L) - r(c_H)] \\ r(c) &= \left[\log\left(\frac{\beta_0 + \beta_1}{1 + \beta_0 + \beta_1} \right) + \log(c) \right] c^{1 + \beta_0 + \beta_1} \end{aligned}$$

$$r'(c) = \left[\log\left(\frac{\beta_0 + \beta_1}{1 + \beta_0 + \beta_1}\right) + \frac{1}{1 + \beta_0 + \beta_1} + \log(c) \right] (1 + \beta_0 + \beta_1) c^{\beta_0 + \beta_1}$$

$$1 + \beta_0 + \beta_1 < 0, -\frac{(\beta_0 + \beta_1)^{\beta_0 + \beta_1}}{(1 + \beta_0 + \beta_1)^{1 + \beta_0 + \beta_1}} > 0$$

Appendix B shows that $r'(c) > 0$ in the region of conflict. Therefore, $r(c_L) > r(c_H)$, which immediately results in the fact that $\Delta\Pi$ increases with β_1 .

RESULT 3: The consumer surplus is maximized when the firm places the product with the highest cost in the non-salient slot if $c_L > G$ with

$$G = \frac{1 + \beta_0 + \beta_1}{\beta_0 + \beta_1} \left[\frac{\beta_0 (1 + \beta_0 + \beta_1)}{(1 + \beta_0)(\beta_0 + \beta_1)} \right]^{\frac{1 + \beta_0}{\beta_1}}$$

PROOF:

$$CS(p_H^*, p_L^*; s_H = 1, s_L = 0) - CS(p_H^*, p_L^*; s_H = 0, s_L = 1) = g(c_H) - g(c_L)$$

$$g(c) = -\frac{(\beta_0 + \beta_1)^{1 + \beta_0 + \beta_1}}{(1 + \beta_0 + \beta_1)^{2 + \beta_0 + \beta_1}} c^{1 + \beta_0 + \beta_1} + \frac{\beta_0^{1 + \beta_0}}{(1 + \beta_0)^{2 + \beta_0}} c^{1 + \beta_0}$$

$$g'(c) = -\left(\frac{\beta_0 + \beta_1}{1 + \beta_0 + \beta_1}\right)^{1 + \beta_0 + \beta_1} c^{\beta_0} \left[c^{\beta_1} - \left(\frac{\beta_0 (1 + \beta_0 + \beta_1)}{(1 + \beta_0)(\beta_0 + \beta_1)}\right)^{1 + \beta_0} \left(\frac{1 + \beta_0 + \beta_1}{\beta_0 + \beta_1}\right) \beta_1 \right]$$

Therefore, $g(c)$ is increasing when $c < G$ and decreasing otherwise. Consequently, $g(c_H) > g(c_L)$ when $c_L > G$. Thus, in this case, $CS(p_H^*, p_L^*; s_H = 0, s_L = 1) > CS(p_H^*, p_L^*; s_H = 1, s_L = 0)$ and consumer surplus is maximized when the firm places the product with the lowest cost in the salient slot. Similarly, consumer surplus is maximized when the firm places the product with highest cost in the salient slot when $c_H < G$.

RESULT 4: The allocation of products to slots that maximize profit does not maximize consumer surplus when $G < c_L < c_H < F$.

PROOF: It suffices to show that $F > G$, which follows immediately from the fact that

$$G = F \left[\frac{\beta_0 (1 + \beta_0 + \beta_1)}{(1 + \beta_0)(\beta_0 + \beta_1)} \right]^{\frac{1}{\beta_1}}, \beta_1 > 0, 1 + \beta_0 < 0, 1 + \beta_0 + \beta_1 < 0$$

RESULT 5: The size of the conflict region for consumer surplus increases with β_1 .

PROOF: We want to show that $\frac{\partial F}{\partial \beta_1} > 0$ and $\frac{\partial G}{\partial \beta_1} < 0$.

$$\begin{aligned} \frac{\partial \log(F)}{\partial \beta_1} &= \frac{\beta_0}{\beta_1^2} (\log(1+x) - x), x = \frac{\beta_1}{\beta_0(1+\beta_0+\beta_1)} > 0 \\ \log(1+0) - 0 &= 0, \frac{\partial \log(1+x) - x}{\partial x} = \frac{1}{1+x} - 1 < 0, \forall x > 0 \\ \log(1+x) - x &< 0, \forall x > 0, \frac{\partial \log(F)}{\partial \beta_1} > 0 \\ \frac{\partial \log(G)}{\partial \beta_1} &= \frac{1+\beta_0}{\beta_1^2} (x + \log(1-x)), x = \frac{\beta_1}{(1+\beta_0)(\beta_0+\beta_1)} > 0 \\ 0 + \log(1-0) &= 0, \frac{\partial x + \log(1-x)}{\partial x} = 1 - \frac{1}{1-x} < 0, \forall x > 0 \\ x + \log(1-x) &< 0, \forall 1 > x > 0, \frac{\partial \log(G)}{\partial \beta_1} > 0 \end{aligned}$$

RESULT 6: The loss in consumer surplus when profit is maximized increases with β_1 .

PROOF:

$$\begin{aligned} \Delta CS &= - \left[\frac{(\beta_0 + \beta_1)^{1+\beta_0+\beta_1}}{(1+\beta_0+\beta_1)^{2+\beta_0+\beta_1}} c_L^{1+\beta_0+\beta_1} + \frac{\beta_0^{1+\beta_0}}{(1+\beta_0)^{2+\beta_0}} c_H^{1+\beta_0} \right] \\ &+ \left[\frac{(\beta_0 + \beta_1)^{1+\beta_0+\beta_1}}{(1+\beta_0+\beta_1)^{2+\beta_0+\beta_1}} c_H^{1+\beta_0+\beta_1} + \frac{\beta_0^{1+\beta_0}}{(1+\beta_0)^{2+\beta_0}} c_L^{1+\beta_0} \right] \\ \frac{\partial \Delta CS}{\partial \beta_1} &= - \frac{(\beta_0 + \beta_1)^{1+\beta_0+\beta_1}}{(1+\beta_0+\beta_1)^{2+\beta_0+\beta_1}} [s(c_L) - s(c_H)] \\ s(c) &= \left[\log \left(\frac{\beta_0 + \beta_1}{1+\beta_0+\beta_1} \right) + \frac{1}{(1+\beta_0+\beta_1)(\beta_0+\beta_1)} + \log(c) \right] + (1+\beta_0+\beta_1)c^{1+\beta_0+\beta_1} \\ s'(c) &= \left[\log \left(\frac{\beta_0 + \beta_1}{1+\beta_0+\beta_1} \right) + \frac{1}{\beta_0+\beta_1} + \log(c) \right] + (1+\beta_0+\beta_1)c^{\beta_0+\beta_1} \\ 1+\beta_0+\beta_1 &< 0, - \frac{(\beta_0 + \beta_1)^{1+\beta_0+\beta_1}}{(1+\beta_0+\beta_1)^{2+\beta_0+\beta_1}} > 0 \end{aligned}$$

Appendix B shows that $s'(c) < 0$ in the region of conflict. Therefore, $s(c_L) > s(c_H)$, from which the results follow immediately.

RESULT 7: Total welfare is maximized when the firm places the product with the highest cost in the non-salient slot when $c_L > H$ with

$$H = \left[\frac{(1+2\beta_0)(1+\beta_0+\beta_1)}{(1+\beta_0)(1+2\beta_0+2\beta_1)} \right]^{\frac{1}{\beta_1}} \left[\frac{\beta_0(1+\beta_0+\beta_1)}{(1+\beta_0)(\beta_0+\beta_1)} \right]^{\frac{\beta_0}{\beta_1}} \left(\frac{1+\beta_0+\beta_1}{\beta_0+\beta_1} \right)$$

PROOF:

$$W(p_H^*, p_L^*; s_H = 1, s_L = 0) - W(p_H^*, p_L^*; s_H = 0, s_L = 1) = h(c_H) - h(c_L), h(c) = f(c) + g(c)$$

$$h'(c) = \frac{1 + 2\beta_0 + 2\beta_1}{1 + \beta_0 + \beta_1} \left(\frac{\beta_0 + \beta_1}{1 + \beta_0 + \beta_1} \right)^{\beta_0 + \beta_1} c^{\beta_0}$$

$$\left[c^{\beta_1} - \frac{(1 + 2\beta_0)(1 + \beta_0 + \beta_1)}{(1 + \beta_0)(1 + 2\beta_0 + 2\beta_1)} \left[\frac{\beta_0(1 + \beta_0 + \beta_1)}{(1 + \beta_0)(\beta_0 + \beta_1)} \right]^{\beta_0} \left(\frac{1 + \beta_0 + \beta_1}{\beta_0 + \beta_1} \right)^{\beta_1} \right]$$

Therefore, $h(c)$ is increasing when $c < H$ and decreasing otherwise. Consequently, $h(c_H) > h(c_L)$ when $c_L > H$. Thus, in this case, $W(p_H^*, p_L^*; s_H = 1, s_L = 0) > W(p_H^*, p_L^*; s_H = 0, s_L = 1)$ and total welfare is maximized when the firm places the product with the lowest cost in the salient slot. Similarly, welfare is maximized when the firm places the product with the lowest cost in the salient slot when $c_H < H$.

RESULT 8: The allocation of products to slots that maximize profit does not minimize total welfare when $F > c_H > c_L > H$.

PROOF: It suffices to show that $F > H$, which follows immediately from the fact that

$$H = F \left[\frac{(1 + 2\beta_0)(1 + \beta_0 + \beta_1)}{(1 + \beta_0)(1 + 2\beta_0 + 2\beta_1)} \right]^{\frac{1}{\beta_1}}, \beta_1 > 0, 1 + \beta_0 < 0, 1 + \beta_0 + \beta_1 < 0$$

Appendix B

Loss in Consumer Surplus as a Function of β_1

The “Theoretical Model” section shows that

$$\frac{\partial \Delta CS}{\partial \beta_1} = - \frac{(\beta_0 + \beta_1)^{1 + \beta_0 + \beta_1}}{(1 + \beta_0 + \beta_1)^{2 + \beta_0 + \beta_1}} [s(c_L) - s(c_H)]$$

$$s(c) = \left[\log \left(\frac{\beta_0 + \beta_1}{1 + \beta_0 + \beta_1} \right) + \frac{1}{(1 + \beta_0 + \beta_1)(\beta_0 + \beta_1)} + \log(c) \right] c^{1 + \beta_0 + \beta_1}$$

$$s'(c) = \left[\log \left(\frac{\beta_0 + \beta_1}{1 + \beta_0 + \beta_1} \right) + \frac{1}{\beta_0 + \beta_1} + \log(c) \right] (1 + \beta_0 + \beta_1) c^{\beta_0 + \beta_1}$$

The loss in consumer surplus increases with β_1 if $s'(c) < 0$ in the region of conflict. In this region, $c > G$, with

$$G = \frac{1 + \beta_0 + \beta_1}{\beta_0 + \beta_1} \left[\frac{\beta_0(1 + \beta_0 + \beta_1)}{(1 + \beta_0)(\beta_0 + \beta_1)} \right]^{\frac{1 + \beta_0}{\beta_1}}$$

Therefore, the result follows from

$$\begin{aligned} \log(c) &> \frac{1+\beta_0}{\beta_1} \log\left(\frac{\beta_0(1+\beta_0+\beta_1)}{(1+\beta_0)(\beta_0+\beta_1)}\right) - \log\left(\frac{\beta_0+\beta_1}{1+\beta_0+\beta_1}\right) \\ s'(c) &< \frac{(1+\beta_0)(1+\beta_0+\beta_1)}{\beta_1} c^{\beta_0+\beta_1} [\log(1-w)+w], 1 > w \\ &= \frac{\beta_1}{(1+\beta_0)(\beta_0+\beta_1)} > 0 \\ \log(1-w)+w &< 0, \forall 1 > w > 0, \frac{(1+\beta_0)(1+\beta_0+\beta_1)}{\beta_1} > 0 \end{aligned}$$

The ‘‘Theoretical Model’’ section shows that

$$\begin{aligned} \Delta\Pi &= -\left[\frac{(\beta_0+\beta_1)^{\beta_0+\beta_1}}{(1+\beta_0+\beta_1)^{1+\beta_0+\beta_1}} c_H^{1+\beta_0+\beta_1} + \frac{\beta_0^{\beta_0}}{(1+\beta_0)^{1+\beta_0}} c_L^{1+\beta_0}\right] \\ &+ \left[\frac{(\beta_0+\beta_1)^{\beta_0+\beta_1}}{(1+\beta_0+\beta_1)^{1+\beta_0+\beta_1}} c_L^{1+\beta_0+\beta_1} + \frac{\beta_0^{\beta_0}}{(1+\beta_0)^{1+\beta_0}} c_H^{1+\beta_0}\right] \\ \frac{\partial\Delta\Pi}{\partial\beta_1} &= \frac{(\beta_0+\beta_1)^{\beta_0+\beta_1}}{(1+\beta_0+\beta_1)^{1+\beta_0+\beta_1}} [r(c_L)-r(c_H)] \\ r(c) &= \left[\log\left(\frac{\beta_0+\beta_1}{1+\beta_0+\beta_1}\right) + \log(c)\right] c^{1+\beta_0+\beta_1} \\ r'(c) &= \left[\log\left(\frac{\beta_0+\beta_1}{1+\beta_0+\beta_1}\right) + \frac{1}{1+\beta_0+\beta_1} + \log(c)\right] (1+\beta_0+\beta_1) c^{\beta_0+\beta_1} \end{aligned}$$

In the conflict region $F > c$, with

$$F = \frac{1+\beta_0+\beta_1}{\beta_0+\beta_1} \left[\frac{\beta_0(1+\beta_0+\beta_1)}{(1+\beta_0)(\beta_0+\beta_1)}\right]^{\frac{\beta_0}{\beta_1}}$$

Therefore,

$$\begin{aligned} \log(c) &< -\frac{\beta_0}{\beta_1} \log\left(\frac{(1+\beta_0)(\beta_0+\beta_1)}{\beta_0(1+\beta_0+\beta_1)}\right) - \log\left(\frac{\beta_0+\beta_1}{1+\beta_0+\beta_1}\right) \\ r'(c) &> \frac{\beta_0(1+\beta_0+\beta_1)}{\beta_1} c^{\beta_0+\beta_1} [w-\log(1-w)], 1 > w = \frac{\beta_1}{\beta_0(1+\beta_0+\beta_1)} > 0 \\ w-\log(1-w) &> 0, \forall 1 > w > 0, \frac{\beta_0(1+\beta_0+\beta_1)}{\beta_1} > 0 \end{aligned}$$

Thus, $r'(c) > 0$ in this conflict region.

Appendix C

Numerical Examples of Loss

Figure C1 shows a numerical example in which two products, A and B are shown to one representative consumer in two slots. The representative consumer exhibits different price elasticities of demand for these slots. The assignment of products to slots that maximizes consumer surplus places product B in the first slot and product A in the second slot. The assignment of products to slots that maximizes profit does the opposite and doing so reduces consumer surplus. This case also shows that maximizing profit may also reduce total welfare.

<ul style="list-style-type: none"> • Demand: $d_j = p_j^{-2+0.8 \text{ (slot=1)}}$ • Costs: $c_A = \\$2, c_B = \\1 		
	Maximize Consumer Surplus	Maximize Profits
Slot assignments	B A	A B
Optimal prices	$p_A = \$4; p_B = \6	$p_A = \$12; p_B = \2
Demand	$d_A = 0.06; d_B = 0.12$	$d_A = 0.05; d_B = 0.25$
Profit	\$0.71	\$0.76
Consumer Surplus	\$3.74	\$3.54
Welfare	\$4.45	\$4.30

Figure C1. Example of Loss in Consumer Surplus with One Representative Consumer

Figure C2 shows a numerical example in which two products, A and B are shown to two representative consumers in two slots. A setting with two representative consumers mimics a taste parameter and introduces the opportunity for personalization, that is, the order in which the firm shows the products to one consumer is different from the order in which it shows them to the other consumer. As before, consumers exhibit different price elasticities of demand for the two slots. The assignment of products to slots that maximizes consumer surplus places product A in the first slot and product B in the second slots for both consumers. The assignment of products to slots that maximizes profit does not. Instead, it assigns product B to the first slot and product A to the second slot for one of the consumers, keeping the same assignment of products to slots as before for the other consumer, thus taking advantage of personalization. This reduces consumer surplus. In this case, maximizing profit increases welfare.

<ul style="list-style-type: none"> • Demand: $d_j = p_j^{-2+0.3 \text{ (slot=1)}+0.6 \text{ (consumer=1)}}$ • Costs: $c_A = \\$2, c_B = \\1 				
	Maximize Consumer Surplus		Maximize Profits	
Consumer	1 2		1 2	
Slot assignments	A B	A B	B A	A B
Optimal prices	$p_A = \$5.6; p_B = \4.7		$p_A = \$13.0; p_B = \2.9	
Demand	$d_A = 0.12; d_B = 0.25$		$d_A = 0.07; d_B = 0.39$	
Profit	\$2.32		\$3.52	
Consumer Surplus	\$12.45		\$11.50	
Welfare	\$14.77		\$15.02	

Figure C2 . Example of Loss in Consumer Surplus with Two Representative Consumers

Appendix D

Distribution of Loss

The following figures show the distribution of losses across our 1,000 simulations. Figure D1 shows the loss in consumer surplus when profit is maximized and Figure D2 shows the loss in profit when consumer surplus is maximized. Figure D3 shows the loss in welfare when profit is maximized and Figure D4 shows the loss in welfare when consumer surplus is maximized. The average loss in consumer surplus when profit is maximized is 8.44% and the loss in welfare when profit is maximized is 2.07%. The loss in welfare when consumer surplus is maximized is 1.62% and this statistic is smaller than the latter. These figures also show that the loss in consumer surplus and the loss in welfare when profit is maximized can be substantial.




Figure D1. Distribution of Loss in Consumer Surplus When Profit Is Maximized

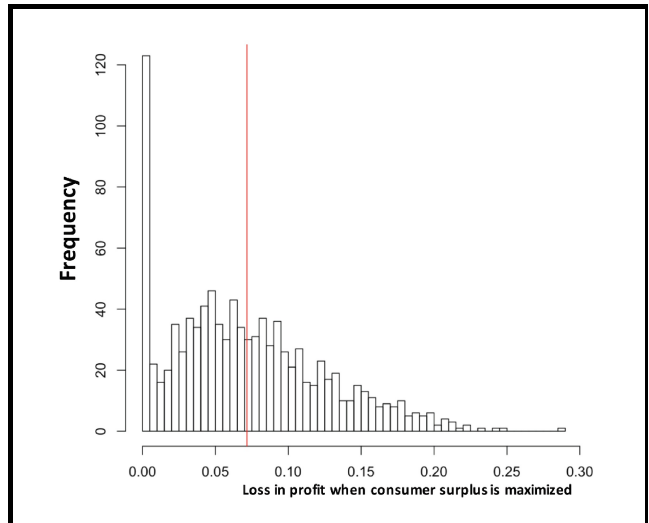


Figure D2. Distribution of Loss in Profit When Consumer Surplus Is Maximized




Figure D3. Distribution of Loss in Welfare When Profit Is Maximized

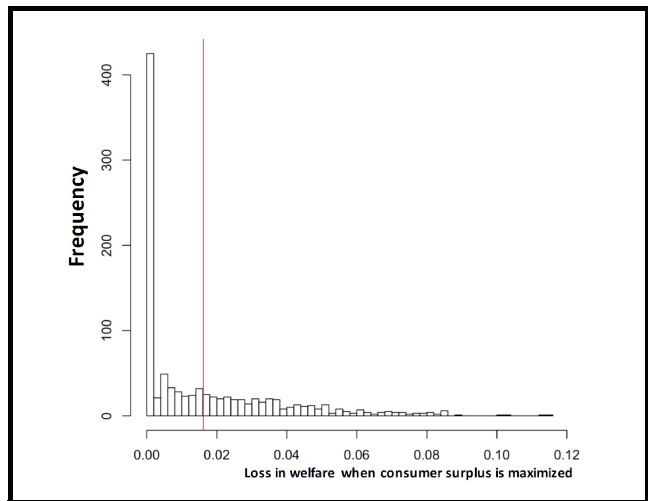


Figure D4. Distribution of Loss in Welfare When Consumer Surplus Is Maximized

Appendix E

Balance for Movie Characteristics Across TV Slots

This appendix uses Tukey family-wise plots to provide additional evidence that the movies placed in different TV slots under the “Good Deals” menu during our experiment were similar on observables, namely price, age, IMDD ratings and IMDb votes. The figures below show comparisons for all pairs of slots.




Figure E1. Tukey Plot for Covariate Price (all contrasts across all movie slots are not statistically different from zero)




Figure E2. Tukey Plot for Covariate IMDb Rating (all contrasts across all movie slots are not statistically different from zero)




Figure E3. Tukey Plot for Covariate IMDb Votes (all contrasts across all movie slots are not statistically different from zero)




Figure E4. Tukey Plot for Covariate Age (all contrasts across all movie slots are not statistically different from zero)

Appendix F

Movie Allocation Across Households

This appendix provides evidence on observables that movies were randomly allocated across households during the experiment. Figure F2 shows that indeed the sampling probability was different across movies, with some movies obtaining more exposure than others. The other figures in this section show that the characteristics of the households that got these movies under the new menu are similar in a number of fundamental observables across movies.




Figure F1. Number of Households Exposed to Each Movie in the Experiment




Figure F2. Average Household ARPU for All Movies Used in the Experiment




Figure F3. Proportion of Households with Active Contract for All Movies Used in the Experiment




Figure F4. Proportion of Households with Premium Sports Channels for All Movies Used in the Experiment




Figure F5. Proportion of Households with Premium Movie Channels for All Movies Used in the Experiment


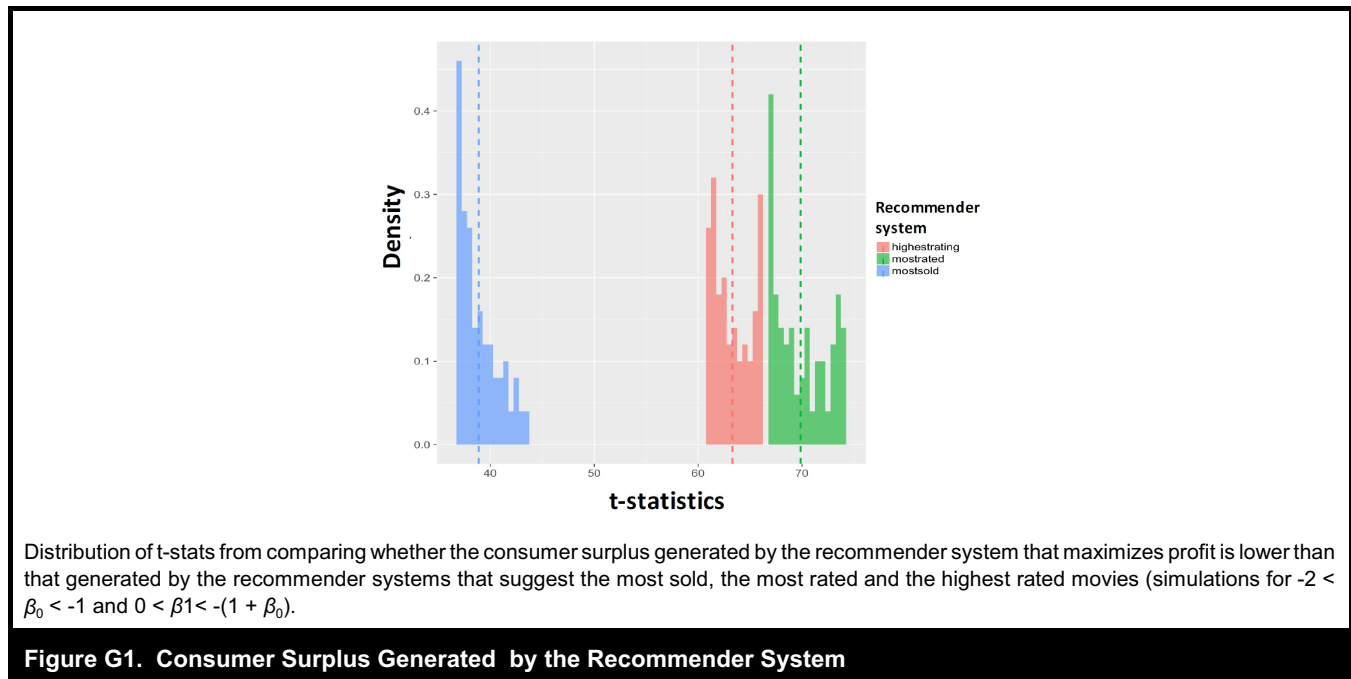



Figure F6. Proportion of Premium Households for All Movies Used in the Experiment

Appendix G

Robustness Check for Different Levels of Price Elasticity on Demand

Figure G1 and Figure G2 show that this finding does not arise because TELCO charges consumers in the inelastic portion of the demand curve. Figure G1 shows the distribution of t-stats obtained from testing whether the recommender system that maximizes profit generates lower consumer surplus than the recommender systems that suggest the most sold, the most rated and the highest rated movies to consumers. Results are obtained for a set of 300 simulations varying between and and between and . For each of these simulations, we perform 1,000 simulations each with a random subset of 15 movies used during the experiment. In each of these simulations the firm now chooses both the order in which movies are recommended to consumers as well as their price in order to maximize profit. In all cases, we conclude that the recommender system that maximizes profit does not generate lower consumer surplus. Figure G2 replicates the analysis for the case of welfare. In this case, we conclude that the recommender system that maximizes profit does not generate lower welfare than other three recommender systems referred above. Therefore, the fact that the recommender system that maximizes profit does not seem to hurt consumers and welfare too much when compared to simpler recommender systems does not seem to be an artifact of the low prices charged by TELCO to consumers.





Distribution of t-stats from comparing whether the welfare generated by the recommender system that maximizes profit is lower than that generated by the recommender systems that suggest the most sold, the most rated, and the highest rated movies (simulations for $-2 < \beta_0 < -1$ and $0 < \beta_1 < -(1 + \beta_0)$).

Figure G2. Welfare Generated by the Recommender System

Copyright of MIS Quarterly is the property of MIS Quarterly and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.