Rev Ind Organ (2014) 44:131–146
DOI 10.1007/s11151-013-9406-8

Inducing Customers to Try New Goods

Alessandro Acquisti

Published online: 30 November 2013

Abstract In recent years, progresses in data mining and business analytics have fostered the advent of recommender systems, behavioral advertising, and other ways of using consumer data to personalize offers and products. We investigate the incentives for sellers to invest in systems that allow the tracking of consumers and then to truthfully report whether potential buyers will enjoy yet untried products. We find that there are two types of equilibria: For some parameter values, sellers will target all potential buyers, hence their targeted ads or purchase recommendations provide no benefit to the consumer. But for other values, ads and recommendations will be accurate. In particular, the incentive for the seller to provide accurate ads and recommendations will be inversely related to the difference between the cost of producing the good and its average market evaluation.

Keywords Recommender systems · Behavioral advertising · Targeted advertising · Business analytics · Data mining · Privacy

1 Introduction

Since many consumer transactions nowadays are computer mediated, merchants have become quite adept at leveraging knowledge of consumers' traits and behavior to predict their preferences, reservation prices, and future choices. The scenario where a seller is able to predict a consumer's future satisfaction with a product better than the consumer can herself may have sounded preposterous a few years ago; but it has become quite reasonable now, due to progress in data mining and business analytics

A. Acquisti (✉)
Carnegie Mellon University, Heinz College, 5000 Forbes Avenue, Pittsburgh, PA 15232, USA
e-mail: acquisti@andrew.cmu.edu

which have fostered the advent of recommender systems, behavioral advertising, and other ways of using consumers' data to target and personalize offers and products.

In the case of so-called "recommender systems" (Resnick and Varian 1997), sellers employ data mining or collaborative filtering to combine information coming from the transaction histories (and degrees of satisfaction with said transactions) of previous customers. With this information, sellers then decide what good or service to offer or recommend to her next.[1] Examples of such systems abound both offline and online (Kohavi and Provost 2001). In the offline world, information from supermarket scanners is used to develop customized coupons based on past and contemporaneous purchases. In the online world, Amazon.com has been a pioneer in developing sophisticated recommender systems to promote books, CDs, and other products based on the contemporaneous purchases of many consumers (Schafer et al. 1999). Netflix has continued to invest to improve algorithms to predict a user's degree of appreciation of a movie, based on movie ratings by the universe of its customers (Bennett and Lanning 2007).

In the case of behavioral advertising (and affine concepts such as behavioral targeting, targeted advertising, or tailored ads), a marketer can combine information from a consumer's online browsing behavior with knowledge or inferences about her traits to predict which ad will likely match her interests. This strategy is more likely to generate a sale or at least a clickthrough. Then, the marketer will show that ad to the consumer. The number of commercial systems that track online consumers' online behavior (often across multiple sites) to offer such targeted ads has been steadily increasing in recent years: from Yahoo! Smart Ads to DoubleClick, from TACODA to NebuAd (Yan et al. 2009).

The academic literature on behavioral targeting has only recently developed (Yan et al. 2009; Goldfarb and Tucker 2011; Beales 2010), and there is little research about situations in which the recommender is, actually, the seller. This paper is concerned with the incentives for a seller to report *accurately* its beliefs about whether or not potential buyers will like untried products, based on the analysis of the consumer's (and her peers') previous behavior. The scenario we consider shares some characteristics with recommender systems, in that the merchant's predictions are based on feedback from multiple consumers. It also shares characteristics with targeted or behavioral advertising, where the merchant must decide whether to use its own predictions truthfully when contacting the consumer with personalized offers.

We investigate a simple scenario for both one-shot and repeated customer-merchant interactions, involving repeated purchase and multiple goods. In the one-shot interaction, we find that there are two types of equilibria, depending on various parameter values: For some values, sellers will target all potential buyers, so such targeted ads and purchase recommendations provide no benefit to the consumer. But for other values,

[1] In this context, data mining (Fayyad et al. 1996) typically refers to the analysis of vast amounts of diverse types of data, searching for interesting patterns and correlations. Collaborative filtering (Resnick et al. 1994) refers to "filtering" information through the collaboration of multiple entities—for instance, predicting a consumer's preferences using the collected preferences of many other consumers. Data mining and collaborative filtering can be used in recommender systems, which try to predict an agent's rating of an item, in order to provide useful item recommendations. See also Kohavi and Provost (2001).

ads and recommendations will be accurate. In particular, the incentive for the seller to provide accurate offers will be inversely related to the difference between the cost of producing the good and its average market evaluation.

With repeated interaction, the merchant is more likely to make "truthful" offers: those that the merchant expects to match the consumers actual preferences. This may be particularly the case when consumers can communicate among themselves easily, as is the case in an on-line environment, since the cost of a bad reputation may be much higher than without communication.

2 Related Work

Research on recommender systems dates to the commercial growth of the Internet (Resnick and Varian 1997; Avery et al. 1999; Dellarocas 2003). Over time, researchers in this field have considered mechanisms and incentives to ensure that feedbacks and recommendations are truthful—i.e., that consumers' feedback about products is honestly reported, and that merchants' recommendations accurately reflect the expected level of consumer's satisfaction with a new product (Resnick and Sami 2007; Miller et al. 2005; Victor et al. 2009). Many of these efforts have focused on the incentives for consumers to participate in the system as recommenders. Relatively less attention has been allocated to the case of recommendation systems where the recommender is the seller itself. For instance, consider the scenario in which the seller uses collaborative filtering of consumers' data to predict the expected satisfaction that a certain customer will derive from a new good. Now, the seller must decide whether to offer her that good or not. What are the incentives for a seller to use its analysis accurately and offer an honest recommendation to the consumer? As a consequence, under which conditions will such technologies benefit the seller as well as its customers?

In order to determine whether or not a firm will target honestly, we develop a flexible model that borrows from several streams of literature. From a modeling perspective, different streams in the microeconomic, information systems, and marketing literature can be compared to the formal analysis we present here. First, our analysis has aspects in common with the literature on coupons (Shaffer and Zhang 1995; Krishna and Zhang 1999; Zhang et al. 2000) and the supply of information by sellers in order to elicit purchase and facilitate price discrimination (Lewis and Sappington 1994; Blume 1998). It also has commonalities with the empirical literature on customer tracking and the value of a customer's information (McCulloch et al. 1996; Rossi and Allenby 1998). Since we consider the case of repeat interaction between the seller and buyers, the literature on trust, reputation, and recommendation systems is also related to this study (Crawford and Sobel 1982; Shapiro 1983; Allen 1984; Sobel 1985; Moorthy and Srinivasan 1995; Che 1996; Ansari et al. 2000).

The main difference between our model and those of previous studies is that we are considering a unique type of signaling game (Maskin and Tirole 1990, 1992). In our approach, a non-peer principal (seller) knows more than does each agent (consumer) about what satisfy the agent's tastes, because the principal has access to information about the buying habits and satisfaction levels of agents of similar type. This is related to the work of Prendergast (1993) and Ottaviani and Sørensen (2006a). It is also related to the literature on experts found in Crawford and Sobel (1982), Benabou and Laroque (1992), Morgan and Stocken (2003), Krishna and Morgan (2001), Morris (2001), Ely and Valimaki (2003), and Ottaviani and Sørensen (2006b); as well as to the literature on social learning and experimentation found in Liebeskind and Rumelt (1989), McFadden and Train (1996), and Schlee (2001).

3 Setup

For the rest of the paper we will suppose that the seller is a monopolist that produces a good at marginal cost c. This good provides a utility $v > 0$ to a fraction π (with $0 < \pi < 1$) of the consumer population of buyers and a utility of zero to the rest of the population. The fraction π is common knowledge to both the buyers and the seller, but the consumer does not know whether she will like the product or not. The monopolist, on the other hand, may apply data mining or collaborative filtering tools to consumer data to estimate whether a certain customer will like a yet untried good, and will decide whether to advertise it and recommend it to the user or not. These recommendations are "personalized." For instance, the seller may have 100 products to sell, but among those 100 products it will choose which one to recommend to each of its consumers. Therefore, different customers may get different recommendations.

As we noted above, the assumption that a seller may be able to predict a consumer's *ex post* satisfaction with a product better than the consumer can predict that herself *ex ante* may have sounded preposterous a few years ago. Nowadays, however, a merchant such as Amazon.com can compare the history of purchases of a given customer i to the history of purchases of large numbers of other customers. Once it finds similarities between patterns of purchases, Amazon may be able to predict that customer i may be interested in product x (which the customer has never tried before), because customers similar to i have also purchased x. The patterns may be not obvious. That is, customer i may have a history of purchasing garden tools, and yet product x could be a clothing item. Similarly, Netflix can compare the ratings that its large base of users have assigned to movies they have watched; then, based on similarities across the ratings of a subset of customers, it can estimate a given customer's future rating (hence, the expected satisfaction) of a movie she has not yet watched.

For the rest of the model's description, we will simply refer to the seller's "recommendation," by which we refer to any means—such as ads, targeted offers, coupons, "You may also like this" messages—that the seller can employ to induce the consumer to purchase a good. We assume that the consumer is aware that the coupon, recommendation, or ad is in fact coming from the same merchant that is selling the product.

In what follows we consider several variations of a dynamic model with more than one period. We distinguish whether the goods being recommended over time are the same repeated purchase good or a series of different, multiple goods. Before purchasing the product for the first time, the consumer is uncertain about the product's desirability. The first scenario corresponds to situations in which the merchant is offering a good that may be consumed repeatedly (for instance, a new type of pasta sauce launched in the marketplace, which the consumer could try and then purchase repeatedly over time). The second scenario corresponds to situations in which the merchant is offering a series of different experience goods, each of them being consumed just once (for instance, a different book that is offered for purchase at each period).

Within each scenario, we consider the case in which the merchant is unable to obtain useful information from the analysis of the consumers (baseline case), and the case in which the merchant obtains perfectly accurate information about the consumers' tastes. We then consider various possible extensions of the basic model.

4 Repeated Purchase

We first consider the case when the merchant recommends to the customer a good that may be repeatedly bought over several periods.[2]

4.1 Symmetric Incomplete Information

As a baseline scenario, we begin by assuming that neither the buyer nor the seller have access to data mining or collaborative filtering technologies, and therefore cannot know whether the buyer will like a particular good (symmetric incomplete information).

There are two periods (defined as units of time during which one transaction can take place between the merchant and the consumer): In the first period the merchant sets the price, and the potential buyers may choose to try the good. In the second period, the consumers who liked the good will purchase it again, as long as the price is less than or equal to its value to them. Let p_1 be the price set in the first period and p_2 the price set in the second period. We assume that the seller cannot commit to the second-period price.

There are, essentially, two cases. When the expected value exceeds cost ($\pi v > c$), any first-period price between πv and c will induce the consumers to try the good. Profit maximization implies that the monopolist will set the first-period price equal to πv and the second-period price equal to v.[3] This yields the monopolist a profit of $(\pi v - c)$ in the first period and $\pi(v - c)$ in the second period.

On the other hand, when the expected value is less than cost ($\pi v < c$), the monopolist will try to follow the same strategy, setting $p_1 = \pi v$ and $p_2 = v$, making a profit of $(\pi v - c) + \pi(v - c)$.[4] However, if π is sufficiently small, this strategy will not be viable, since the losses in the first period may not be recouped by second-period profit, and the seller will not offer the product.

[2] We focus on a single product recommendation per period. This approach is justified not just by the fact that certain recommender systems work that way (for instance, Amazon.com may send its customers emails recommending a specific new product), but also because, even when the merchant is offering multiple products, each recommendation can be analyzed in isolation from the others. For instance, the consumer may choose to try one among the products being offered; her reaction to that single product will influence her future interactions with that merchant, as detailed in the model, even if the merchant had recommended other products as well.

[3] As usual, one can resolve the indifference in favor of the seller, since the seller could charge a price slightly smaller than v and create a strict decision to purchase.

[4] We ignore the equality case, as it has probability zero of occurring.

In this model, the consumer is being offered a first-period price that compensates him or her for the risk of trying the good. It is trivial to show that the results do not change when we extend the analysis from 2 to $n = 3, \ldots, \infty$ periods, although as n increases the seller will be able to sustain higher initial costs.

Since we are considering a repeated purchase scenario, it is worthwhile to ask whether commitment changes the dynamics of the model relative to the no commitment case that we have considered above. Specifically: Can the seller charge a higher price in the first period by promising to keep lower prices in the following periods, or vice-versa? The answer is no, which means that the seller's inability to commit does not penalize it.

Imagine there are n periods. For the seller, for commitment to be optimal it must be:

$$\pi v + \epsilon - c + n\pi(\pi v + \alpha - c) > \pi v - c + \pi n(v - c) \tag{1}$$

where $\pi v + \epsilon$ is the price charged in the first period, and $\pi v + \alpha$ is the price charged in the subsequent n periods.

This simplifies to:

$$\epsilon + n\pi\alpha > (1 - \pi)\pi n v. \tag{2}$$

However, for the consumer to accept a higher price in the first period it must be that:

$$\pi(v - \pi v - \epsilon) + n\pi(v - \pi v - \alpha) > (1 - \pi)(\pi v + \epsilon) \tag{3}$$

since there is a probability π that she will like the good and a probability $1 - \pi$ that she will not. This reduces to:

$$(1 - \pi)n\pi v > \alpha n\pi + \epsilon. \tag{4}$$

Adding the two inequalities gives a contradiction.

4.2 Asymmetric Perfect Information

Let us now assume that the producer knows which consumers will like a new good, thanks to data mining or collaborative filtering tools. As discussed earlier, this information may come from comparing the purchase history of the current customer to similar purchase histories of her peer consumers. Each consumer, on the other hand, believes that there is a probability π that she will like the product, just as before, but does not know for certain whether she will or not.

The game now becomes a signaling game, albeit of a somewhat unusual form. In this game the *seller* knows the buyer's type, but the *buyer* is not sure about her own type—which in this model is the buyer's taste relative to the attributes of the good. This is the scenario we have highlighted above, which novel technologies have made

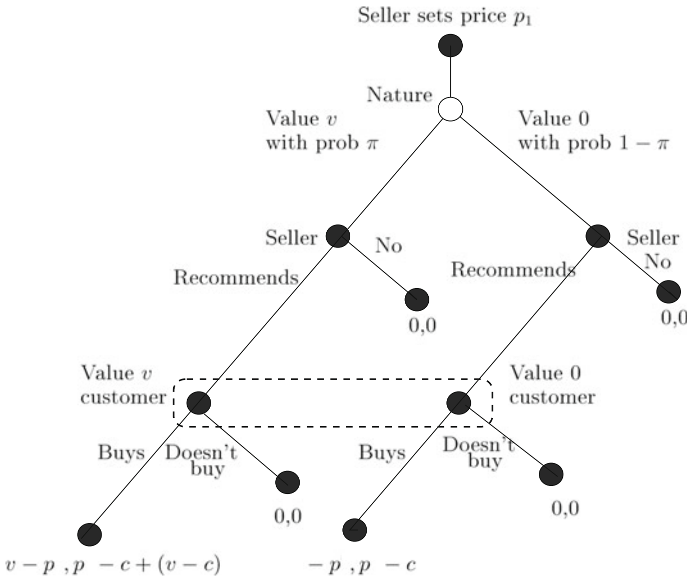

Fig. 1 First period of the asymmetric information, uniform price game. The seller moves first, deciding at which price to sell, then observes the consumer's type. The consumer has value v or value 0, but does not know which type he is, as indicated by the *dashed* information set. The payoffs for the seller and the consumer refer to the transactions with a single consumer over two periods of the game

possible, such that a seller may be able to predict a consumer's *ex post* satisfaction with a product better than the consumer can, *ex ante*. The consumer has, of course, some expectations (v and π are common knowledge). However, the consumer cannot predict precisely what her satisfaction with the good will be; thanks to the wealth of data it possesses, the seller's prediction is more accurate.

We consider two variants of the game. The first is called the "uniform price" variant. The game tree for the uniform price game is given in Fig. 1 (for simplicity, the tree shows only the first period of the game, as the actions in the second period follow directly from the first; the payoffs refer to the entire game). The seller first chooses the price, which is publicly announced and is applied uniformly to all customers during a given period of the game. The seller then observes the agent's type (represented, in the figure, by Nature's choice) and decides whether or not to signal the agent. The signal is typically an ad or a recommendation, but it could, in principle, be any action taken by the seller that the buyer can observe. The buyer is unaware of which type she is (as indicated by the dashed information set), and has to decide whether to purchase or not based on the signal. The question of interest is: Under what conditions will the seller send an "honest" signal?

In the second period (not shown in the figure), the seller will always set a price equal to v and sell only to the type v consumers. Hence the seller always wants at least the type v consumers to purchase in the first period.

As before, we have two cases: If $\pi v > c$, the seller will set a first-period price πv and sell to everyone. In the second period, it will sell only to those who like the product. The profit from this will be $(\pi v - c) + \pi(v - c)$, as before. In this case, even if the producer knows who will like the good, there is no reason for it to reveal this, since it benefits from the consumer experimentation.

Note that when the seller knows who will like the good, this equilibrium is inefficient. The fully efficient outcome would involve the seller telling the type v consumers to buy, realizing profit of $2\pi(v - c)$. However, this is not an equilibrium, since if the seller believed that consumers would purchase based on its recommendation, it would optimally recommend to all consumers, making the signal uninformative. If the signal is uninformative, the consumers would not be willing to pay v.

The second case, where $\pi v < c$, is more interesting. We will consider strategies for the buyer of the form "Buy if receive signal and $Eu \geq p$," where Eu is the expected utility for the buyer. We seek a perfect Bayesian equilibrium, which requires that the subjective posterior probabilities in the expectation be consistent with Bayes' Law.[5]

Proposition 1 *Assume that $v > c$, but $\pi v < c$. Then in a perfect Bayesian equilibrium $p_1 = c$, the seller always signals honestly, and the buyer purchases on receipt of the signal.*

Proof See "Appendix."

In this equilibrium, the *ex ante* expected utility is $\pi(v - c)$ for the buyers and $\pi(v - c)$ for the seller, whereas in the symmetric incomplete information case, the *ex ante* expected utility of the buyers is 0 and the expected utility of the seller is $(\pi v - c) + \pi(v - c)$. Hence, this equilibrium Pareto dominates the no information case. It is, of course, Pareto efficient, since the sum of the utilities is $2(v - c)$, which is the maximum possible.

Intuitively, since experimentation is costly when $\pi v < c$, it is in the interest of the monopolist to minimize experimentation, so it chooses truthfully to reveal the types. An alternative interpretation is that when the market evaluation, πv, is high relative to the cost of the good, the seller will tend to recommend the good to everybody, hence it will provide also "dishonest" recommendations: Goods that are highly desirable will induce the seller to act dishonestly by recommending the good to those who will not necessarily find it attractive. This equilibrium is inefficient, because the seller actually knows which customers will like the good but decides to ignore that information. On the other hand, when producing the good is costly compared to its market evaluation, seller recommendations will be honest, and the outcome will be Pareto efficient. Thus, recommendation systems may be socially desirable for niche goods (with small π) for which experimentation would otherwise be too costly.

It is straightforward to prove that these results do not change when the number of periods increases from 2 to an infinite horizon.

The second variant of the game is the "personalized price" case, where the seller can condition the price on the consumer type. The game tree is depicted in Fig. 2 (again, the tree shows only the first period of the game, from which the actions in the second period follow directly, but the payoffs refer to the entire game). The buyer is

[5] We need to restrict the set of allowable strategies. Otherwise strategies such as "buy only if the price is greater than $Eu/2$ and less than or equal to Eu" would be possible.

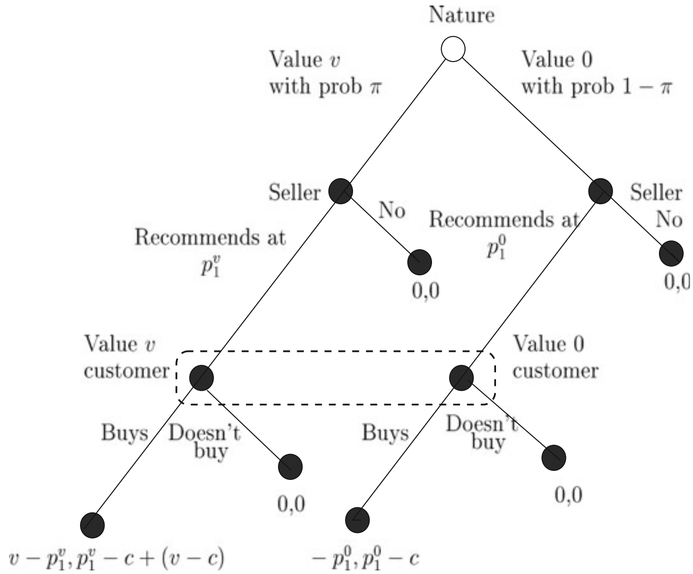

Fig. 2 First period of the asymmetric information, personalized price game. The seller observes the buyer's type, then sets price. The consumer has value v or value 0, as indicated by the *dashed* information set. The payoffs for the seller and the consumer refer to the transactions with a single consumer over two periods of the game

still unaware of which type she is, and has to decide whether to purchase or not based on the signal. In this case, however, the seller gets to set a price *after* it has observed the customer's type (chosen by Nature).

As before, we consider two cases. When $\pi v \geq c$, the profit maximizing strategy for the seller is to sell to both types at πv. In the second period, it will sell only to the customers who like the product. The profit for the seller will be $(\pi v - c) + \pi(v - c)$. Also in this case, even if the producer knows who will like the good and can offer different prices to the different types, it prefers to charge the same price to all and benefit from the consumer experimentation, as we show in the following proposition.

Proposition 2 *If $v > \pi v \geq c$, then it is a perfect Bayesian equilibrium for the seller to set the same price for both types, $p_1 = \pi v$ and to signal both types, with both types choosing to buy. This equilibrium is unique in pure strategies except in the knife-edge case where $\pi v = c$.*

Proof See "Appendix."

When $\pi v < c$, the seller will never charge the same price to both types. If it did so by setting $p_1 = c$, the expected utility for the customer would be negative and no transaction would take place. If it did so by setting $p_1 < c$, it would be making more losses than by selling only to the v type. In fact, the seller will choose to sell only to the v type and will set a price $p_{1\_v} = c$. We formalize these observations in the following fact.

Proposition 3 *Assume that $v > c$, but $\pi v < c$. Then in a perfect Bayesian equilibrium $p_1 = c$, the seller always signals only to the type that will like the good, and the buyer purchases on receipt of the signal.*

Proof See "Appendix."

This equilibrium mirrors the one that we have found for the uniform price case under the same conditions. Similarly to that case, the *ex ante* expected utility of the buyers is $\pi(v - c)$ and the expected profit for the seller is $\pi(v - c)$ (which is higher than its expected profit in the symmetric incomplete information case). In essence, when the costs are sufficiently high, the merchant has more to lose from recommending goods at a price above the consumer's ex post reservation price, than it has to gain from personalizing the prices. Also in this case, the results are not qualitatively affected when the number of periods increases from two to an infinite horizon.

To summarize, two results emerge from this analysis: First, the seller will tend to give honest recommendations when the cost of the good is high relative to its market evaluation. Alternatively, when the cost is relatively low, the seller is more likely to provide a dishonest recommendation. This is true whether the seller offers a constant price or a different price to different customers.

5 Multiple Goods

Now consider a market where a new product (e.g., a new CD) is introduced in every period. The seller must now decide whether to recommend it or not to its customers; each customer has to decide whether to purchase the good or not. If the customer is recommended a good that after experimentation she does not like, we assume that she will adopt a trigger strategy where the customer retaliates by discontinuing to purchase from the seller for an extended length of time. For simplicity, we assume that the unit cost c is the same for all products. In addition, only one new product is offered in the market in every period, and all of the new products have the same value v for a π share of customers and a 0 value for the remaining $1 - \pi$ customers. Again, we consider different levels of accuracy of the merchant's prediction of the customer's preferences.

5.1 Symmetric Incomplete Information

This trivial case is presented as a baseline scenario. In a $n = 2, \ldots, \infty$ period model with n goods (one in each period) with symmetric incomplete information, the seller can only sell in any period in which $\pi v > c$ at $p = \pi v$. In any period in which $\pi v < c$, the seller will not recommend the good.

5.2 Asymmetric Information

We assume now that the seller is using data mining or collaborative filtering tools that lead to an accurate prediction of a customer's preferences. We can visualize this game

Table 1 Multiple goods

Nature	Buyer/seller	Buy	Do not buy
Nature draws a v type	Offer, offer	$p-c, v-p$	0, 0
Nature draws a 0 type	Offer, offer	$p-c, -p$	0, 0
Nature draws a v type	Offer, not offer	$p-c, v-p$	0, 0
Nature draws a 0 type	Offer, not offer	0, 0	0, 0
Nature draws a v type	Not offer, offer	0, 0	0, 0
Nature draws a 0 type	Not offer, offer	$p-c, -p$	0, 0

as similar to the one in Fig. 1, where the tree of the game represents one period of the repeated game and is reiterated indefinitely following a purchase by a customer that likes the good. Note that payoffs must be altered for the seller, such that the $v-c$ margin coming from the second period sale can no longer be taken for granted. This is important in the repeated purchase scenario, since a new good is sold in each period of the game. This allows the consumer quickly to learn her preferences and whether or not to trust the seller. Now, the seller offers different goods and values repeat purchases, and so will have a greater incentive to gain and keep the trust of each consumer.

We represent the payoffs for the players in the one-shot game in the following normal form in Table 1, where the seller's strategies must be contingent to all possible Nature's moves.

We have studied a very similar game in the previous section (the only change is in the payoff of the seller). In the one-shot game, if Nature plays only a mixed strategy $(v, 0)$ with probabilities π and $1-\pi$, and $\pi v > c$, the sellers' best response will be to choose (Offer, Offer) at a price of πv, that the customer will accept. If $\pi v < c$, no transaction would take place (because we are considering a one-shot game, and the seller can no longer cut its losses from the first period in the second period).

Now, consider the 2 period case. By backward induction, we know that in the last period the seller will be tempted to recommend to all. Therefore it cannot price higher than $p = \pi v$. In the first period, can the seller sustain an honest strategy of recommending only to those who will like the good, at $p = v$? Even if the trigger strategy of scorned consumers is never to purchase again from that merchant, the answer is no. For the answer to be yes, we should have $\pi(v-c) + \pi v - c > v - c + \pi(\pi v - c)$ (where we sum the undiscounted profits from period 1 and 2), which is never satisfied. So, the seller can either earn $\pi(\pi v - c) + \pi v - c$ by offering only to selected consumers in the first period, or earn $\pi v - c + \pi(\pi v - c)$ by offering to all in the first period (which will cause $1-\pi$ consumers to abandon the seller in the second period). The undiscounted sum of the profits under the two strategies is identical. But under any positive discount rate, the seller will prefer to cash in early and will choose to offer to all at $p = \pi v$ in the first period. The seller will not provide an accurate recommendation in the finite horizon case when there is a positive discount rate.

However, in the infinite horizon case, it becomes possible to sustain a price of $p = v$ and recommend the good only to the customers selected by the filtering tool. Under certain parameter values (especially when π is large), the seller has no incentive to deviate because the customer would buy at every period if she is always recommending the right good. Accurate information may be provided in the infinite horizon case.

As usual, with infinitely repeated games, many equilibria are actually possible—but we are interested in one in particular. Consider the following seller's strategy: Sell only to the customers who will like the product. Since there are π such customers in every period, if $\delta = \frac{1}{1+r}$ is the discount rate, the profit from this strategy would be $\pi\,(p-c) + \frac{\delta}{1-\delta}\pi\,(p-c)$. What values of p will be possible in equilibrium? The following proposition answers this question:

Proposition 4 *If* $\frac{3r\pi+3\pi-2\pi^2-1-r}{1+r-\pi}\,(v-c) > 0$ *then it is an equilibrium for the seller to set* $p = v$ *and offer the good only to the customers who will like it, and for those customers to accept it, as long as* $v > c$ *but regardless of whether* $\pi v \lesseqgtr c$.

Proof See "Appendix."

How can we interpret these results? When the relationship between buyer and seller can extend in the long-term, the comparison between the cost of producing the product and its expected market evaluation becomes less important (as long as the condition in Eq. 5 holds). With repeated entry of experience goods, the seller has to gain and keep the trust of the customer during each single good—which makes him, in a sense, more likely to be truthful. On the other hand, if Eq. 5 is not satisfied, the seller might again offer dishonest recommendations, and the market might even end up with no transactions taking place at all.

6 Conclusion

Recommendation systems and behavioral targeting rely on similar principles. They use data mining and/or collaborative filtering to compare a consumer's traits and behaviors with the traits and behaviors of her peers. Of social concern is that sellers may use these data to make dishonest recommendations that benefit the seller at the expense of the customer.

To investigate the conditions under which seller's recommendations will tend or not tend to be truthful, we have constructed a simple model of seller-buyer interaction where the seller has more information than does the buyer about the buyer's tastes relative to the attributes of the product, based on collaborative filtering of consumer data. The model captures the essence of how merchants can use consumer peer-based recommendation systems and peers evaluations to induce customers to try new goods. We find that there are two types of equilibria, depending on parameter values of the product: For some parameter values, sellers recommend to all potential buyers, so recommendations have no value; but for other values, recommendations will be accurate. In particular, the incentive for the seller to provide accurate recommendations will be inversely related to the difference between the cost of producing the good and its average market evaluation.

Acknowledgments The author is grateful to Hal Varian for the inspiration for the model, and to Veronica Marotta as well as the editors and anonymous reviewers for extremely helpful comments. This research is partly supported by the Singapore National Research Foundation under its International Research Centre Singapore Funding Initiative and administered by the IDM Programme Office.

7 Appendix

Proof of Proposition 1 Let π_s be the posterior belief of the buyer about its type upon receiving the signal. There are three cases to consider:

Case 1 $\pi_s v > c$. The seller must set the price $p_1 \leq \pi_s v$ in order to induce purchase. If $p_1 > c$, then the seller would profit from signaling to both types since

$$\pi_s(p_1 - c) < p_1 - c.$$

But then the signal is uninformative, so $\pi_s = \pi$. Since $\pi v < c$, this contradicts the premise of this case. If $p_1 \leq c$, it will be profitable to signal only type v. The seller wants to choose the largest such price, so it picks $p_1 = c$. This is consistent with any π_s such that $\pi_s v > c$. The seller would rationally want to signal all type v consumers, since this would maximize second-period sales. Hence a rational buyer must believe $\pi_s = 1$.

Case 2 $\pi_s v < c$. Note that this includes the case $\pi_s = \pi$, in which the buyer regards the seller's signal as uninformative. In order to induce purchase, the seller must set $p_1 \leq \pi_s v$. Since $p_1 \leq \pi_s v < c$,

$$\pi_s(p_1 - c) > p_1 - c,$$

so it is not in the seller's interest to sell to both types. In this case the signal is completely informative, so the buyer should revise its posterior probability to $\pi_s = 1$. Since $v > c$, this contradicts the premise. Hence this case cannot arise.

Case 3 $\pi_s v = c$. Consider the strategy where the seller sets a price $p_1 \leq \pi_s v = c$ and signals both types. The signal would then be uninformative, so $\pi_s = \pi$ which contradicts $\pi_s v = c$ since $\pi v < c$. On the other hand, if the seller only signals type v, then the signal is informative so $\pi_s = 1$, which implies $v = c$, which is a contradiction. □

Proof of Proposition 2 The profit maximizing strategy for the seller is to sell to both types at the highest price that they are willing to accept. This price is πv for both types. If the seller's strategy was to offer the good at a higher price $p_1 > \pi v$ to all, then the expected utility of the buyer would be negative, and no purchase would result. If the seller's strategy was to offer the good only to the v type at a price $p_{1\_v} = v > \pi v$, then every buyer upon receiving the signal would accept, in which case the best response of the seller would be to offer that price to all types, thereby making the signal uninformative and the expected utility of the buyer negative once again. Hence, no price above πv is sustainable in equilibrium. On the other hand, no price $p_1 < \pi v$ and no combination of prices p_{1_0}, p_{1_v} (with $p_{1\_0} \neq p_{1\_v}$) for the v and 0 types are profit maximizing for the seller. The linear combinations $p_{1\_0} = \frac{\pi}{1-\pi} e$, $p_{1\_v} = v - e$

(where $e = [0, v]$ and $v - e \neq \frac{\pi}{1-\pi}e$) give both types nonnegative expected utilities but guarantee a lower profit for the seller than selling at a common price $p_1 = \pi v$. In the knife-edge case where $\pi v = c$, as before the seller makes zero profit so it is again an equilibrium to signal only type v, and for type v to buy. $\square$

Proof of Proposition 3 We have already shown that the seller will never charge the same price to both types when $\pi v < c$. The seller will neither sell to both types at different prices. The following linear combinations of prices, $p_{1\_0} = \frac{\pi}{1-\pi}e$, $p_{1\_v} = v - e$ (where $e = [0, v]$ and $v - e \neq \frac{\pi}{1-\pi}e$, and the seller practically gives away the good to one type of customer), give the buyer a nonnegative expected utility, but guarantee the seller only a profit of $2\pi v - \pi c - c$, which is less than selling only to the v type at $p_1 = c$. Only the strategy of selling to the v type at $p_{1\_v} = c$ satisfies Bayes's Law. Let π_s be the posterior belief of the buyer about its type upon receipt of the signal. When $\pi_s v > c$, if the seller sets $p_1 > c$, then it could sell to both types, thereby making the signal uninformative and $\pi_s = \pi$, which contradicts $\pi v < c$. When $\pi_s v > c$ and the seller sets $p_1 = c$ and sells to all, the same contradiction arises. Only if the seller sets $p_1 = c$ and sells to the v type and $\pi_s v > c$, then the signal is informative and $\pi_s = 1$, which satisfies $\pi_s v > c$. $\square$

Proof of Proposition 4 Imagine that p is set to v. To understand whether the customers can trust that the seller will only recommend a good sold at this price to those who will like it, we must compare the profit from giving good recommendations, to the profit from recommending the good to all. In the first case the present discounted value of the profit, over an infinite horizon, is $\pi\,(p - c) + \frac{\delta}{1-\delta}\pi\,(p - c)$. In the second case, if the customer's trigger strategy is never to buy again from a merchant that recommended a bad product to him, and the supply of customers is finite, the seller's profits is $(p - c) + \frac{\theta}{1-\theta}\,(p - c)$, where $\theta = \frac{\pi}{1+r}$ because at each period $1 - \pi$ customers will abandon the seller after discovering that they have been recommended a good that they do not like. Hence, if $\pi\,(p - c) + \frac{\delta}{1-\delta}\pi\,(p - c) > (p - c) + \frac{\theta}{1-\theta}\,(p - c)$, the honest strategy for the seller is a best response to the customer choosing to purchase. We can rewrite the inequality as:

$$\frac{3r\pi + 3\pi - 2\pi^2 - 1 - r}{1 + r - \pi}\,(v - c) > 0 \tag{5}$$

by setting $\frac{\delta}{1-\delta} = \frac{1}{r}$ and $\frac{\theta}{1-\theta} = \frac{\pi}{1+r-\pi}$ and $p = v$. Given that it is assumed that $v > c$, and $(1 + r) > \pi$ (see denominator in Eq. 5), then this condition is satisfied simply as long as $3r\pi + 3\pi - 2\pi^2 - 1 - r > 0$. If this is the case, then the honest strategy is sustainable *regardless* of whether $\pi v \lesseqgtr c$ (but as long as $v > c$). $\square$

References

Allen, F. (1984). Reputation and product quality. *RAND Journal of Economics*, *15*(3), 311–327.

Ansari, A., Essegaier, S., & Kohli, R. (2000). Internet recommendation systems. *Journal of Marketing Research*, *37*(3), 363–376.

Avery, C., Resnick, P., & Zeckhauser, R. (1999). The market for evaluations. *American Economic Review*, *89*(3), 564–584.

Beales, H. (2010). The value of behavioral targeting. *Network Advertising Initiative*. http://www.networkadvertising.org/pdfs/Beales_NAI_Study.pdf.
Benabou, R., & Laroque, G. (1992). Using privileged information to manipulate markets: Insiders, gurus, and credibility. *Quarterly Journal of Economics*, *107*(3), 921–958.
Bennett, J., & Lanning, S. (2007). The Netflix prize. In *Proceedings of KDD Cup and Workshop*.
Blume, A. (1998). Contract renegotiation with time-varying valuations. *Journal of Economics and Management Strategy*, *7*(3), 397–433.
Che, Y.-K. (1996). Customer return policies for experience goods. *Journal of Industrial Economics*, *44*(1), 17–24.
Crawford, V., & Sobel, J. (1982). Strategic information transmission. *Econometrica*, *50*(6), 1431–1452.
Dellarocas, C. (2003). *The impact of online opinion forums on competition and marketing strategies*. Mimeo, MIT Sloan School of Management.
Ely, J., & Valimaki, J. (2003). Bad reputation. *Quarterly Journal of Economics*, *118*(3), 785–814.
Fayyad, U., Piatetsky-Shapiro, G., & Smyth, P. (1996). From data mining to knowledge discovery in databases. *AI Magazine*, *17*(3), 37.
Goldfarb, A., & Tucker, C. (2011). Privacy regulation and online advertising. *Management Science*, *57*(1), 57–71.
Kohavi, R., & Provost, F. (2001). Applications of data mining to electronic commerce. In *Data Mining and Knowledge Discovery* (pp. 1–7). Boston, MA: Kluwer.
Krishna, A., & Zhang, Z. J. (1999). Short or long-duration coupons: The effect of the expiration date on the profitability of coupon promotions. *Management Science*, *45*(8), 1041–1056.
Krishna, V., & Morgan, J. (2001). A model of expertise. *Quarterly Journal of Economics*, *116*(2), 747–775.
Lewis, T. R., & Sappington, D. E. M. (1994). Supplying information to facilitate price discrimination. *International Economic Review*, *35*(2), 309–327.
Liebeskind, J., & Rumelt, R. P. (1989). Markets for experience goods with performance uncertainty. *RAND Journal of Economics*, *20*(4), 601–621.
Maskin, E., & Tirole, J. (1990). The principal-agent relationship with an informed principal, I: The case of private values. *Econometrica*, *58*(2), 379–409.
Maskin, E., & Tirole, J. (1992). The principal-agent relationship with an informed principal, II: Common values. *Econometrica*, *60*(1), 1–42.
McCulloch, R. E., Rossi, P. E., & Allenby, G. M. (1996). The value of purchase history data in target marketing. *Marketing Science*, *15*(4), 321–340.
McFadden, D. L., & Train, K. E. (1996). Consumers' evaluation of new products: Learning from self and from others. *Journal of Political Economy*, *104*(4), 683–703.
Miller, N., Resnick, P., & Zeckhauser, R. (2005). Eliciting informative feedback: The peer-prediction method. *Management Science*, *51*, 1359–1373.
Moorthy, S., & Srinivasan, K. (1995). Signaling quality with a money-back guarantee: The role of transactions costs. *Marketing Science*, *14*(4), 442–466.
Morgan, J., & Stocken, P. (2003). An analysis of stock recommendations. *RAND Journal of Economics*, *34*(1), 183–203.
Morris, S. (2001). An instrumental theory of political correctness. *Journal of Political Economy*, *109*(2), 231–265.
Ottaviani, M., & Sørensen, P. (2006a). Reputational cheap talk. *Rand Journal of Economics*, *37*(1), 155–175.
Ottaviani, M., & Sørensen, P. (2006b). The strategy of professional forecasting. *Journal of Financial Economics*, *81*(2), 441–466.
Prendergast, C. (1993). A theory of "yes men". *American Economic Review*, *83*(4), 757–770.
Resnick, P., & Sami, R. (2007). The influence limiter: Provably manipulation-resistant recommender systems. In *Proceedings of the 2007 ACM Conference on Recommender Systems* (pp. 25–32). ACM.
Resnick, P., & Varian, H. (1997). Recommender systems. *Communications of the ACM*, *40*(3), 56–58.
Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., & Riedl, J. (1994). Grouplens: An open architecture for collaborative filtering of netnews. In *Proceedings of the 1994 ACM Conference on Computer supported cooperative work* (pp. 175–186). ACM.
Rossi, P. E., & Allenby, G. M. (1998). Marketing models of consumer heterogeneity. *Journal of Econometrics*, *89*(1–2), 57–78.
Schafer, J. B., Konstan, J., & Riedl, J. (1999). Recommender systems in e-commerce. In *Proceedings of the 1st ACM Conference on Electronic commerce* (pp. 158–166). ACM.

Schlee, E. E. (2001). Buyer experimentation and introductory pricing. *Journal of Economic Behavior and Organization*, *44*(3), 347–362.

Shaffer, G., & Zhang, Z. J. (1995). Competitive coupon targeting. *Marketing Science*, *14*(4), 395–416.

Shapiro, C. (1983). Premiums for high quality products as returns to reputation. *Quarterly Journal of Economics*, *98*(4), 659–679.

Sobel, J. (1985). A theory of credibility. *The Review of Economic Studies*, *52*(4), 557–573.

Victor, P., Cornelis, C., De Cock, M., & Pinheiro da Silva, P. (2009). Gradual trust and distrust in recommender systems. *Fuzzy Sets and Systems*, *160*(10), 1367–1382.

Yan, J., Liu, N., Wang, G., Zhang, W., Jiang, Y., & Chen, Z. (2009). How much can behavioral targeting help online advertising?. In *Proceedings of the 18th International Conference on World Wide Web* (pp. 261–270). ACM.

Zhang, Z. J., Krishna, A., & Dhar, S. (2000). The optimal choice of promotional vehicles: Front-loaded or rear-loaded incentives? *Management Science*, *46*(3), 348–362.