



Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Decision Support Systems xx (2004) xxx–xxx

Decision Support
Systems

www.elsevier.com/locate/dsw

Intelligent agents in electronic markets for information goods: customization, preference revelation and pricing

Ravi Aron^{a,*}, Arun Sundararajan^b, Siva Viswanathan^c

^a*The Wharton School, University of Pennsylvania, United States*

^b*Stern School of Business, New York University, United States*

^c*Robert H. Smith School of Business, University of Maryland, United States*

Abstract

Electronic commerce has enabled the use of intelligent agent technologies that can evaluate buyers, customize products, and price in real-time. Our model of an electronic market with customizable products analyzes the pricing, profitability and welfare implications of agent-based technologies that price dynamically based on product preference information revealed by consumers. We find that in making the trade-off between better prices and better customization, consumers invariably choose less-than-ideal products. Furthermore, this trade-off has a higher impact on buyers on the higher end of the market and causes a transfer of consumer surplus towards buyers with a lower willingness to pay. As buyers adjust their product choices in response to better demand agent technologies, seller revenues decrease since the gains from better buyer information are dominated by the lowering of the total value created from the transactions. We study the strategic and welfare implications of these findings, and discuss managerial and technology development guidelines.

© 2004 Published by Elsevier B.V.

Keywords: Ecommerce; Digital goods; Dynamic pricing; Automated pricing; Software agents; Price discrimination; First-degree price; Discrimination; Mass customization; Shopbot; Pricebot; Information filtering

1. Introduction and motivation

Electronic markets and web-based technology have made it possible for online sellers to obtain accurate information on individual buyers, and to provide products, services and prices tailored pre-

cisely to each consumer's individual preferences. This is typically achieved by the deployment and use of *intelligent agents* [20]. These software agents can estimate buyer preferences, valuations and product tastes by combining consumer purchase histories with individual and site demographics. Merchants can then use this information to customize, price and recommend products to these buyers, simultaneously increasing the fit of the product to the buyer, and the amount of surplus they extract from the buyer. The success of companies like Amazon.com that were

* Corresponding author.

E-mail addresses: raviaron@wharton.upenn.edu (R. Aron), asundara@stern.nyu.edu (A. Sundararajan), sviswana@rhsmith.umd.edu (S. Viswanathan).

early adopters of collaborative filtering agents like Firefly and NetPerceptions suggests huge potential gains for that the companies who lead the way in exploiting more advanced agent technology in electronic commerce.

In this paper, we study software agents that determine *pricing* based on customer preferences (commonly called dynamic pricing agents, demand agents or preference-based pricing agents). Although the technology is still fairly nascent, there are deployable products available—early software vendors for preference-based pricing agents include BlueMartini, Manugistics, ProfitLogic, Retek and Zilliant—and many companies are beginning to realize their potential. For instance, Ford plans to move towards pricing its automobile financing products dynamically, based on customer profiles and choices, and expects to cut its \$10 billion spending on blanket promotions significantly as a consequence; Baker et al. [2] also describe an online electronics components retailer which is using customer preference data related to demand immediacy to charge differentially on the products it sells; Ticketmaster has successfully experimented with differential ticket pricing based on customer specifications [7]. A number of financial services companies are using similar technology to price products and deliver service to their banking and credit card customers. NextCard uses a self-segmentation system that prices terms for credit cards based on a product specification by potential customers; Capital One uses profiles based on hundreds of variables to tailor products and prices for specific clients [23]. Other examples of financial services products using preference-based pricing and service include those offered by United California Bank and First Union Bank [32]. This form of pricing is more closely related to first-degree price discrimination, in contrast with quality or usage-based second-degree price discrimination, which has been studied extensively in the theoretical industrial organization literature [1,22,38].

As the revenue models of online retailing companies and web portals evolve, they are well positioned to use agent-based dynamic pricing. Much like financial services companies, their core product is a highly customizable information good (an online shopping experience customized to the buyer, infor-

mation customized to the viewer's browser) with variability in perceived value, a tangible cost to the buyer from lack of customization, and low or zero marginal cost of customization. Currently, leveraging preference information in these markets has taken on indirect forms like cross-selling, targeted advertising and selling segmented marketing information. Each of these strategies has the same objective—to extract as much surplus as possible from each consumer, based on information the consumer reveals to the seller while customizing a product. These companies may have to acquire a significant degree of market power and customer lock-in before they can successfully gain customer acceptance for direct pricing agent technologies—the mandate by Jupiter Research that loyal, quality-seeking customers are the ones who should be offered personalized dynamic pricing [7] supports this observation.

However, when a seller actively infers buyer preferences and valuations, the quality-seeking consumer faces a trade-off between giving up their personal information, and getting products or information customized better to their tastes. Trade-press attention has focused primarily on the issue of personal privacy. According to Weise [37], 'lying when Web sites ask for personal information is a common tactic to protect privacy. But on such sites, it destroys the quality of the recommendations you receive. Answer honestly, and these sites quickly learn a great deal about what you read, listen to and like to watch. Whether users will be willing to trade information about themselves for a more personal experience online remains to be seen'. These concerns about privacy have been heightened by the perceived misuse of information by companies like RealAudio and DoubleClick.

As preference-based pricing agents become more widespread, another crucial trade-off for the consumer—and the subject of our paper—is between the price paid, and the level of customization obtained. Put simply, the more a demand agent infers about one's ideal product, the more it will know about one's willingness-to-pay. Intuitively, it appears that this kind of agent technology is likely to help sellers extract more value from their buyers. However, it is possible that consumers may change their behavior and choices in a manner that counters these potential losses in consumer surplus. This trade-off is illustrated qual-

itatively in Fig. 1. The value of a product increases as the buyer customizes the product more, as illustrated by the value curve above the x -axis. However, since the seller is also able to infer more about the buyer's willingness to pay, the price paid by the buyer increases as well, as illustrated by the (negative) price curve below the x -axis.

As consumers begin to make these price–product trade-offs, the profitability and welfare implications of these inferencing technologies in an electronic market are not intuitively evident. We address these issues in this paper, by asking the following questions:

- In an electronic market for customizable information goods, how does the presence of intelligent agents that can infer buyer preferences affect product pricing and consumer choices?
- What are the relative benefits of intelligent agents to buyers and sellers, when consumers have heterogeneous valuations for products, and value product customization and quality differentially?
- Given these relative benefits to buyers and sellers, what characteristics of intelligent agents are likely to benefit sellers and consumers?

Our paper enhances the recent stream of work on price discrimination and differentiation enabled by IT, and in IT-enabled markets. Recent papers studying customization and price discrimination include Ulph and Vulcan [35], who examine a firm's incentives to offer mass customization in conjunction with first-degree price discrimination. They find that a firm which first-degree price discriminates is also better off if it mass-customizes. Interestingly, they also conclude that if mass customization strategies are chosen by two competing firms, then the profits of each firm are independent of the price discrimination strategies chosen by the other firm; this result indicates that our analysis of a monopolist's pricing strategies may generalize well to a competitive environment. Dewan et al. [14] analyze pricing strategies for firms that sell both a generic and a differentially customized product, and find that a firm which mass-customizes must raise the price of its generic product in order to protect margins in its custom-product market. In other work on price discrimination for information goods, Bakos and Brynjolfsson [3] study the use of bundling strategies by multiproduct monopolist as a price discrimination strategy. They find that when consum-

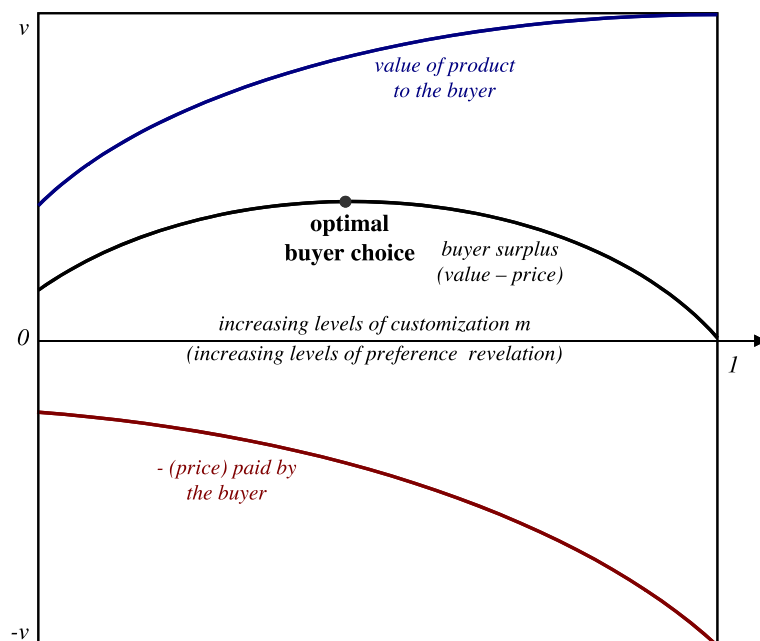


Fig. 1. The buyer's fundamental trade-off, when facing a demand agent.

ers differ systematically in their valuation for goods, then offering a menu of different bundles for each market segment increases the power of traditional price discrimination strategies. As in our study, their focus is on information goods with zero marginal costs; however, we model a market in which individualized pricing is feasible, making aggregation a suboptimal strategy.

IT-enabled quality differentiation has been studied by Nault and Dexter [27], Nault [26], Barua et al. [5] and Clemons and Kleindorfer [12]. Nault [26] considers the case where IT facilitates quality differentiation by enhancing the value of the good/commodity and enables the seller to charge a price premium. In contrast, consumers in our model have a choice of a continuum of qualities (level of customization), and our model derives optimal levels of quality for each customer (represented by the degree of customization chosen). While adoption of technology is costless to consumers in our model, Nault (1995) investigates the adoption of IOS, where consumers incur a fixed cost of adopting the technology and finds that although adoption of these technologies (such as IOS and EDI) benefit both the seller as well as the consumers, the seller may have to provide subsidies to overcome the barriers created by the cost of adoption. They also find that in the case of a duopoly, both firms as well as consumers are better off with a segmented market with only one firm adopting the value-enhancing technology.

The unique nature of differentiation in information goods is also highlighted by the study by Bhargava and Choudhury [6] who analyze the vertical differentiation of information goods. Contrary to the case of traditional physical goods analyzed by Shaked and Sutton [31] and Mussa and Rosen [25], they find that vertical differentiation is not an optimal strategy when the highest quality product has the best benefit-to-cost ratio; the monopolist is better off offering just the highest quality. In contrast, Mukhopadhyay et al. [24] study quality competition among Internet search engines, finding that zero marginal costs enable the survival of lower-quality products. Other work in the area of pricing and differentiation of information goods includes Varian [36], Jones and Mendelson [18] and Sundararajan [33], who demonstrates that transaction costs can lead to the simultaneous optimality

of both fixed-fee and usage-based pricing. While adding to this stream of research, our focus is different—rather than abstracting away the details of the interaction between the seller, buyer and agent, we explicitly model the optimal customization choice of the buyer, and the possible incentives to trade-off lower customization for better pricing as agent technology becomes more effective.

In related work on pricing network services based on user preferences, Konana et al. [19] present a nonparametric estimation technique that enables the estimation of consumer demand characteristics from their observed behavior. Their choice of delay cost parallels our choice of customization levels by consumers. Similarly, in the responsive pricing mechanism in MacKie-Mason et al. [21], price-sensitive users adjust their traffic inputs based on the price and how valuable the network service is to them, an outcome that is consistent with buyer choices of customization in our model. We contextualize work done on agent-based technologies [9,28] to a dynamic pricing setting, and provide a stronger economic framework for computational work on dynamic pricing agents [30].

The rest of the paper is organized as follows. In Section 2, we provide an overview of our model, and explain its parameters through a simple example. Section 3 characterizes buyer and seller actions and derives expressions for their respective surplus. Some preliminary implications of these results are also illustrated. Section 4 presents comparative statics, using general functional forms as far as possible and also using a specific set of functional forms to describe agent inferences. We use these results to describe the revenue and welfare implications of these agents, for different types of products, consumers and agent inference rates. Section 5 discusses the business implications of these results, and concludes the paper with a summary of our ongoing research. An extended mathematical appendix details our proofs, and analyzes the sensitivity of our results to altering three sets of key assumptions.

2. Model overview

Our market has six basic building blocks—the sellers, the buyers, the intelligent agent, the seller's

knowledge about the intelligent agent, the buyer's knowledge and method of choice, and the sequence of buyer–seller interaction. We list the elements and assumptions of each of these below. We then summarize our notation, and provide an example that illustrates the different variables in our model.

2.1. Seller

We model a monopolist seller in a market for a highly customizable information good, with zero marginal cost and zero cost of customization. The seller interacts independently with every potential buyer, and can set a price p independently for each buyer. The seller can customize the product over a continuum of *customization levels* $m \in [0,1]$, independently for each specific buyer.

2.2. Buyers

There are a set of heterogenous buyers for the customizable product. Each buyer wishes to buy one unit of the good. Each buyer has a (different) ideal product, for which he/she has a valuation v . Different buyers have different valuations v . The ideal product is the one that is perfectly customized for the buyer, or has a customization level $m=1$. If the product has a customization level $m<1$, then the buyer's valuation for the (suboptimally) customized product is $v-t(1-m)^2$, where t is referred to as the unit *cost of commoditization*. All buyers share a common unit cost of commoditization t , which is known to the seller. This is analogous to the concept of a transportation cost used in spatial models of product differentiation, and these assumptions, including the quadratic cost function and common unit transportation cost, are very commonly used in such models (for instance, Ref. [13], and the literature following that paper). The buyer is rational in his/her choice of m , and consequently chooses only those values of m , such that $v-t(1-m)^2 \geq 0$.

2.3. Intelligent demand agent

The seller uses an intelligent demand agent (IA), which makes imperfect inferences about the valuations v of these buyers, based on an analysis of the buyers' specifications of their desired customization.

If the buyer provides the seller with a set of product specifications that correspond to a customization level m , and the seller provides this information to the IA, then the IA provides the seller with an interval estimate of possible values $[\varepsilon, \varepsilon+\theta(m)]$ of the buyer's valuation v for that buyer's ideal product.¹ The 'window' $[\varepsilon, \varepsilon+\theta(m)]$ always contains v . This is illustrated in Fig. 2(a). ε can range anywhere from $\max[t(1-m)^2, v-\theta(m)]$ to v .

The width of the range $\theta(m)$ is determined by how accurate the IA is, and by the level of customization m chosen by the buyer. $\theta(m)$ is always *strictly* decreasing in m , since a higher level of customization implies a higher level of information about the customer's preferences, and consequently, a better estimate of v . In addition, we assume that $\theta(m)$ is convex, reflecting diminishing returns to this information. This is illustrated in Fig. 2(b). These assumptions are consistent with empirical results about the returns to information in intelligent agent learning (for instance, Ref. [28] or Ref. [16]). The 'level of curvature' of the function $\theta(m)$ is qualitatively proportional to how rapidly the agent learns.

We denote $\theta(1)=\theta_{\min}$, and $\theta(0)=\theta_{\max}$ as the narrowest (best) and widest (worst) widths of range for any buyer, corresponding, respectively, to the buyer choosing the highest and lowest possible levels of customization m .

2.4. Seller's knowledge about the IA

We assume that the seller cannot infer any additional strategic or distributional information about the buyer's valuation from this interval estimate $[\varepsilon, \varepsilon+\theta(m)]$. This is motivated by the fact that in reality, software agents of the kind we model make their inferences based on product choices, click-stream data and demographic data, rather than on strategic buyer representations. The algorithms commonly used by these agents are typically based on neural networks and fuzzy logic [15,17]; in addition, more recent descriptions of demand agents use Q-learning, which draws from

¹ Alternately, one could think of this as the IA providing the seller with a point estimate, and the point estimate being within $1/2\theta(m)$ of the buyer's true valuation v .




Fig. 2. (a) The 'sliding window' of the intelligent agent. (b) The rate of learning of the intelligent agent.

look-up tables with no underlying model of the environment [34]. The agents are effectively 'black boxes' that provide a seller with outputs after processing a series of inputs, and do not reveal the underlying functions that resulted in mapping the inputs to the output. Consequently, sellers can get an estimate of the buyer's valuation within a window, but cannot actually see or explicitly recalibrate the process that led to the creation of the window. This translates to the following restrictions in our economic model—that the seller does not use the information that a buyer with a specific v would find it optimal to choose a specific level of m , and that the buyer is aware that the seller does not make this kind of strategic inference.

Our assumption—modeling intelligent agents as nonstrategic 'interval choosers'—allows us to focus on our problem of interest—the changes in consumer behavior, the corresponding pricing and revenue implications, and the welfare effects of having agents that make imperfect inferences. It also implies that the seller's prior on v is that it is uniformly distributed² in

² Since the buyer will not choose a value of θ such that $v < \tau(\theta)$, it leads to the question of whether the agent can simply choose $\varepsilon \in [v - \theta, v]$, and then narrow its interval if $\varepsilon < \tau(\theta)$. We investigate this modification in part B of our extended appendix, and show that this modification does not change the optimal buyer choices predicted by Proposition 1 (and consequently, does not change any of our subsequent results).

$[\varepsilon, \varepsilon + \theta(m)]$. This choice of distribution corresponds to a noninformative prior, since the seller has no distributional information beyond the support of v .

2.5. Buyer choice and knowledge

The buyer is aware that if he/she chooses a customization level m , the intelligent agent will give the seller an interval estimate of width $\theta(m)$. Since $\theta(m)$ is assumed strictly monotonic, it is invertible. Therefore, we can treat the buyer's choice of m as being equivalent to the buyer choosing the width of the IA's interval $\theta(m)$. When we say that 'the buyer chooses an interval width $\hat{\theta}$ ', what we mean is that the buyer chooses the customization level m that will result in an interval width $\hat{\theta} = \theta(m)$.

For subsequent notational convenience, we define the function $\tau(x) = t[1 - \theta^{-1}(x)]^2$. If the buyer chooses an interval width $\theta(m)$, the value of the product to the buyer is $v - t[1 - m]^2 = v - \tau(\theta(m))$, and we refer to $\tau(\theta)$ as the corresponding cost of commoditization (consistent with our definition of unit cost of commoditization in Section 2.2).

The buyer does not know the actual interval $[\varepsilon, \varepsilon + \theta(m)]$ that the IA will provide the seller. We assume that the buyer does not make inferences about the intelligent agent's behavior from prior interaction. However, the buyer does know what possible sets of interval estimates the IA could give the seller.

Table 1
Sequence of events

Event	Buyer knows	Seller and IA know
Buyer is endowed with v	$v, \tau(\cdot)$	$\tau(\cdot)$
Buyer chooses θ	$v, \theta, \tau(\theta)$	$\tau(\cdot)$
Buyer reveals choice of θ to seller	$v, \theta, \tau(\theta)$	$\theta, \tau(\theta)$
Intelligent agent gives seller interval estimate $[\varepsilon, \varepsilon + \theta]$	$v, \theta, \tau(\theta)$	$\theta, \tau(\theta), [\varepsilon, \varepsilon + \theta]$
Seller chooses and reveals price p^*	$v, \theta, \tau(\theta), p^*$	$\theta, \tau(\theta), [\varepsilon, \varepsilon + \theta], p^*$

Consistent with our model of the IA, after the buyer chooses an interval width $\theta(m)$, the prior that the buyer has on ε is that it is uniformly distributed in $[\alpha, v]$ where $\alpha = \max[v - \theta(m), \tau(\theta(m))]$.

This prior on ε causes the buyer to form a corresponding prior on the seller's price; the price expected by the buyer³ is denoted $p(\theta)$. The distribution of $p(\theta)$ is derived in Lemma 2. In order to compute the prior on $p(\theta)$, the buyer 'solves' the seller's pricing problem for each value of $\varepsilon \in [\alpha, v]$, and estimates the price that she will face at that level of customization. The buyer then chooses the θ that maximizes expected consumer surplus $\psi(\theta)$, which is:

$$\psi(\theta) = E[\max(v - p(\theta) - \tau(\theta), 0)]$$

2.6. Sequence of buyer–seller interaction

Each buyer interacts independently with the seller. The sequence of interaction is as follows:

- (1) The buyer is endowed with an ideal-product valuation v .
- (2) A buyer chooses a set of product specifications, which correspond to a customization level m , and an interval width $\theta(m)$.
- (3) The buyer specifies these product specifications to the seller, thereby communicating m , and consequently $\theta(m)$ to the seller.
- (4) The seller's intelligent agent generates the interval estimate $[\varepsilon, \varepsilon + \theta(m)]$, which always contains the buyer's true valuation v .

- (5) The seller sets a price p^* for the product.
- (6) The buyer purchases the product if he/she gets nonzero surplus, i.e., if $v - \tau(1 - m)^2 \geq p^*$.

The sequence of information that each party (the buyer and seller) has at each decision stage is summarized in Table 1.

2.7. Illustration

To illustrate our model of customization and agent inference, consider the following example. The purpose of this example is to explain our model—the actual dynamics of customization and pricing at the company mentioned may be different.

eDisc is an Internet music service that allows individuals to create their own personalized compact discs. The service features over 250,000 songs from a wide variety of eras and musical genres. A well-structured site enables easy navigation, making it easy for customers to find music that fits their tastes exactly. A customer can customize a CD with the exact songs of his/her choice, complete with a unique title and cover art. eDisc charges different prices depending on the choices made by each customer. In addition, the site also offers generic collections of songs. The firm uses personalization agents to make inferences about its customers based on individual profiles and past buyer choices.

Each CD is, therefore, a highly customizable product. The *ideal product* ($m=1$) for each buyer is a CD containing exactly the songs the buyer wants. The generic product ($m=0$) is any CD with one of the generic collections of songs. In addition, a buyer could choose an intermediate level of customization m , by resorting to a combination of unique choices and standard compilations provided by the seller. For instance, a buyer wishing to compile a 70-min CD of sentimental country favorites, could handpick each song (one from each of his/her favorite artists) for his/her CD or choose a combination of a few tracks that he/she is particularly interested in and an assorted bundle of country/folk tracks recommended by the seller. In the latter case, the product is no longer the buyer's ideal product and a buyer opting for partial customization incurs a cost of commoditization. As mentioned earlier, the only way a buyer can receive

³ $p(\theta)$ is a random variable, not an expected value.

his/her ideal product is to individually select all the tracks that would constitute his/her CD.

Based on the choices made by these buyers, eDisc is able to obtain an interval estimate of their valuations. Consider a buyer who is choosing a CD of 10 songs from eDisc's selection of 1960's folk music tracks, and who values his/her ideal product at \$25. This would correspond to $v=\$25$. If this buyer chooses all the songs that he/she wants ($m=1$), eDisc's intelligent agent is able to infer her valuation within a margin of \$2 (or $\theta_{\min}=2$). A sample interval estimate here would be that the buyer's valuation for the buyer's ideal product is in $[\$23.50, \$25.50]$ (which would correspond to $\varepsilon=\$23.50$).

On the other hand, if this buyer were to purchase a precompiled assorted collection of 1960's folk songs ($m=0$), the buyer's valuation is \$15. This means that $v-t(1-0)^2=\$15$, or $t=\$10$. In addition, suppose that in this case, the agent, which has much less information, can only infer her valuation within a margin of \$5 (which means that $\theta_{\max}=5$). A sample estimate of the buyer's valuation v here is that v is in $[\$23, \$28]$. Remember that v is the valuation the buyer has for his/her *ideal* product. Consequently, the seller knows that the buyer will be willing to pay between \$13 and \$18 for the generic product.

If the buyer hand-picks, say, 5 out of 10 tracks, and chooses a precompiled set of 5 others (for simplicity, let's say that this corresponds to $m=5/10$), suppose the intelligent agent can estimate v within a margin of \$3 [which makes $\theta(5/10)=3$]. A sample interval estimate here would be $[\$23, \$26]$. The buyer's cost of commoditization for this partially customized product is $t(1-5/10)^2=\$2.50$, which means that $\tau(3)=\$2.50$ ($\theta=3$ is the width of the interval 'chosen' by the buyer through her choice of $m=5/10$).

Consequently, the seller knows that the buyer would be willing to pay between $(\$23-\$2.50)$ and $(\$26-\$2.50)$, or between \$20.50 and \$23.50 for this partially customized product. The buyer's true valuation for this partially customized product is, of course, $\$25-\$2.50=\$22.50$.

Note here that m refers to a level of customization, not a specific product. *Two buyers who get the same level of customization are not buying the*

same product—they are merely buying products which are at the same 'distance' from their ideal product. It is precisely this difference in buyers' ideal products that makes it possible for the intelligent agent to make inferences about the buyers' valuations, based on their product specifications. For instance, a buyer who chooses a rare classic from 1913, and another who chooses the latest single by today's teen music sensation could both be getting the same level of customization. However, their valuations v could be very different. In addition, a buyer who chooses a combination of (i) eight of her favorite classical tunes and (ii) an assorted bundle of rare classics provided by the seller, and another buyer who chooses a combination of (i) eight of her favorite jazz tunes and (ii) an assortment of 1980's jazz tunes, are modeled as choosing the same level of customization m .

We have chosen this approach to modeling customization for three reasons. First, while we have an inherent underlying model of product differentiation (and a corresponding mapping from different products to valuation intervals), treating product choice implicitly as a level-of-customization variable, instead of explicitly as a vector in a product space, enables us to capture the aspect of customization pertinent to our research questions, without explicitly dealing with a complex product space. Secondly, the focus of our model is not on product differentiation—it is on the pricing, consumer choice and welfare implications of inferences made by an intelligent agent about buyers' underlying valuations, given that buyers who have certain attribute preferences (and therefore a preference for particular instances of the differentiated product) have valuations that are associated with their product attribute preferences. Finally, using a continuum of values for customization in one dimension enables us to use continuous optimization techniques.

3. Analysis: preference revelation and pricing

In this section, we characterize the optimal customization choice of the buyer, optimal pricing by the seller, and the corresponding buyer surplus and seller profit functions.

3.1. Optimal seller pricing

The solution to the seller’s pricing problem is presented in Lemma 1. The mathematical proofs of all our results are in an extended mathematical appendix, available online as indicated in the references, or upon request from the authors.

Lemma 1. *If the estimate about the buyer’s valuation that is provided by the seller (demand) agent is that the buyer’s true valuation v is uniformly distributed in $[\varepsilon, \varepsilon + \theta]$, then the optimal price p^* for the seller is:*

$$(i) p^* = \frac{\varepsilon + \theta - \tau(\theta)}{2} \text{ if } \varepsilon \leq \theta + \tau(\theta)$$

$$(ii) p^* = \varepsilon - \tau(\theta) \text{ if } \varepsilon > \theta + \tau(\theta)$$

The intuition behind this pricing scheme is illustrated in Fig. 3. At a low enough interval width θ from the intelligent agent, the price is set

to capture the sale, irrespective of the buyer’s true valuation—at this level of information accuracy, the gains from closing the sale with probability 1 outweigh the losses from a lower price. As the level of error increases, the seller has to trade-off losing the customer with lower profits from pricing at a level that will ensure that the sale takes place. The pricing problem becomes identical to that of a monopolist facing a set of customers whose composite demand function is linear and downward sloping, and the solution—pricing at half the upper bound on the distribution of valuation—is the familiar optimal monopoly price.

3.2. Optimal buyer choice

At specific values of ε and θ , the buyer knows that the seller sets prices according to the price schedule in Lemma 1. Given the buyer’s prior on ε , the price distribution the buyer expects for a specific choice of θ is derived in Lemma 2.




Fig. 3. (a) Optimal seller pricing when $\varepsilon < \theta + \tau(\theta)$. (b) Optimal seller pricing when $\varepsilon > \theta + \tau(\theta)$.

Lemma 2. *The buyer's prior distribution over the price p set by the seller, at a level of customization $m(\theta)$ that induces an error θ , has a density function $f(p)$ which is as follows:*

- (a) *If $\theta \leq \frac{v-\tau(\theta)}{2}$: p is uniformly distributed in $[v-\tau(\theta)-\theta, v-\tau(\theta)]$, and $f(p)=1/\theta$ in this interval.*
- (b) *If $\frac{v-\tau(\theta)}{2} \leq \theta \leq v-\tau(\theta)$: $f(p)$ has support $[\frac{v-\tau(\theta)}{2}, v-\tau(\theta)]$, and :*
 - $f(p) = \frac{2}{\theta}$ for $p \in [\frac{v-\tau(\theta)}{2}, \theta]$
 - $f(p) = 1/\theta$ for $p \in [\theta, v-\tau(\theta)]$
- (c) *If $\theta \geq v-\tau(\theta)$: p is uniformly distributed in $[\frac{\theta}{2}, \frac{v+\theta-\tau(\theta)}{2}]$, and $f(p) = \frac{2}{v-\tau(\theta)}$ in this interval.*

Fig. 4 illustrates the result of this lemma. In case (a), the buyer's prior on ε is such that ε is always in the range of values for which it is optimal for the seller to charge the price of Lemma 1 (ii). Hence, the prior of the buyer on price is simply the prior on ε shifted to the left by $\tau(\theta)$. In case (c), the same logic applies, but for the price of Lemma 1 (i). Since this

admits a lower range of prices for the same range of ε , the buyer's prior on price has a narrower support, and more density on each point of this support. In case (b), either price from Lemma 1 is possible, depending on the value of ε . Note that $\theta+\tau(\theta)$ is an increasing function of θ so long as $\tau(\theta)$ is nondecreasing, and therefore, these successive intervals correspond to increasing values of θ . This is formally established in the proof of Lemma 3.

3.3. Buyer surplus

The buyer's decision problem is to choose the level of $m(\theta)$ that maximizes his/her surplus $\psi(\theta)$. At any price $p(\theta)$, the buyer's expected surplus $\psi(\theta)$ is the expected value of $\max[v-p(\theta)-\tau(\theta), 0]$, with the expectation taken over the buyer's distribution over price $p(\theta)$. The functional form of the buyer's expected surplus is derived in Lemma 3, and the buyer's optimal choice of product customization is characterized in Proposition 1.

Lemma 3. *If the buyer chooses a level of customization $m(\theta)$, with a corresponding interval width θ , then the expected surplus $\psi(\theta)$ of the buyer is:*

- (a) *If $\theta \leq \frac{v-\tau(\theta)}{2}$, then $\psi(\theta) = \psi_1(\theta) = \frac{\theta}{2}$;*

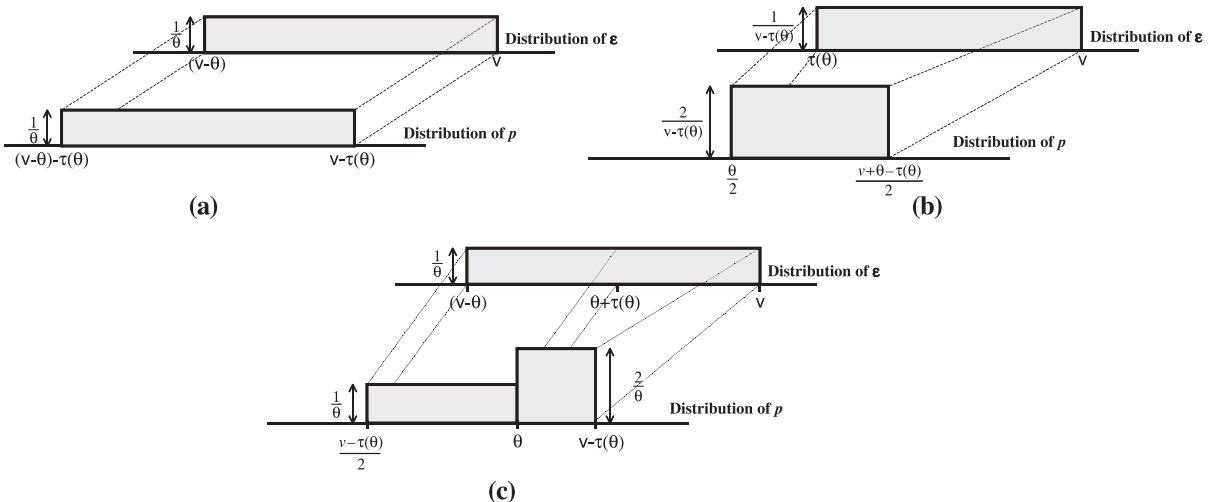


Fig. 4. (a) For values of $\theta \leq \frac{v-\tau(\theta)}{2}$, yielding $p = \varepsilon - \tau(\theta)$. (b) For values $\theta \geq v - \tau(\theta)$, yielding $p = \frac{\varepsilon + \theta - \tau(\theta)}{2}$. (c) For values of $\frac{v-\tau(\theta)}{2} \leq \theta \leq v - \tau(\theta)$, yielding: $p = \varepsilon - \tau(\theta)$ for $\varepsilon \leq \theta + \tau(\theta)$, $p = \frac{\varepsilon + \theta - \tau(\theta)}{2}$ for $\varepsilon \geq \theta + \tau(\theta)$.

- (b) If $\frac{v - \tau(\theta)}{2} \leq \theta \leq v - \tau(\theta)$, then $\psi(\theta) = \psi_2(\theta) = v - \tau(\theta) - \frac{\theta}{2} - \frac{[v - \tau(\theta)]^2}{4\theta}$;
- (c) If $v - \tau(\theta) \leq \theta \leq 2[v - \tau(\theta)]$ then $\psi(\theta) = \psi_3(\theta) = \frac{(2[v - \tau(\theta)] - \theta)^2}{4[v - \tau(\theta)]}$, and
- (d) If $\theta \geq 2[v - \tau(\theta)]$, then $\psi(\theta) = 0$.

Part (a) of Lemma 3 indicates that for high values of m (or low values of θ), decreasing customization (or increasing θ) actually benefits the buyer, since the buyer’s surplus increases linearly in θ . Intuitively, this is a consequence of the fact that at very high levels of customization, the seller ‘knows too much’ about the buyer’s valuation to make near-perfect customization worthwhile. This lemma indicates the nature of the central trade-off of the model—between the benefits (better customization) and costs (higher price) of preference revelation.

Proposition 1 examines what the optimal choice of θ will be. This is a crucial proposition of the paper, since it proves that buyers will almost always choose a less than ideal product (rather than opting for the ideal product), or even a generic product. In fact, it shows that buyers *never* choose their ideal products.

Proposition 1. *The buyer chooses a level of customization $m(\theta^*)$ which induces an interval width θ^* such that:*

- (a) If $\theta_{\max} \leq \frac{v - \tau(\theta_{\max})}{2}$, implying that $v \geq 2\theta_{\max} + \tau(\theta_{\max})$, then $\theta^* = \theta_{\max}$.
- (b) If $\theta_{\max} > \frac{v - \tau(\theta_{\max})}{2}$, implying that $v \leq 2\theta_{\max} + \tau(\theta_{\max})$, then θ^* is the solution to the following optimization problem:

$$\max_{\theta \in \mathfrak{R}^+} \psi_2(\theta) = v - \tau(\theta) - \frac{\theta}{2} - \frac{[v - \tau(\theta)]^2}{4\theta},$$

subject to:

$$v - \tau(\theta) \leq 2\theta,$$

$$v - \tau(\theta)\theta.$$

The solution θ^* is always interior. Furthermore, if the technology $\theta(m)$ of the seller’s intelligent agent is such that $\tau(\theta)$ is convex (i.e., if $\tau''(\theta) \geq 0$), then the unique optimal choice θ^* of the seller solves $\psi'_2(\theta) = 0$.

The proof of Proposition 1 defines the following two variables θ_1 and θ_2 :

$$\theta_1 \text{ solves } v - \tau(\theta) = 2\theta$$

$$\theta_2 \text{ solves } v - \tau(\theta) = \theta.$$

The results of Proposition 1 are illustrated in Fig. 5. At low values of θ (i.e., values of θ less than θ_1) which correspond to high values of customization m , the price set by the intelligent agent is very close to the buyer’s true valuation, since the margin of error θ in the agent’s estimate is fairly low. Intuitively, when the buyer decreases his/her level of customization, she gets a better price, but also gets a worse product. The price effect dominates in this region—at high levels of customization, the gains from a better price strictly outweigh the losses from a less suitable product. This causes the buyer to steadily increase the interval width of the agent. It also confirms that the buyer never chooses his/her ideal product, however poor the agent’s inferences are.

For high values of θ (i.e., values of θ greater than θ_2), the buyer’s surplus is strictly decreasing as the width of the agent’s interval increases, or as the level of customization decreases. For the seller, a higher value of θ corresponds to a less precise estimate from the intelligent agent, and a higher level of uncertainty about the buyer’s true valuation. This favors the buyer, since the seller is more likely to price lower than the buyer’s true valuation. However, higher uncertainty also makes the seller more likely to price too high. The buyer is shut out

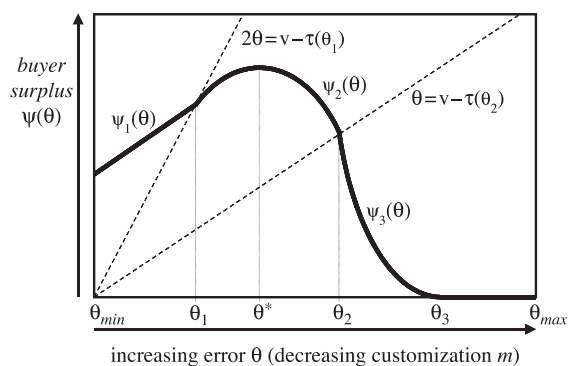


Fig. 5. Buyer surplus as a function of θ .

of the market in these cases, and gets no surplus. This reduces the desirability of higher values of θ —coupled with the fact that there are steadily increasing consumer surplus losses from a less customized product; this results in surplus strictly decreasing in this region.

The buyer is therefore pushed towards the middle. Facing the trade-off between better prices and better products, Proposition 1 shows that the buyer always finds herself in the region $[\theta_1, \theta_2]$.

3.4. Seller profits

We now solve for the seller’s expected profit when selling to a buyer with valuation v .

Proposition 2. *If the buyer’s optimal choice is θ^* , the seller’s expected profits $\pi(\theta^*|v)$ from a buyer of valuation v , are given by:*

(a) *If $v \leq 2\theta_{max} + \tau(\theta_{max})$, then*

$$\pi(\theta^*|v) = \frac{\theta^*}{2} + \frac{(v - \tau(\theta^*))^2}{4\theta^*}.$$

(b) *If $v \geq 2\theta_{max} + \tau(\theta_{max})$, then*

$$\pi(\theta^*|v) = \frac{2v + \theta^* - 2\tau(\theta^*)}{4}.$$

Apart from the last part of Proposition 1(b), which requires the convexity of $\tau(\theta)$, all of the preceding results are valid for any specification of $\theta(\cdot)$ that is decreasing and convex, and for any appropriate specification of $\tau(\theta)$. Further analysis of the buyer’s choices are needed to make stronger statements about the revenue and welfare effects of these intelligent agents require us to choose concrete functional forms for $\theta(\cdot)$, which we do in Section 4.

4. Results and comparative statics

This section focuses on a specific family of agent technologies, characterized by the inference function $\theta(m)=(1-m)^k$, for a range of values of $k \geq 1$. These satisfy the properties of $\theta(\cdot)$ described in Section 2. They also all yield

the same range of values of $\theta(\theta_{min}=0, \theta_{max}=1)$, allowing a comparison of buyer choices and seller profits across different values of k , a measure of the effectiveness of the technology of the agent, or the rate at which it infers. In Fig. 2(b), the more ‘convex’ curves correspond to higher values of k . In addition, this family of functions yields $\tau(\theta)=t\theta^{2/k}$ —an analytically appealing form.

We first prove one more set of results that are independent of the functional form of $\theta(m)$:

Proposition 3. *If the agent technology is such that $\tau(\theta)$ and $\tau'(\theta)$ are increasing in the parameter t , and the buyer’s choice of customization is such that $\theta^* < \theta_{max}$, then:*

- (a) *At the optimal level of customization $m(\theta^*)$, the buyer surplus $\psi(\theta^*)$ is increasing in v and decreasing in t .*
- (b) *The optimal level of customization $m(\theta^*)$ is decreasing in v and increasing in t .*
- (c) *The ‘marginal rate of substitution’ between v and t , which is $-\frac{d\psi(\theta^*)}{dt} / \frac{d\psi(\theta^*)}{dv}$, is equal to $\frac{d\tau(\theta^*)}{dt}$.*

These results are discussed shortly. We now focus on the specific family of agent technologies characterized by the inference function $\theta(m)=(1-m)^k$. In the figures that follow, we vary k between 1 and 3, using the closed-form solutions of $k=2$ from Proposition 4 as the benchmark midpoint:

Proposition 4. *If $\theta(m)=(1-m)^2$, then for a given level of v and t :*

- (a) *For a buyer with valuation v and unit commoditization cost t , the optimal level of customization $m(\theta^*)$ chosen by the buyer induces an interval width $\theta^*(v,t)$ such that:*

$$\theta^*(v,t) = \frac{v}{\sqrt{2+4t+t^2}}.$$

- (b) *The resulting consumer surplus is*

$$\psi^*(v,t) = \frac{(2+t - \sqrt{2+t(4+t)})v}{2}.$$

(c) The seller's expected profits are

$$\pi^*(v, t) = \frac{v[(2 + t(2 + t)) - t\sqrt{2 + t(4 + t)}]}{2\sqrt{2 + t(4 + t)}}$$

and are increasing in v , and decreasing in t .

Propositions 3 and 4 demonstrate that surplus and profits are higher with higher valuation customers, and lower for customers with higher costs of commoditization. Proposition 3 also establishes that as the value a buyer places on their ideal product increases, the final product chosen is further and further away from that ideal product.

Fig. 6 further illustrates⁴ the results of Propositions 3 and 4 (for a value of $k=2$). Fig. 6(a) summarizes the shape of the surplus function $\psi(v, t)$ in its two arguments v and t , confirming that ψ is increasing in v and decreasing in t . Since the curves are progressively closer together along the t -axis as v increases, the figure also indicates that a unit increase in the cost of commoditization t has a more negative marginal surplus impact on a higher valuation customer. In other words, high valuation customers are more adversely affected by increases in the cost of commoditization; an observation confirmed by the fact that the cross-partial of ψ with respect to v and t is negative.

Proposition 4 shows that the profits of the seller and the surplus of the buyer move in similar directions, when v and t change. Fig. 6(b) strengthens these observations, by indicating that the shape of ψ and π are very similar. The seller's profits are more sensitive to the cost of commoditization at higher values of buyer valuation and at higher revenue values. Profits are jointly (weakly) convex in v and t in the region plotted, and are a little over twice the buyer surplus.

Fig. 6(c) illustrates how optimal customization levels chosen by the buyers vary with buyer valuation and cost of commoditization. This figure illustrates the trade-off between withholding personal information in order to get a better price, and revealing this information in order to get a better product. Clearly, as v

decreases and t increases, the optimal level of customization increases. This indicates that as the ideal product valuation v increases, at a constant cost of commoditization t , the buyer gains more at the margin from withholding information (from a better price) than he/she loses (due to a less customized product); hence, the choices of product become increasingly commoditized for higher valuation buyers.

This observation leads us to an interesting conclusion—that *high* valuation buyers seem to incur a *greater* loss of surplus than lower valuation buyers, which is in contrast with standard results from models of second-degree price discrimination. This effect is partially a consequence of the fact that t is constant across buyers. However, Proposition 3(c) indicates that at a constant level of consumer surplus, the level of increase in v required to compensate for a unit increase in t is $\frac{d\tau(\theta^*)}{dt}$. This departure from standard results will therefore sustain even when t increases in v , so long as its rate of increase is less than $\frac{d\tau(\theta^*)}{dt}$.

We now investigate the effects of varying the inference rate of the intelligent agent.

Proposition 5. *If $\theta(m)=(1-m)^k$, an increase in the inference rate k of the agent causes a strict reduction in buyer surplus.*

Differentiating the expression for seller profits from Proposition 2(a), yields the following expression for the sensitivity of π^* to variation in k

$$\frac{d\pi^*}{dk} = -\frac{\partial\theta^*}{\partial k} \left(\frac{(v - \tau(\theta^*))^2}{4(\theta^*)^2} - \frac{1}{2} \right) - \left(\frac{\partial\tau}{\partial k} + \frac{\partial\theta^*}{\partial k} \tau'(\theta^*) \right) \frac{v - \tau(\theta^*)}{2\theta^*} \quad (A)$$

for $v \leq 2\theta_{\max} + t$.

This decomposition(A) has an interesting interpretation. Using the first-order conditions to the consumer surplus maximization problem, one can show that $\frac{(v - \tau(\theta^*))^2}{4(\theta^*)^2} > \frac{1}{2}$, and that $\frac{v - \tau(\theta^*)}{2\theta^*} > 0$. Note that $\frac{\partial\tau}{\partial k} + \frac{\partial\theta^*}{\partial k} \tau'(\theta^*)$ is simply the total derivative of $\tau(\theta^*)$ with respect to k when θ^* depends on k , and the functional form of τ also depends on k . This will be positive when an improvement in technology (increase in k) causes a net increase in misfit cost

⁴ The iso-function curves of Figs. 6–8 are a succinct way of depicting the rate of change and shape of a function of two variables. Iso-profit curves are commonly used in economic analysis; our iso-surplus and iso-product curves are loosely analogous to indifference curves. The curves are plotted by projecting the iso- (v, t) points (i.e., points on the function's surface which have equal values of the function) onto the (v, t) plane.

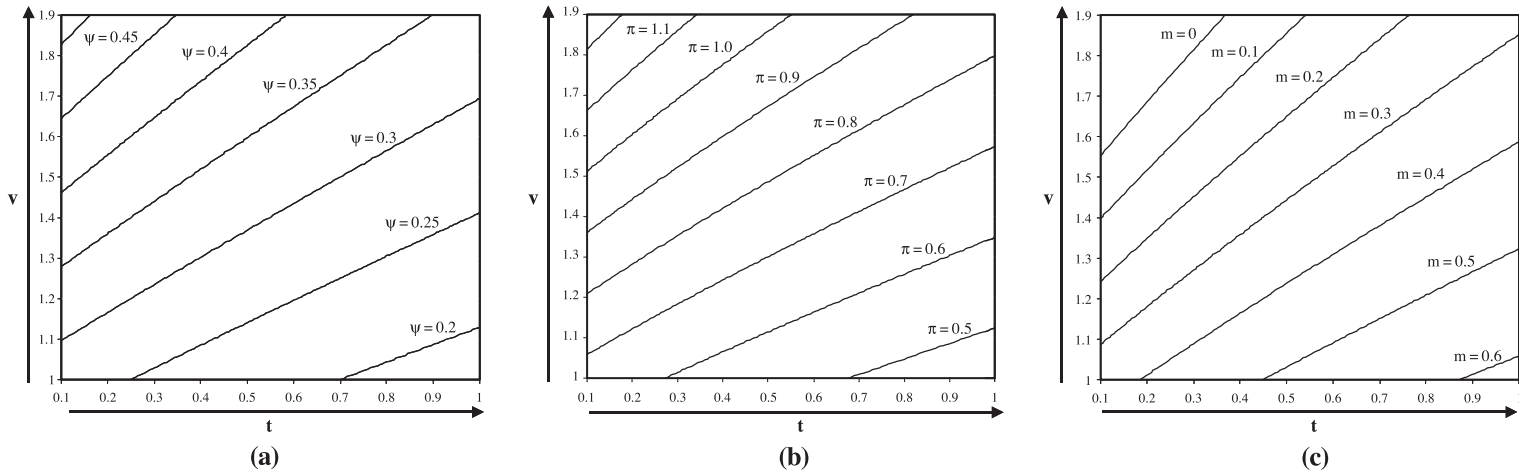


Fig. 6. (a) Iso-surplus curves for the buyer, as a function of v and t . (b) Iso-surplus curves for the seller, as a function of v and t . (c) Iso-product (constant levels of customization m) curves for the buyer, as a function of v and t .

$\tau(\theta^*)$ borne by the consumer, which happens as the consumer reduces customization in exchange for a better price. Consequently, the second half of the expression, $-\left(\frac{\partial \tau}{\partial k} + \frac{\partial \theta^*}{\partial k} \tau'(\theta^*)\right) \frac{v - \tau(\theta^*)}{2\theta^*}$, which represents the *change in profits from the loss in customer information*, is negative.

However, the term $\frac{\partial \theta^*}{\partial k}$, measuring the change in the net level of error from the increase in accuracy is likely to be negative, since it is unlikely that the withholding of information by the buyer is so drastic that the actual width of the IA's window reduces. The first half of the expression, $-\frac{\partial \theta^*}{\partial k} \left(\frac{(v - \tau(\theta^*))^2}{4(\theta^*)^2} - \frac{1}{2}\right)$, therefore measures the *change in profits from superior agent inference*, and is likely to be positive. Consequently, the seller's profits could either increase or decrease, depending on the relative magnitude of these expressions.

The other profit expression derived in Proposition 2(b) is for $v \geq 2\theta_{\max} + t$, and as shown in Proposition 1, the optimal choice of the buyer is constant at $\theta^* = \theta_{\max}$, implying that seller profits are unchanged by an increase in k .

We investigate the region $v \leq 2\theta_{\max} + t$ further by solving the optimization problem of Proposition 1 numerically for values of k varying between 1 and 3, with our analytical benchmark $k=2$ as the midpoint. The lower limit $k=1$ describes a situation where the agent's error is linear in the level of customization chosen.

Figs. 7 and 8 illustrate some of the results of our analysis. Consistent with Proposition 5, increasing the rate of inference of the intelligent agent makes the buyer's surplus increasingly lower. Surprisingly, seller profits consistently displayed the same trend—a higher rate of agent inference, rather than helping

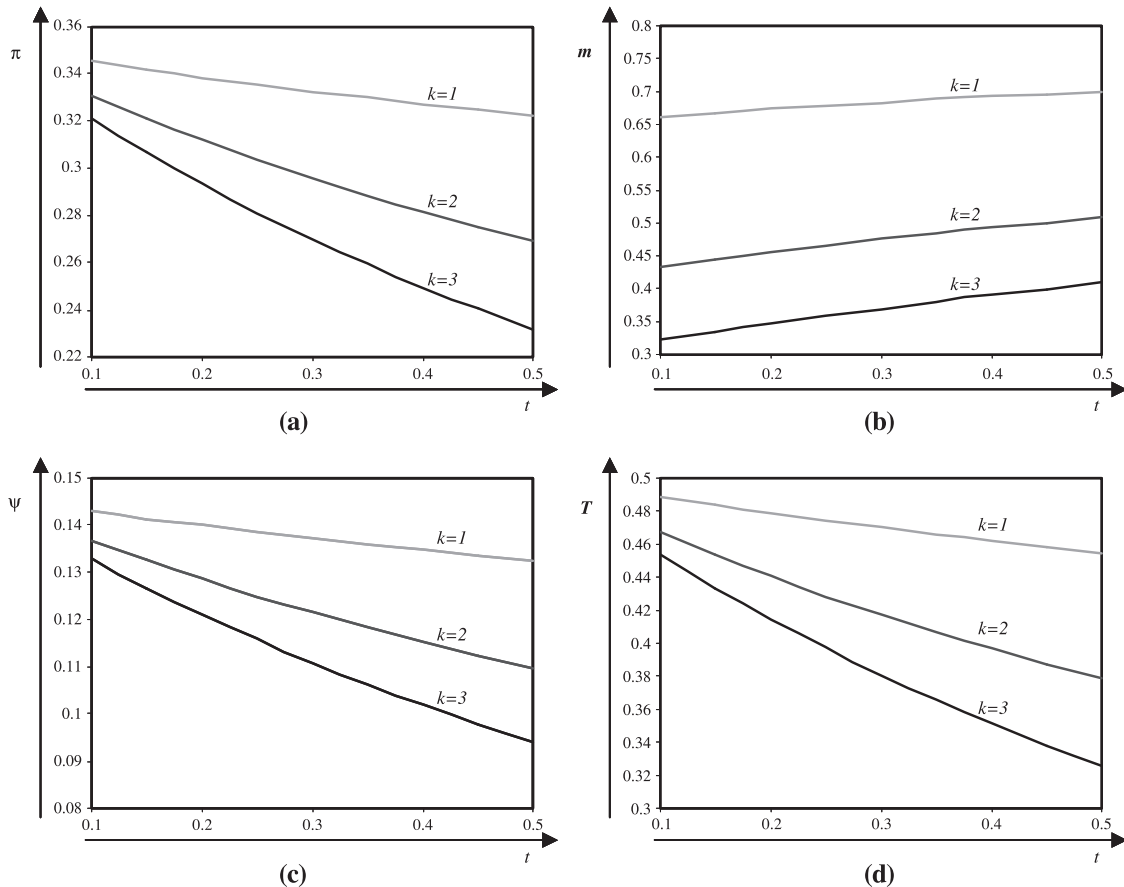


Fig. 7. Low valuation customers ($v=0.5$). (a) Profits. (b) Level of customization. (c) Consumer surplus. (d) Total surplus.

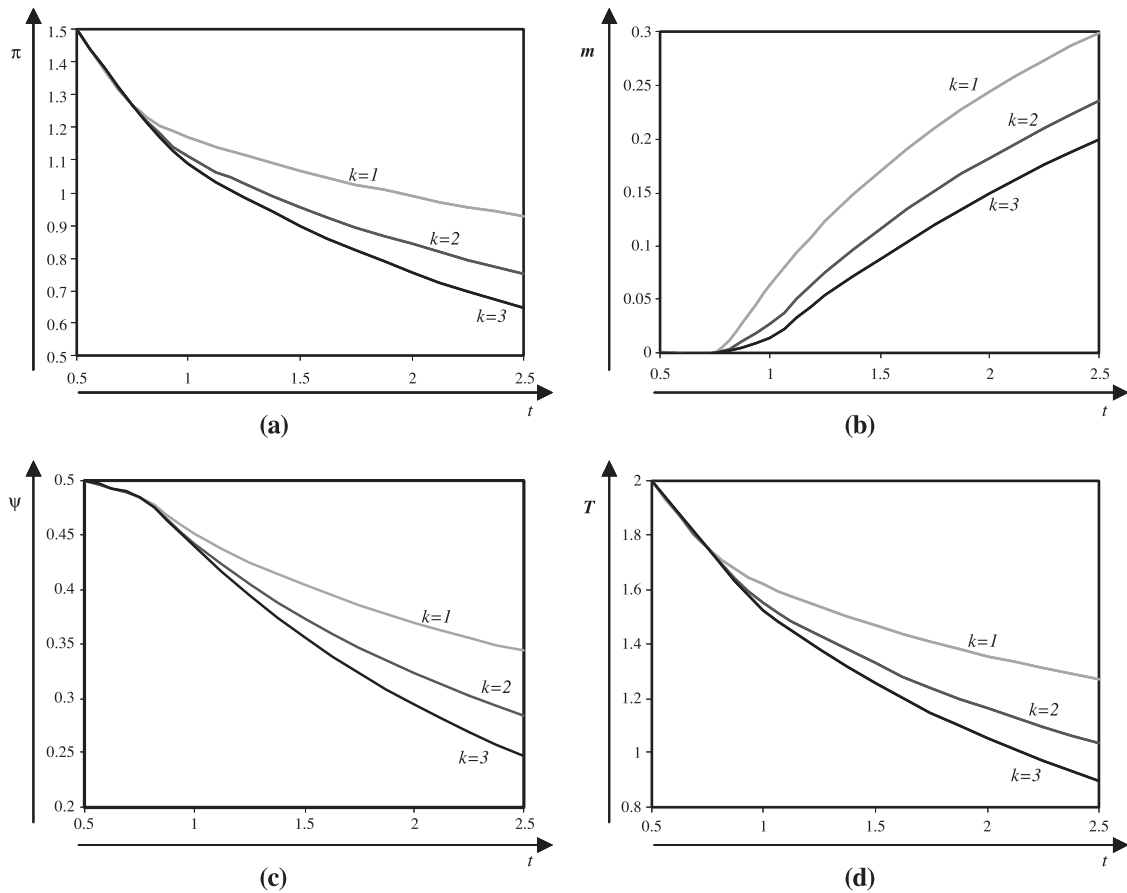


Fig. 8. High valuation customers ($v=2.5$). (a) Profits. (b) Level of customization. (c) Consumer surplus. (d) Total surplus.

the seller, actually reduces the seller's profits steadily. This result holds across a wide range of v and t values—we have depicted two sample ranges in Figs. 7 and 8. If one examines the level of customization charts in Figs. 7(b) and 8(b), the total surplus charts of Figs. 7(d) and 8(d), and relates them back to expression(A), this result can be explained. As indicated in Figs. 7(b) and 8(b), since the optimal level of customization chosen by the buyer drops dramatically as the agent's ability to make inferences improves, so does the total surplus $[v - \tau(\theta^*)]$ that the buyer and seller split. The revenue results simply indicate that the buyers adjust their behavior enough in such a way that the seller shares in this loss in surplus. In other words (referring back to Expression A), for the seller, the gains from superior agent inference are outweighed by the losses from the resulting information withholding by the buyer.

This observation leads us to a key insight: as the monopolist's ability to price discriminate improves, the high valuation buyer withholds a *disproportionately higher* level of information about herself thus resulting in her choosing a product that is a poorer fit with her ideal product. Thus, the high valuation buyer does indeed lose a greater proportion of her surplus than the low valuation buyer by withholding information. However, the surplus lost by the high valuation buyer is not captured by the monopolist, and most of it is deadweight loss that is reflected in lower total surplus.⁵

⁵ The range of values of t is different for different parts of Figs. 7 and 8—since we have normalized the value of m to lie between 0 and 1, fixed ranges of t result in some t values that are significantly higher or lower than the values of v . We investigate ranges of t that are between 20% and 100% of values of v .

Fig. 9 illustrates the reductions in optimal customization levels as k increases, highlighting the differential effect these changes have on low and high valuation buyers. As m decreases and k increases, higher valuation buyers choose increasingly lower levels of customization. In addition, Fig. 9 shows that improving agent technology has a much higher marginal effect on lower valuation buyers than on higher valuation buyers. The intuition behind this is that the effect of improved inference rates has a much higher net effect on buyers choosing higher values of m than on those choosing lower values of customization (the net reduction in m required to maintain the same value of θ is higher at higher values of m), and, consequently, the low valuation buyers, who choose higher values of m , are more adversely affected.

A related issue of interest is how the *split* in total surplus is affected by changes in the rate of inference.

Fig. 10 illustrates that as buyer valuations increase, sellers universally attract increasingly larger portions of the total surplus, across different values of k . This is illustrated by the fact that as v increases, the seller percentage revenue curves of Fig. 10(a) move up, while the buyer percentage surplus curves of Fig. 10(b) move down. As v increases, it does appear that sellers extract a smaller fraction of the total surplus, as the inference rates of their agents increase.

5. Conclusion

A number of products sold in electronic markets are consistent with our model of marginal-cost products which are highly customizable at minimal cost, whose buyers reveal information about their valuations through their product choices, and where

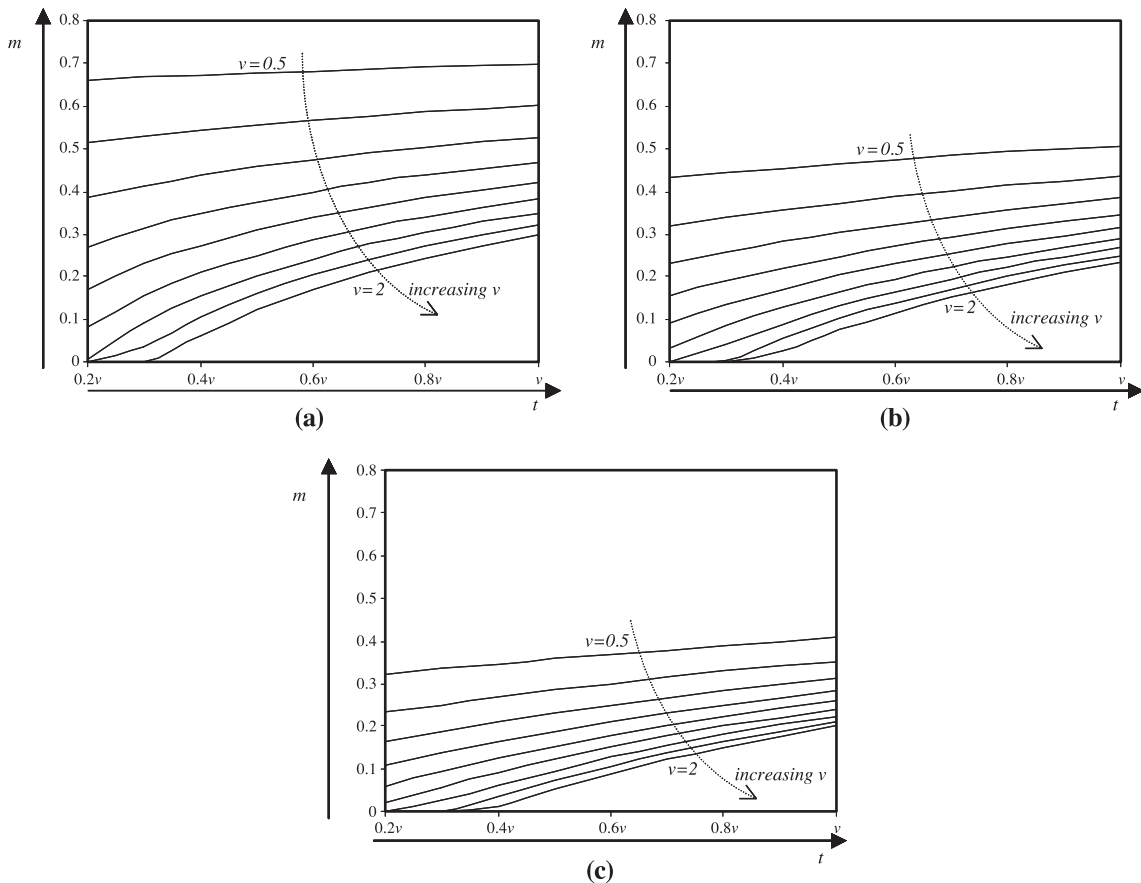


Fig. 9. (a) Slow pace of IA learning ($k=1$). (b) Moderate pace of IA learning ($k=2$). (c) Rapid pace of IA learning ($k=3$).

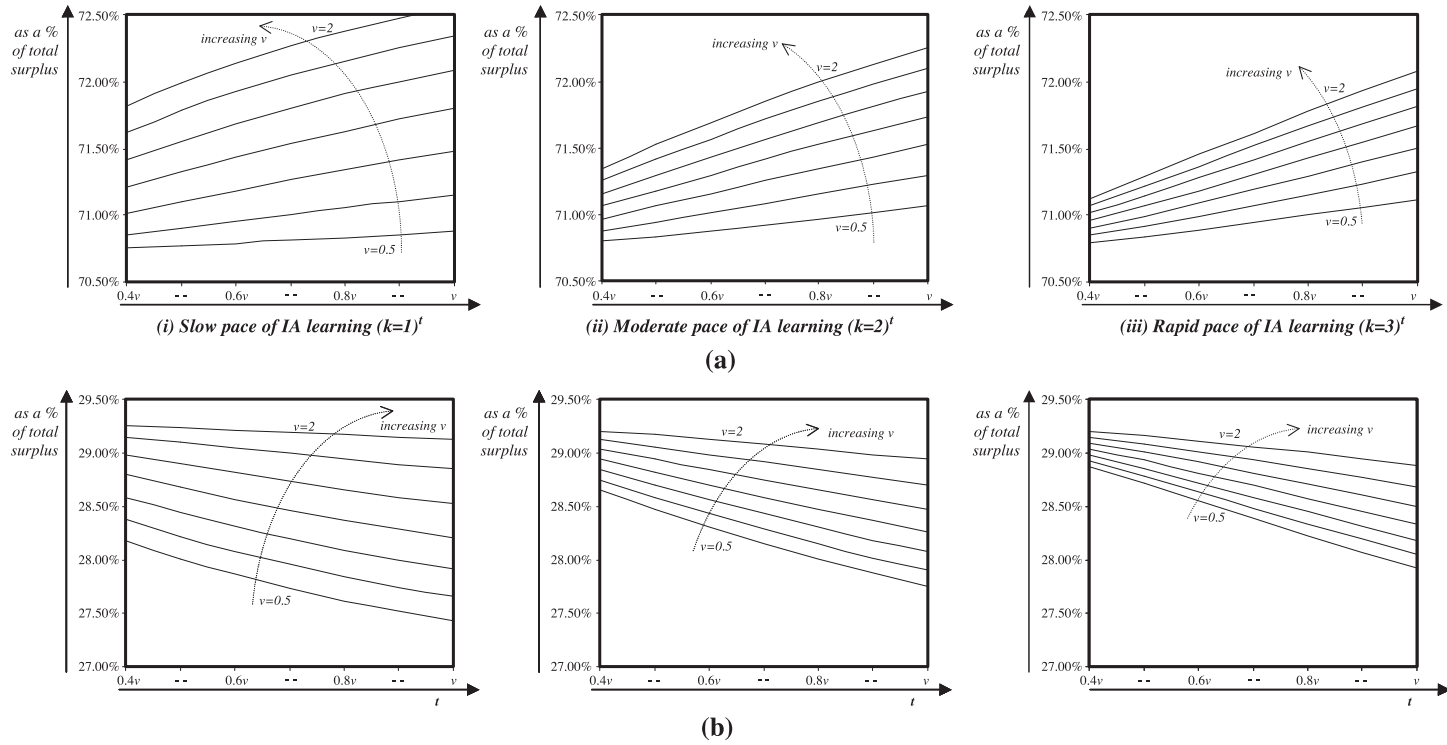


Fig. 10. (a) Seller profits as a percentage of total surplus. (b) Buyer surplus as a percentage of total surplus.

sellers can potentially price each sale. These include online retailing services, customizable gateway pages offered by portals, online financial services products, travel products, event tickets, credit card and banking services. As mentioned in Section 1, it is already common for online merchants of these products to indirectly extract surplus based on an analysis of user preferences. For instance, many web portals use self-reported customer information to place targeted advertisements, at a premium of 100% to 300%. As the level of customization increases, and the buyer reveals more about his/her preferences, the buyer ‘pays’ more attention to the advertisements (increasing the premium paid by the advertiser), and will eventually face the trade-off between higher ‘prices’ and better customized content.

Furthermore, as these merchants gain more monopoly power and user lock-in, they will evolve towards direct one-on-one dynamic pricing. Our analysis offers them the following key managerial insights:

- Intelligent agents cause buyers on the higher end of the market to move away from customizing their product choices, despite the fact that they actually value these ideally customized products more than lower-end buyers. This result holds for both fixed and value-proportionate unit costs of commoditization.
- As these buyers adjust their optimal product choices in response to better demand agent technologies, sellers may experience diminishing profits, since the gains from better buyer valuation information are countered by the lowering of the total surplus that the seller eventually extracts a portion of.

Consequently, sellers may actually benefit from limiting their use of buyer preference information to infer willingness-to-pay, so long as they credibly inform their customers that they are doing so. This kind of behavior is already widely observed in the context of consumer privacy. Clearly, on the face of it, companies could benefit by extracting as much surplus as possible from their customers’ personal information. However, a number of them willingly choose to assure their customers that they will not use or sell this information, and they make these statements credible through the endorsement of organizations like TRUSTe and BBB Online. Rather than on account of privacy ethics on the sellers’ parts, this choice is often because unless they promise not to use the information too much, they will not get the information at all. Our analysis shows that sellers using intelligent demand agents will face exactly the same trade-off. One might therefore anticipate the evolution of a similar structure, wherein sellers credibly promise not to extract too much of the informational rents they can get from their buyers’ preference descriptions. It is possible that technological watchdog agencies analogous to TRUSTe may emerge as demand agents become more popular, and that multiple sellers will seek the services of a single, well-known intelligent agent technology, which buyers understand and trust; this indicates significant market potential for a company that can establish itself as the trusted agent intermediary.

The differential impact that intelligent agents have on high and low valuation buyers has interesting implications. Figs. 11 and 12 illustrate the nature of the sources of revenue and surplus in markets driven

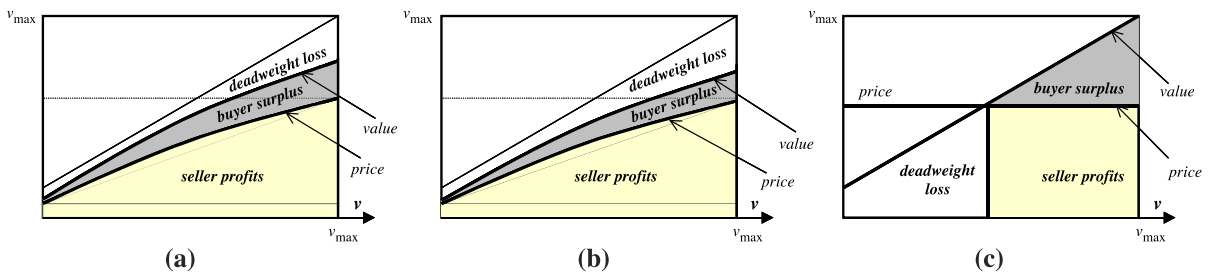


Fig. 11. Low cost of commoditization t . (a) IA: Slow pace of learning ($k=1$). (b) IA: Rapid pace of learning ($k=3$). (c) No IA: Fixed price.




Fig. 12. High cost of commoditization t . (a) IA: Slow pace of learning ($k=1$). (b) IA: Rapid pace of learning ($k=3$). (c) No IA: Fixed price.

by intelligent agents. They indicate that the desirability of these agents is higher in markets where the average cost of commoditization is lower. When customers are more product quality sensitive, using a pricing agent has a potentially adverse effect on seller profits. On the other hand, seller profits may increase vis-a-vis fixed pricing, if buyers in the market are not very sensitive to customization.

In addition, higher valuation buyers are more adversely affected by demand agents, and the use of these agents shifts deadweight loss from the low end of the market to the high end of the market. When a seller uses a demand agent, while the magnitude of total consumer surplus often is reduced, the distribution of surplus is far more even between buyers of varying valuation, which is illustrated by a comparison of the agent-driven surplus distributions to those with a fixed price.⁶ One strategy for sellers in this context is to credibly commit to using the intelligent agent on just lower valuation buyers, and one way of doing this is by committing to a fixed maximum price. This way, buyers on the low-end of the market, who may actually have been shut out with a fixed price, can enter the market. Simultaneously, buyers with high valuations can choose their ideal product with the assurance of a price cap.

We do not explicitly model seller strategies in a competitive market. Qualitatively, however, one effect of competition could be to induce agent-based differentiation among sellers, forcing them to choose between low and high valuation buyers. Some sellers

may offer limited customization with no real ceiling price, and attract a relatively high proportion of low value sellers. Others may attract high valuation buyers by offering a combination of superior customization technology and a credible commitment to a price ceiling. Since the price ceiling would determine the threshold valuation above which buyers must value products to transact with a seller, this suggests that there could be a continuum of ceiling prices along which sellers could initially differentiate themselves (based on the quality of their agent technology). In contrast, if all sellers were identical and there really existed no viable basis for differentiation, as more sellers entered the market, we would expect to see that the gains from employing intelligent agents be competed away, and, in a free-entry model, equilibrium customization levels being driven by the fixed costs of deploying agent-based technology.

The presence of an outside good which is an imperfect substitute for the customizable product could have varying effects on our results. Since buyers can customize products until the candidate product is identical to their ideal product, this implies that all versions of the product less customized than the ideal one can be treated as imperfect substitutes of the ideal product, and in that sense, our model deals adequately with the presence of imperfect substitutes, at least indirectly. The direct presence of noncustomizable outside goods (such as preconfigured music CD's, or integrated research reports) could have other effects—it may limit the set of buyers who are interested in customizing, and consequently, redefine the support of v without changing the results of the model otherwise. In contrast, the presence of several sellers that offer

⁶ These comparisons are qualitative. A precise comparison requires us to assume an upper bound on consumer valuations, which changes the behavior of a demand agent close to this bound.

imperfect substitutes at low prices could limit the extent to which the sellers that wish to sell to high valuation buyers can price-discriminate.

As mentioned earlier, a third party that certifies the extent to which various technologies can be used for price discrimination would allow the buyer to reveal more preference information, thereby reducing the deadweight loss from the use of intelligent agents. However, a combination of competition and repeat purchases buyers could achieve the same outcome as a third-party certifying agency, if there is a mechanism for recording and reporting reputation. Sellers that commit *ex ante* to a self-enforced limit on their extent of price discrimination and who do not renege on the commitment would attract higher repeat sales, especially from high valuation buyers. On the other hand, sellers that do not make or adhere to such commitments would lose customers in future periods (especially the high valuation buyers). A combination of competition and repeat purchasing could therefore lead to self-regulated agent use, although the optimal technology deployment might well be different from that induced by an intermediary.

We also prescribe the use of intelligent pricing agents only after a seller has acquired a significant degree of market power, which explains our focus on a monopolist seller. The value of pricing agents in an environment of undifferentiated price competition is minimal, especially if the product sold has zero marginal cost. However, once consumers are locked-in due to the brand-specific training, sunk information, and personal search costs—all significant in the context of portals, online retailing and financial services—the firm essentially acts as the monopoly described in our model. This is supported by evidence that customer loyalty is high in electronic commerce [8,10], and also by results from work by Ulph and Vulcan [35] analyzing competitive models in similar, albeit simplified versions of our setting, where they conclude that if mass customization strategies are chosen by two competing firms, then the profits of each firm are independent of the price discrimination strategies chosen by the other firm.

While sellers do not necessarily benefit from better IA technologies, the threat of a rival's

adoption of demand agents might force sellers to adopt these technologies. Similar results have been shown in studies on interorganizational systems and EDI by Barua and Lee [4] and Riggins et al. [29]. As noted by Clemons and Kimbrough [11], sometimes these IT investments become strategic necessities, rather than yielding competitive advantage. This represents an interesting direction for future research. Another interesting extension may be studying the effects of sequential learning by the demand agent. The technological implications are likely to be better inference over time, or in the language of our paper, a narrower window around the buyer's true valuation with each successive interaction. Our model already has the elements to capture the effects of this improvement—via a rescaling of the inference function, by an increase in the rate of learning, or by a combination of the two. This enables us to hypothesize on the overall effects of learning, based on our current results—if the buyer is aware that with each interaction, the agent's accuracy increases at some exogenously specified rate, learning is likely to strengthen one of the basic results of our paper—that demand agents can impact seller profits adversely, and that this effect is more pronounced for higher valuation buyers. A buyer will withhold more preference information in early periods, and will also progressively choose lower levels of customization as the agent's performance improves. It is likely that this will result in lower total surplus, but a higher percentage of the split going to the seller. The complexity of a model of this sort makes its analysis beyond the scope of this paper; however, it is an extension we may consider for future research.

This paper is the first systematic analysis of the economic implications of intelligent agent technologies that infer buyer valuations in a market for customized information products. Our results are consistent with contemporary business trends, and prescribe business strategies for the many companies who will be faced with the decision of whether to use this technology, and if so, how to design and target it effectively. We hope that our paper will serve as the foundation for research that further enhances our understanding of the future of agent-driven commerce.

6. Summary of notation

v	Buyer valuation of ideal product.
m	Level of customization chosen by the buyer. $m \in [0,1]$. $m=0$ implies a choice of a generic product, and $m=1$ implies a choice of the buyer's ideal product.
t	Unit cost of commoditization. A choice of a level of customization m results in a cost of commoditization $t(1-m)^2$, and a net value of $v-t(1-m)^2$ to the buyer.
$\theta(m)$	Width of the interval estimate of a buyer's valuation by the intelligent agent, when the buyer's level of customization is m . $\theta'(m) < 0$, $\theta''(m) \geq 0$.
$\theta_{\max} = \theta(0)$	Width of the interval estimate at the lowest level of commoditization.
$\theta_{\min} = \theta(1)$	Width of the interval estimate, when the buyer chooses his/her ideal product.
ε	Lower support of the interval estimate.
$\tau(\theta)$	Cost of commoditization borne by the buyer choosing a level of customization m corresponding to an interval width θ . $\tau(\theta) = t(1-m(\theta))^2$.
α	Lower support of the buyer's prior on ε . $\alpha = \max[\tau(\theta), v - \theta]$.
$p(\theta)$	Price expected by the buyer choosing a level of customization corresponding to an interval width θ . $p(\theta)$ is a random variable, not an average value.
$f(p)$	Density of $p(\theta)$.
p^*	Optimal price chosen by seller.
$\psi(\theta)$	The net consumer surplus expected by the buyer when choosing a customization level corresponding to an interval width θ .

References

- [1] M. Armstrong, Multiproduct nonlinear pricing, *Econometrica* 64 (1) (1996) 51–75.
- [2] W. Baker, M. Marn, C. Zawada, Pricing smarter on the net, *Harvard Business Review*, February 2001 (2001) 22–27.
- [3] J.Y. Bakos, E. Brynjolfsson, Bundling information goods: pricing, profits and efficiency, *Management Science* 45 (12) (1999) 1613–1630.
- [4] A. Barua, B. Lee, An economic analysis of the introduction of an electronic data interchange system, *Information Systems Research* 8 (4) (1997) 398–422.
- [5] A. Barua, C.H. Kriebel, T. Mukhopadhyay, An economic analysis of the strategic impacts of information technology investments, *MIS Quarterly* 15 (3) (1991) 313–331.
- [6] H.K. Bhargava, V. Choudhury, Information goods and vertical differentiation, *Journal of Management Information Systems* 18 (2) (2001) 89–106.
- [7] J. Blank, D. Daniels, J. Gibs, M. Guinle, M. May, R. Leathern, T. Clark, V. Sehgal, Personal value pricing: using dynamic pricing tools to maximize revenues throughout a product's life cycle, *Vision Report Jupiter Research* (2001 June) 5–6.
- [8] E. Brynjolfsson, M. Smith, The Great Equalizer? Consumer Choice Behavior at Online ShopBots, Working Paper, MIT Sloan School of Management, 2000.
- [9] T. Bui, J. Lee, An agent-based framework for building decision support systems, *Decision Support Systems* 25 (3) (1999) 225–237.
- [10] P. Chen, L. Hitt, Switching Costs and Brand Loyalty in Electronic Markets: Evidence from Online Retail Brokers. Working Paper, University of Pennsylvania, 2000.
- [11] E.K. Clemons, S.O. Kimbrough, Information Systems and Business Strategy: A Review of Strategic Necessity. Working Paper, The Wharton School, Philadelphia, PA, 1987.
- [12] E.K. Clemons, P. Kleindorfer, An economic analysis of interorganizational information systems, *Decision Support Systems* 8 (5) (1992) 431–446.
- [13] C. d'Aspremont, J. Gabszewicz, J.-F. Thisse, On Hotelling's 'stability in competition', *Econometrica* 47 (5) (1979) 1145–1150.
- [14] R. Dewan, B. Jing, A. Seidmann, Adoption of internet-based product customization and pricing strategies, Proceedings of the 31st Hawaii International Conference on System Sciences, 2000.
- [15] P. Fingar, 1999. Intelligent Agents: The Key to Open eCommerce Building The Next-Generation Enterprise. Available at <http://home1.gte.net/pfingar/csAPR99.html>.
- [16] L. Frey, D. Fisher, Modeling decision tree performance with the power law, Proceedings of the Seventh International Workshop on Artificial Intelligence and Statistics, 1999.
- [17] C.R. Johnson, 1997. Intelligent Agents Breathe Life Into Internet. Available at <http://www.eet.com/news/97/946news/intelligent.html>.
- [18] R. Jones, H. Mendelson, Product and Price Competition for Information Goods. Working Paper, University of Rochester, 1998.
- [19] P. Konana, A. Gupta, A.B. Whinston, Integrating user preferences and real-time workload in electronic commerce, *Information Systems Research* 11 (2) (2000) 177–196.
- [20] P. Maes, Agents that reduce work and information overload, *Communications of the ACM* 37 (7) (1994) 31–40.
- [21] J.F. MacKie-Mason, L. Murphy, J. Murphy, The role of responsive pricing in the Internet, in: J. Bailey, L. McKnight (Eds.), *Internet Economics*, MIT Press, 1997, pp. 279–303.
- [22] E. Maskin, J. Riley, Monopoly with incomplete information, *Rand Journal of Economics* 15 (2) (1984) 171–196.
- [23] S. McDonnell, Microsegmentation, *Computerworld* (2001 January 29) <http://www.computerworld.com/databasetopics/data/story/0,10801,56969,00.html>.
- [24] T. Mukhopadhyay, U. Rajan, R. Telang, Competition Between Internet Search Engines. Working Paper, Carnegie-Mellon University, 2000.

- [25] M. Mussa, S. Rosen, Monopoly and product quality, *Journal of Economic Theory* 18 (1978) 301–317.
- [26] B.R. Nault, Quality differentiation and adoption costs: the case for interorganizational information systems pricing, *Annals of Operation Research* 71 (1997) 115–142.
- [27] B.R. Nault, A.S. Dexter, Added value and pricing with information technology, *MIS Quarterly* 19 (4) (1995) 449–463.
- [28] F. Provost, D. Jensen, T. Oates, Efficient progressive sampling, *Proceedings of the Fifth International Conference on Knowledge Discovery and Data Mining*, 1999.
- [29] F.J. Riggins, C.H. Kriebel, T. Mukhopadhyay, The growth of interorganizational systems in the presence of network externalities, *Management Science* 40 (1994) 984–998.
- [30] J. Sairamesh, J.O. Kephart, Price dynamics and quality in information markets, *Decision Support Systems* 28 (2) (2000) 35–47.
- [31] A. Shaked, J. Sutton, Relaxing price competition through product differentiation, *Review of Economic Studies* 49 (1982) 3–14.
- [32] J. Stephanek, Weblining: companies are using your personal data to limit your choices—and force you to pay more, *Business Week* (2000 April 3) http://www.businessweek.com/2000/00_14/b3675027.htm.
- [33] A. Sundararajan, Nonlinear pricing of information goods, *Management Science* 50 (12) (2004).
- [34] G. Tesauro, J.O. Kephart, Pricing in agent economies using multi-agent Q-learning, *Proceedings of the Game Theoretic and Decision Theoretic Agents Workshop*, 1999.
- [35] D. Ulph, N. Vulkan, e-Commerce, Mass Customisation and Price Discrimination. Working Paper, ESRC Center for Economic Learning and Social Evolution, London, U.K., 2001.
- [36] H. Varian, 1995. Pricing Information Goods. Available at <http://www.sims.berkeley.edu/~hal/Papers/price-info-goods.pdf>.
- [37] E. Weise, Turning the Net into your personal shopping valet, *USA Today*, 1997, Nov. 25 04D.
- [38] R. Wilson, *Nonlinear Pricing*, Oxford University Press, 1993.
- Extended Appendix. Available at: <http://oz.stern.nyu.edu/agent/appendix.pdf>.
- [2] R. Schmalensee, Output and welfare implications of monopolistic third-degree price discrimination, *American Economic Review* 71 (1981) 242–247.
- [3] G.J. Stigler, The economics of information, *The Journal of Political Economy* 69 (1961) 213–225.
- [4] H. Varian, C. Shapiro, *Information Rules*, Harvard Business School Press, 1999.

Dr. Ravi Aron is Assistant Professor of Operations and Information Management at the Wharton School. He is also a member of the Information Systems: Strategy and Economics (ISSE) group at the Wharton School. Dr. Aron's research addresses the pricing and sourcing of information-rich services and the global sourcing of IT-enabled services by firms. His research also addresses the design of business-to-business markets and the global sourcing of products by firms through electronic markets. Dr. Aron has published several articles in research journals on these topics and has presented his work at a number of conferences to both practitioners and researchers in academia.

Dr. Arun Sundararajan is an assistant professor at New York University's Stern School of Business. His research on the economics of information technology has been widely published in journals such as *Management Science*, *Information Systems Research*, *Decision Support Systems*, *Journal of Management Information Systems*, and *Economics Letters*. His recent work has focused on topics that include price discrimination for digital goods, digital piracy, competition in IT industries, and the structure of economic networks. He serves on the editorial board of *Management Science* as an associate editor and as the director of the IT economics track in NYU's Center for Digital Economy Research.

Dr. Siva Viswanathan is an Assistant Professor in the Department of Decision and Information Technologies at the Robert H. Smith School of Business, University of Maryland, College Park. He holds a PhD from New York University, an MBA, and a bachelor's degree in engineering from India. Dr. Viswanathan's current research interests include studying the competitive and strategic implications of emerging technologies and the impact of online intermediaries. In addition to his recent publication in *Management Science*, he also has several research presentations, book chapters, and articles published in conference proceedings.

Further reading

- [1] R. Aron, A. Sundararajan, S. Viswanathan, 2004. *Intelligent Agents in Electronic Markets for Information Goods*: