

OPTIMAL SEARCH FOR PRODUCT INFORMATION*

FERNANDO BRANCO

(Universidade Católica Portuguesa)

MONIC SUN

(Stanford University)

J. MIGUEL VILLAS-BOAS

(University of California, Berkeley)

December, 2010

* Comments by Dmitri Kuksov, and seminar participants at the 2010 Choice Symposium and Hong Kong University of Science and Technology are gratefully appreciated. E-mail: fbranco@ucp.pt, monic@stanford.edu, and villas@haas.berkeley.edu.

OPTIMAL SEARCH FOR PRODUCT INFORMATION

ABSTRACT

Consumers often need to search for product information before making purchase decisions. We consider a parsimonious model in which consumers incur search costs to learn further product information, and update their expected utility of the product at each search occasion. We characterize the optimal stopping rules to either purchase, or not purchase, as a function of search costs and the informativeness of each attribute. The paper also characterizes how the likelihood of purchase changes with the ex-ante expected utility, search costs, and informativeness of each attribute. We discuss optimal pricing, the impact of consumer search on profits and social welfare, and how the seller chooses its price to strategically affect the extent of the consumers' search behavior. We show that lower search costs can hurt the consumer because the seller may choose to then charge higher prices. The paper also considers the impact of searching for signals of the value of the product, of discounting, and of endogenizing the intensity of search.

1. INTRODUCTION

Most consumers search for product information before making their purchasing decisions. They examine the product, touch it, ask for information from other consumers, and maybe read about it. As they do this, they update their beliefs on how much they would enjoy the product. At a certain point, consumers decide not to search any more, and make a decision as to whether or not to purchase the product.

As the Internet has made information gathering seemingly easier than ever, consumers still end up making purchase decisions without full information because there are still search and processing costs (see, e.g., Chen and Sudhir 2004). Furthermore, with the Internet it has become easier to monitor a consumer's search behavior. Given this environment, it has become more important to understand a consumer's optimal search strategy for information. In purchasing situations consumers gather product information for the purpose of trying to decide whether to make a purchase, updating the product's expected valuation each time new information is obtained.

This question also has important implications for managers. By trying to understand the factors that affect information search, a manager could potentially influence the extent of search. *A priori*, it is not clear whether it benefits the seller when consumers spend more time searching for product information – the information obtained could either increase or decrease consumer valuation of the product. As a result, the extent of search is closely linked with the optimal price and profit. Given these considerations, we study the following questions in this paper. First, when would a consumer stop searching for more information and make a decision on whether to purchase a product? Second, given an understanding of the optimal stopping rule, how can a seller manipulate price to influence the extent of search? Third, what factors determine the consumer's search behavior and the firm's pricing decisions in this context?

In order to answer these questions we consider a framework in which consumers can sequentially gather information about the product, updating their expected valuation for the product at each search occasion. We consider a parsimonious preference model in which the valuation of the product is the sum of the utility from multiple product attributes. As a consumer searches, he obtains information from one attribute at a time.¹ The information contained in an attribute could either increase or decrease the consumer's expected valuation of the product. With an infinite number of attributes, each providing an infinitesimal amount of information, we can obtain that the expected valuation follows a Brownian motion while searching occurs.

¹Throughout the paper we refer to the consumer as a 'he,' and to the seller as a 'she'.

We show that the optimal stopping rule consists of an upper bound and a lower bound, which we refer to as “purchase threshold” and “exit threshold,” respectively, on the consumer’s expected valuation. When the consumer’s expected valuation of the product hits the purchase threshold, the consumer stops searching and purchases the product. When the expected valuation hits the exit threshold, the consumer stops searching and does not purchase the product. When the expected valuation is in between the two bounds, the consumer continues to search.

We derive closed-form expressions for these optimal two bounds. One can obtain how they are determined by the marginal cost of search and the importance of each product attribute, i.e., the amount of information contained in each product attribute. If it is more costly for the consumer to keep searching, consumers search less, and the purchase threshold becomes lower. This result suggests that the purchase threshold may actually become higher as the Internet drives down consumer’s search costs. One implication of this is that the likelihood to purchase can decrease for some purchase occasions with the development of the Internet. The purchase threshold increases with the amount of information contained in each product attribute: If the consumer’s expected valuation of the product is more volatile as he searches through different product attributes, he wants to be more confident and positive about the product before deciding to purchase, hence the purchase threshold is also higher.

Depending on the ex-ante valuation, the seller either charges a low price to eliminate search for information or a high price to stimulate search. When the ex-ante expected valuation (prior to any search) is high, the seller maximizes her profit by charging a relatively low price and having the consumer buy the product right away. Stimulating search in this case has a big cost (driving down purchase likelihood) and a small benefit (potentially increasing the consumer’s willingness to pay). When the ex-ante expected valuation is low, the seller encourages the consumer to search for more product information. In the latter case, the price would be too low if the seller insists on getting the consumer to purchase without any search. In equilibrium, the purchase likelihood depends on the ratio of the ex-ante valuation and the purchase threshold.

The profit always increases with the ex-ante expected valuation of the product. The search cost, however, enters into the profit equation in a more sophisticated manner. Essentially, lower search costs lead to more search, or in other words, a higher purchase threshold. When this happens, the optimal price increases, but the purchase likelihood can decrease. As these two forces trade off against each other, the profit may either increase or decrease with the purchase threshold. Overall, the seller wants to encourage search at the margin, or increase the purchase likelihood, if the ex-ante expected valuation is low, as information search may lead to both higher purchase likelihood and higher price. She wants to hinder search if the consumer’s ex-ante expected valuation of the

product is high, and have the consumer buy the product with little additional information. The seller may be able to influence the purchase threshold by changing how much product information is made available and how such information is communicated to the consumer.

We consider a variety of extensions of the base model. We show that the model is qualitatively similar to one in which consumers try to infer the “true” valuation from a series of independent signals. Technically, the latter model generalizes our baseline model in making the amount of information contained in each attribute vary through the search process. We also discuss what happens when the mass of attributes that the consumer could check is finite. In another extension, we investigate the effect of allowing consumer discounting and show that in that case an indifferent consumer is more likely to purchase the product. Finally, we allow for the possibility of consumers choosing the search intensity per unit of time. In this case we show that consumers search with greater intensity when their valuation of the product is closer to the purchase threshold.

The current paper contributes to the literature on information gathering. Papers in this literature often study searching for the lowest price (e.g., Diamond 1971, Stahl 1994, Kuksov 2004), or searching for the best-matched alternative (e.g., Weitzman 1979, Wolinsky 1986, Bakos 1997, Anderson and Renault 1999, Armstrong et al. 2009) with complete learning after one alternative/price is searched.² Our paper, on the other hand, focuses on the fundamental properties of gradual, cumulative private learning through search prior to the decision of whether to purchase a single product. We present a tractable and rich model of consumer search for information that includes uncertainty on which information is obtained at each search stage, and optimal stopping rules for search that lead to either purchase or no purchase. Given the optimal stopping rules, we derive the optimal pricing decision by a firm, taking into account the effect of price on information search.³

²Weitzman (1979) studies Pandora’s problem: if there are multiple closed boxes, each containing an uncertain reward, what would be the optimal sequence of opening the boxes? That paper is similar to this paper as costs are cumulative while rewards are uncertain, but information in that paper is not cumulative, as the other search papers cited above: the reward in each box has a probability distribution that is independent of the other rewards. In the current paper, there is only one ultimate valuation of the product, and all search has the purpose of figuring out what that valuation is. Hauser et al. (1993) consider the rules that a consumer can use on deciding the order to check different sources of information. See also Moorthy et al. (1997) for an investigation of the empirical implications from the results in Weitzman. For competition with search costs see also, for example, Varian (1980), Narasimhan (1988), Kuksov (2006). For the effect of search costs creating a potential hold-up problem see, for example, Lal and Matutes (1994), Wernerfelt (1994), Rao and Syam (2001), Villas-Boas (2009). Search costs can also have effects on consideration sets (e.g., Hauser and Wernerfelt 1990, Mehta et al. 2003), and channel structure (e.g., Gal-Or, Geylani and Dukes). Consumers can also obtain information by learning from other consumers (e.g., Ellison and Fudenberg 1995, Zhang 2010), or experimenting across multiple products (e.g., Bergemann and Välimäki 1996, Villas-Boas 2004). For empirical estimates of search costs see, for example, Honka (2010), De los Santos et al. (2010).

³Also related to this paper in terms of modeling of continuous learning of information are Bolton and Harris (1999) and Bergemann and Välimäki (2000). Both of these papers model the evolution of beliefs of agents about

One related paper in terms of model setup is Roberts and Weitzman (1981). That paper models Sequential Development Projects (SDP) that feature additive costs and uncertain benefits. While that paper focuses on how firms can choose to optimally decide whether to continue researching a sequential project, we focus on how a consumer gathers information on a product before deciding whether to purchase it. One distinct difference between our model and Roberts and Weitzman is that in that model one key feature is the information eventually running out as the outcome of the Research and Development (R&D) becomes clear. Consequently, the boundaries of search in his model would always shrink as the firm moves toward the end of uncertainty. Our model, on the other hand, assumes that the consumer can always search for more information before he makes a purchase, motivated by the observation that online search engines now provide boundless opportunities for a consumer to research a product. This allows us to obtain sharper results on the consumer search thresholds and search behavior. In addition we also model in detail how consumer beliefs evolve through the search process given the ex-ante priors and expectations. The focus on the consumer, combined with the departure from the limited-information assumption, enables us to get closed-form results, provide insights into optimal consumer search behavior for information, generate implications for pricing, and discuss the impact of discounting and the choice of search intensity.

The remainder of the paper is organized as follows. The next section presents the model, and Section 3 characterizes the optimal consumer search. Section 4 considers the optimal decisions of a firm, and Section 5 discusses extensions of the model. Section 6 presents concluding remarks.

2. THE MODEL

Consider a consumer who gathers information sequentially on the attributes of a particular product to decide whether or not to buy it.⁴ The total utility that he ultimately derives from the product is the sum of utilities obtained from each of the T product attributes. This can be written as the ex-ante expected utility prior to learning about any attribute, v , plus the sum over attributes of the difference between the realized utility from an attribute and his expected utility.

whether a product has high or low quality, while here we model the beliefs for any level of quality, which is important to understand when a consumer stops searching to buy or stops searching not to buy a product. Neither of these papers looks at optimal consumer search for information, considering instead strategic experimentation in Bolton and Harris and experimentation with social learning in Bergemann and Välimäki.

⁴Even in situations where information on multiple attributes is displayed to the consumer simultaneously, the consumer may have to process these attributes one after another. We do not model the seller's incentive to add or delete product attributes. See Ofek and Srinivasan (2002) for a discussion on that issue.

For attribute i , we denote this difference between the realized utility from that attribute and the expected utility as x_i . The realized utility from a product is then

$$U = v + \sum_{i=1}^T x_i.$$

The term v can also be seen as the stand-alone valuation v of the product, as utility from the most basic function of the product, such as the “picture-taking” function of a camera, the “hunger-reducing” function of a meal, or the “getting-you-around” function of a car. Without searching for information, the only thing the consumer knows is v . In this interpretation, besides the most basic function, there exist product attributes that additively affect the total utility.

Multiple attributes of a product are often considered before a purchase is made. For example, consumers can check the color, pattern, material, and size of the product. The realization of each attribute (e.g., a car being yellow) either increases or decreases the total expected utility. That is, x_i , the deviation from the expected utility of the attribute, can be either positive or negative. We assume that the x_i ’s are independent across i . One interpretation is that after checking t attributes, checking attribute i changes the expected utility by x_i with what is unexpected given the information on the prior t attributes checked. Given that x_i is unexpected it is then independent of the previous x_i ’s that were checked.

Suppose that x_i can take the value of z or $-z$ with equal probability, such that the expected value of x_i is zero. The value of z can be interpreted as the weight of a product attribute on the consumer’s utility. That is, the change in utility from the attribute, x_i , is equal to $z \cdot y_i$, where y_i takes the value of 1 or -1 with equal probability. We consider that z is the same across all attributes, but in general, the analysis remains the same as long as consumers do not know the value of z prior to checking an attribute.⁵ We assume that the number of attributes, T , is infinity. This allows us to get sharp results, and captures the idea that the number of attributes is large enough, such that the set of attributes is never fully searched.⁶

The consumer does not know *a priori* the value of x_i for any i . When examining an attribute i , a consumer pays a search cost c and learns the realization of x_i . His expected utility goes either up or down by z . As it will become clear in the next section, the consumer essentially trades off the cost of search with the possibility of gathering positive information which eventually leads to a

⁵Alternatively, one could model that at each search occasion the consumer gets an independent signal of the product quality. This would lead to similar general ideas, with less sharp results. In Section 5 below we explore the features of such a model.

⁶In Section 5 below we discuss what happens when the set of attributes that consumers can check is bounded.

purchase.

After examining t attributes, the “current expected valuation” of the product becomes

$$u(t) = v + \sum_{i=1}^t x_i + \sum_{i=t+1}^T E(x_i|t) = v + \sum_{i=1}^t x_i.$$

In the presentation, when it is clear, we write u in place of $u(t)$ for ease of notation. One can see u as a state variable that changes stochastically with the number of attributes checked. Note that by the assumptions above, $E(x_i|t) = 0$ for all i . Note also that the search cost is assumed to be the same for all attributes. The search cost could take many forms in reality: the time and effort it takes to use search engines to gather relevant information on the product attribute, the mental costs of processing information on an attribute (e.g., Shugan 1980), the travel cost to the store to get a feel for the product, or even the monetary cost paid to a consultant if the product is rather complicated.

In order to simplify the solution of the dynamic programming problem we work with a continuous rather than discrete process.⁷ Imagine that there are many small-valued attributes, such that the change in utility from each attribute, z , gets infinitely smaller as the number of attributes goes to infinity. The u process then becomes continuous. In the limit, it becomes a Brownian motion (Cox et al., 1979). That is

$$du = \sigma dw,$$

where w is a standardized Brownian motion and $\sigma > 0$ is the standard deviation of the Brownian motion. By definition, $z = \sigma\sqrt{dt}$. The parameter σ therefore indicates how informative learning about an attribute is, or more generally, the amount of updating the consumer can do when checking each product attribute. Figure 1 below gives an example of how u can change over time.

After checking t units of attributes, the consumer has to make two decisions. First, does he want to continue examining attributes, or stop? Second, if he stops gathering information, should he buy the product or not? The answer to these questions is provided in the next section.

⁷While in some cases product attributes are obviously discrete, in other cases the information gathering process can be rather continuous. For example, a consumer might be constantly updating his beliefs on the valuation of a novel or a movie when reading or watching a sample of it.

3. OPTIMAL CONSUMER SEARCH

3.1. Expected Utility of the Search Problem

To analyze the optimal consumer search it helps to characterize the expected utility function of the search problem under the consumer's optimal decision making, $V(u, t)$, given that the expected utility of the product if purchased without further search is u and the consumer has checked t units of attributes. This expected utility of continuing to search, after checking t units of attributes, can be written as

$$V(u, t) = -c dt + EV(u + du, t + dt). \quad (1)$$

By a Taylor expansion (see, e.g., Dixit 1993), valid to terms of first order and less in dt , we can have

$$V(u, t) = -c dt + V(u, t) + V_u E(du) + V_t dt + \frac{1}{2} V_{uu} E[(du)^2] + V_{ut} E(du) dt \quad (2)$$

where V_u is the partial derivative of $V(u, t)$ with respect to u , V_t is the partial derivative of $V(u, t)$ with respect to t , V_{uu} is the second derivative of $V(u, t)$ with respect to u , and V_{ut} is the cross derivative of $V(u, t)$ with respect to u and t . Equation (2) must hold for all sufficiently small dt (using Itô's Lemma).

As $E(du) = 0$ and $E[(du)^2] = \sigma^2 dt$ we have, dividing (2) by dt ,

$$-c + V_t + \frac{\sigma^2}{2} V_{uu} = 0. \quad (3)$$

Since the mass of attributes available to check is unbounded by assumption, the mass of attributes already checked does not matter:

$$V_t = 0. \quad (4)$$

From (3) and (4),

$$V_{uu} = \frac{2c}{\sigma^2}, \quad (5)$$

hence $V(u)$ is quadratic in u .⁸

In addition, we also have a set of boundary conditions, when u is large enough such that the consumer chooses to buy the product, and when u is low enough such that the consumer chooses to stop searching and not buy the product. When u is large enough such that the consumer is

⁸As $V(u, t)$ is independent of t we then just write $V(u, t)$ as $V(u)$.

indifferent between continuing the search process, and stopping the search to purchase the product, we have

$$V(\bar{U}) = \bar{U} \text{ and } V'(\bar{U}) = 1, \quad (6)$$

where \bar{U} is the upper bound such that when u reaches \bar{U} the consumer stops searching and buys the product. Graphically, condition (6) states that the curve of $V(u)$ is tangent to the 45 degree line at $u = \bar{U}$ (see Figure 2 below).

The first part of condition (6) just means that when the expected utility of buying the product reaches $u = \bar{U}$ then the consumer chooses to purchase the product and gets the expected utility \bar{U} .⁹ The second part of condition (6) is known as the “high contact” or “smooth pasting” condition (e.g., Dumas 1991, Dixit 1993). The condition is implied by the fact that the consumer maximizes $V(u, t)$ for all (u, t) .

To get some intuition on the high contact condition, note that by searching extra dt attributes when reaching $u = \bar{U}$, a consumer pays a cost $c dt$ and gets

$$\frac{1}{2}(\bar{U} + E[du|du \geq 0]) + \frac{1}{2}E[V(\bar{U} + du)|du < 0].$$

By a Taylor approximation, and $V(\bar{U}) = \bar{U}$, we have $V(\bar{U} + du) = \bar{U} + V'(\bar{U}^-) du$ for $du < 0$, where $V'(\bar{U}^-)$ represents the first derivative of $V(u)$ just below \bar{U} . Then, using the fact that the distribution of du is symmetric, and $E[du|du > 0] = \sigma \frac{\sqrt{dt}}{\sqrt{2\pi}}$, we have that the benefit over \bar{U} of a consumer searching dt more attributes is $\frac{1}{2}[1 - V'(\bar{U}^-)]\sigma \frac{\sqrt{dt}}{\sqrt{2\pi}} - c dt$. That is, for dt close to zero, it is beneficial for a consumer to keep on searching as long as $V'(\bar{U}^-) < 1$. Since for $u < \bar{U}$, we have necessarily $V(u) > u$ (as otherwise it would be better to stop searching) and consequently, $V'(\bar{U}^-) \leq 1$, it must be that $V'(\bar{U}^-) = 1$. Essentially, the argument states that when the consumer’s expected utility, u , walks away from \bar{U} , in order for him to be indifferent between continuing to search and purchasing right away, the marginal decrease in $V(u)$ when u walks to the left would have to equal the marginal increase in $V(u)$ when u walks to the right.

Next, we write down the second set of boundary conditions, when u is low enough such that when u reaches \underline{U} the consumer stops searching and decides not to buy the product,

$$V(\underline{U}) = 0 \text{ and } V'(\underline{U}) = 0. \quad (7)$$

⁹This is equivalent to exercising an American put option with \bar{U} being the value of the strike price minus the stock price.

Condition (7) is similar to condition (6) for the case in which the consumer decides to stop searching without buying the product. The first part of (7) indicates that at $u = \underline{U}$, the consumer does not expect any positive utility from buying the product. The second part indicates that even if u starts to increase as the consumers continues to search, the change of $V(u)$ will be very slow. It means that the curve $V(u)$ is tangent to the x -axis at $u = \underline{U}$ (see Figure 2 below).

3.2. Optimal Stopping Rule

In order to derive the optimal stopping rule, consider first some intuition of why the consumer would stop searching when u is either high enough or low enough. When u is far above zero, the likelihood of u returning back to zero is low. Even if u comes back to zero, the consumer would have to pay substantial search costs before seeing it happen. Therefore, when u first hits \bar{U} , the consumer buys the product right away. Similarly, when u is near \underline{U} , the expected valuation is far below zero, and is not likely to return above zero again. Rather than paying more search cost for the low probability event, the consumer terminates the search process without buying the product.

Equations (5)-(7) fully determine $V(u)$:

$$V(u) = \frac{c}{\sigma^2}u^2 + \frac{1}{2}u + \frac{\sigma^2}{16c}. \quad (8)$$

From (6)-(8) we can then obtain

$$\bar{U} = -\underline{U} = \frac{\sigma^2}{4c}. \quad (9)$$

We have now derived the consumer's optimal stopping rule.

PROPOSITION 1: *The consumer searches for more product information when $-\frac{\sigma^2}{4c} < u < \frac{\sigma^2}{4c}$. The consumer stops searching and buys the product when $u \geq \frac{\sigma^2}{4c}$. The consumer stops searching without buying the product when $u \leq -\frac{\sigma^2}{4c}$.*

As implied by the proposition, if the stand-alone valuation (ex-ante valuation prior to any search) v is greater than \bar{U} , the consumer should buy the product without any search. If v is lower than \underline{U} , the consumer should exit the market also without any search. If v is between the two thresholds, the consumer would initiate search, in the hope of eventually making a purchase.

We discuss more properties of the optimal stopping rule by looking at a graphical presentation of $V(u)$ in Figure 2.

Two features of $V(u)$ are noteworthy. First, $V(0) > 0$: the consumer expects positive utility from the possibility of buying the product even when his current expected valuation of buying the product is zero. This is a consequence of the option value of the search process. By searching a consumer might find that the product provides a good fit and generates a positive expected utility, and if the product turns out to be a poor fit the consumer can always choose not to buy it.

Second, the upper and lower bounds turn out to be symmetric around zero. The upper bound being above zero and the lower bound being below zero confirms the intuition that the consumer trades off the likelihood of changing his purchase decision with the amount of search costs he needs to pay.¹⁰ This intuition explains why the starting point v does not affect the boundaries. The intuition also explains why, for example, the upper bound increases with σ . When the attributes become more important, a high u can walk back to zero in fewer “steps” and the consumer has more incentive to continue searching. On the other hand, when the search cost, c , decreases, it is easier for the consumer to gather enough information to change his purchase decision. He therefore has more incentive to continue searching, which leads to a higher \bar{U} .

3.3. Purchase Likelihood

It would be interesting to understand how the parameters of the model affect the purchase likelihood. As a next step, we write out the purchase likelihood given any u , $Pr(u)$. (The proof is in the Appendix.)

PROPOSITION 2: *The consumer’s purchase likelihood $Pr(u)$, for any $u \in [\underline{U}, \bar{U}]$, is*

$$Pr(u) = \frac{u - \underline{U}}{\bar{U} - \underline{U}} = \frac{1}{2} \left(1 + \frac{u}{\bar{U}} \right). \quad (10)$$

Proposition 2 states that when search occurs, the purchase likelihood is the ratio of the distance from u to the exit threshold, $u - \underline{U}$, and the total distance between the two thresholds, $\bar{U} - \underline{U}$. Therefore, prior to any search, the consumer buys the product with probability

$$Pr(v) = \frac{1}{2} \left(1 + \frac{v}{\bar{U}} \right) = \frac{1}{2} + \frac{2cv}{\sigma^2} \quad (11)$$

¹⁰The symmetry around zero depends on the symmetric properties of the Brownian motion. As shown below in Section 5, with discounting we no longer have the symmetry of \bar{U} and \underline{U} around zero, as with discounting the consumer would prefer to purchase earlier and not delay the product benefit, leading to \bar{U} being closer to zero than \underline{U} .

for intermediate values of v . If $v > \frac{\sigma^2}{4c}$ the consumer buys with probability one, without incurring any search costs. If $v < -\frac{\sigma^2}{4c}$ the consumer buys with probability zero and also does not incur any search costs.

One can make several interesting observations using equation (11). First, making search more informative, i.e., increasing σ , makes the consumer less likely to buy the product if $v > 0$, and more likely to buy the product if $v < 0$.

To gain intuition on this result think of the extreme case of $\sigma \rightarrow 0$. In this case, the consumer has no incentive to search and immediately buys the product if $v > 0$, and does not buy the product if $v < 0$. As σ increases, the consumer is more likely to search, hence he is more likely to gather negative information if $v > 0$, and positive information if $v < 0$. Consequently, the purchase likelihood would be reduced if $v > 0$, and increased if $v < 0$. If *a priori* a consumer has a good expectation about the value of a product, providing more information may only make it worse. On the other hand, if *a priori* a consumer has a poor expectation of the value of a product, the only possibility for that consumer to potentially want to buy the product is if the consumer receives more information.¹¹

This observation has the strategic implication that the seller who aims to maximize the purchase likelihood should increase the difficulty for the consumer to search if the product brings the consumer a positive ex-ante expected utility, and facilitate consumer search if the product brings the consumer a negative ex-ante expected utility prior to any search. The seller could, for example, decide on how much information to provide. This conclusion for the positive ex-ante utility condition can be seen as fitting real-world observations where a consumer would start researching a product, being excited and optimistic, and after quickly reading some online reviews, decide that the product does not really fit his preferences and hence forego the idea of buying it.¹² On the other hand, if consumers would not buy the product without further information (the case when the ex-ante utility is negative), the firm wants to provide information such that with some probability the information on fit gained by a consumer is positive and the consumer decides to buy the product.

Consider now the effect of search costs. Note that if the ex-ante expected utility is negative, increasing search costs makes the consumer less likely to buy the product, while if the ex-ante expected utility is positive, increasing search costs makes the consumer more likely to buy the product. The intuition is similar: if the consumer starts out with a negative ex-ante expected utility, as search becomes easier the consumer is more likely to gather positive information, hence

¹¹See Ottaviani and Prat (2001) and Johnson and Myatt (2006) for similar points without consumer search.

¹²Sun (2010) has a similar result that a seller whose product has a high vertical quality should avoid disclosing the product's horizontal location in consumers' taste space.

more likely to buy the product. These two observations suggest that, as the Internet drives down the cost of gathering product information, purchase rate may increase among consumers with a negative ex-ante expected utility and decrease among consumers with a positive ex-ante expected utility.

In terms of social welfare note that a consumer prefers lower search costs, c , and greater informativeness of the attributes checked, σ^2 . To see this one can just differentiate the expected utility of the search process, $V(v)$, with respect to c and σ^2 for $v \in (\underline{U}, \bar{U})$ to obtain that $V(v)$ is decreasing in c and increasing in σ^2 , and note that for v outside this interval the consumer can only benefit from lower search costs or a greater informativeness of the attributes checked.

3.4. Extent of Search

The number of attributes checked during the search process can vary depending on the results of the search.¹³ It is interesting to investigate what is the expected number of attributes checked. To obtain this note that the expected utility of the search process, $V(u)$, is equal to the purchase likelihood times the expected utility of buying the product when the decision to buy the product is made, \bar{U} , minus the expected costs of search:

$$V(u) = Pr(u)\bar{U} - cE[\text{number of attributes searched}|u].$$

Using the expressions above for $V(u)$ and $Pr(u)$ one can obtain that

$$E[\text{number of attributes searched}|u] = \frac{\sigma^2}{16c^2} - \frac{u^2}{\sigma^2}.$$

Note that the expected number of attributes searched is increasing in the informativeness of each attribute, σ^2 , and decreasing in the costs of search, c . Note also that the expected number of attributes searched is greater the further away the consumer is from both deciding to buy the product and stopping the search process and deciding not to buy the product, the case when u is close to zero.

4. FIRM'S PRICING DECISION

Consider now that the seller chooses a price, p , observable by consumers before they decide

¹³Given the continuous framework, the number of attributes checked means the mass of attributes checked.

whether or not to engage in the search process. The consumer's initial valuation, prior to any search, becomes $v - p$. Conceptually, it is $u - p$ that is moving now as a consumer searches. Recall from (9) that the starting point does not enter the expression of the optimal stopping boundaries. Therefore, the consumer buys the product when $u - p$ first hits \bar{U} , and exits without buying when $u - p$ first hits \underline{U} , and otherwise continues to search. In what follows we restrict attention to the case where v is homogenous across consumers, but a similar analysis could be considered with heterogenous v , and then in the expected profit one would have to integrate out across all v .

Based on (11), the ex-ante likelihood of purchase becomes

$$Pr(v - p) = \begin{cases} 1, & \text{if } v - p \geq \bar{U}; \\ \frac{1}{2} + \frac{2c(v-p)}{\sigma^2}, & \text{if } \underline{U} < v - p < \bar{U}; \\ 0, & \text{if } v - p \leq \underline{U}. \end{cases} \quad (12)$$

Given (12), and a marginal cost g , the seller maximizes her expected profit, $(p - g) \cdot Pr(v - p)$, which leads to the following result.

PROPOSITION 3: *The seller can make a positive profit if $v \geq g + \underline{U}$. The optimal price is*

$$p^* = \begin{cases} v - \frac{\sigma^2}{4c} = v - \bar{U}, & \text{if } v \geq g + 3\bar{U}; \\ \frac{v+g}{2} + \frac{\sigma^2}{8c} = \frac{1}{2}(v + g + \bar{U}), & \text{if } g + \underline{U} \leq v < g + 3\bar{U}. \end{cases} \quad (13)$$

It is intuitive that when v is high, as in the first case, the seller would rather that the consumers buy the product without any search.

When v is not that high, as in the second case in Proposition 3, the optimal price turns out to be increasing in both v , the ex-ante expected utility, and \bar{U} , the final valuation when the consumer decides to purchase the product. Note that the optimal price will move in the opposite direction with the search cost: if search were easier, the consumer has a higher purchase threshold which leads to a higher equilibrium price. Lower search costs lead to the possibility of consumers learning about their higher valuation, which leads the firm to charge higher prices.

The expression for the optimal price in the case with consumer search is noteworthy. If it were not possible for the consumer to search at all, the seller would simply charge v ; and if the seller can choose the price after the consumer decides to purchase the product, he would charge \bar{U} . The model is in between the two scenarios above: the consumer is given the option to search, and the seller needs to commit to a price before the consumers starts searching. Correspondingly, the seller takes these two prices into account in its pricing. The intuition behind this result is based on

standard monopoly pricing. The demand, or probability of buying, is proportional to $v + \bar{U} - p$, and the marginal cost is g . Therefore, profit is maximized when marginal revenue equals marginal cost, implying $p^* = \frac{1}{2}(v + g + \bar{U})$.

The ex-ante expected valuation of buying the product, or the starting point of search, now becomes

$$v - p^* = \begin{cases} \bar{U}, & \text{if } v \geq g + 3\bar{U}; \\ \frac{1}{2}(v - g - \bar{U}), & \text{if } g + \underline{U} \leq v < g + 3\bar{U}. \end{cases}$$

Consequently, the ex-ante purchase likelihood is

$$Pr(v - p^*) = \begin{cases} 1, & \text{if } v \geq g + 3\bar{U}; \\ \frac{1}{4}(1 + \frac{v-g}{\bar{U}}), & \text{if } g + \underline{U} \leq v < g + 3\bar{U}. \end{cases} \quad (14)$$

Note that whenever v is not too high, the seller chooses the price so that the consumer would conduct some search before making the purchase decision. The equilibrium purchase likelihood depends on the level of $\frac{v-g}{\bar{U}}$. Based on (13) and (14), we can then obtain the equilibrium profit:

$$\Pi(v) = \begin{cases} v - g - \bar{U}, & \text{if } v \geq g + 3\bar{U}; \\ \frac{(v-g+\bar{U})^2}{8\bar{U}}, & \text{if } g + \underline{U} \leq v < g + 3\bar{U}. \end{cases} \quad (15)$$

Naturally, profit always increases with the stand-alone utility v . By taking the derivative of $\Pi(v)$ with respect to \bar{U} , one can also see that the expected profit increases with \bar{U} if and only if $v < g + \bar{U}$. Intuitively, as \bar{U} increases the optimal price goes up and the purchase likelihood goes up or down depending on whether $v < g$ or $v > g$, respectively. When \bar{U} goes up by one unit, the price always increases by half a unit. The purchase likelihood, however, decreases faster when v is big: when v is near g , the purchase likelihood is near $1/4$ and is not affected much when \bar{U} increases by a unit; when v is high, the purchase likelihood is almost one and can be driven down significantly when \bar{U} increases by one unit.

Therefore, a profit-maximizing seller should aim to increase $\bar{U} = \frac{\sigma^2}{4c}$ when v is small, i.e., $v < g + \frac{\sigma^2}{4c}$, and decrease \bar{U} when v is big.¹⁴ In most cases, the seller cannot fully control the level of either the informativeness of search σ or the search cost c . At best, she can increase

¹⁴We consider a situation where it is increasingly costly for a seller to change the search costs or the informativeness of search, such that the change in \bar{U} does not affect the region of the parameter space where we are in terms of the comparison of v with $g + \frac{\sigma^2}{4c}$. If it is costless to change the search costs or the informativeness of search, a seller would prefer to have \bar{U} as high as possible (through lower search costs or greater informativeness of search, being in the region of v small), as in that case the expected profit converges to $\frac{\bar{U}}{8}$, while if the seller reduced \bar{U} the highest profit she could get would be $\max[v - g, 0]$.

the informativeness of search by trying to provide better answers to consumers' questions and/or making specifications of product attributes, such as color and size, vastly different. Similarly, the seller can try to decrease the search costs by, for example, circulating online video introductions of the product. We state the results above in the following proposition.

PROPOSITION 4: *The profit is increasing (decreasing) in the search costs and decreasing (increasing) in the informativeness of search if $v > (<)g + \frac{\sigma^2}{4c}$. The highest profit is obtained when $\frac{\sigma^2}{4c} \rightarrow \infty$, and, given v , the lowest profit is obtained when $\frac{\sigma^2}{4c} = \max[v - g, 0]$.*

Now consider the effects of search costs and informativeness of search on social welfare, given that the firm chooses the price to charge. To see this, first consider the effect on the consumer surplus, and let us focus the analysis on the case in which a firm chooses its price such that consumers decide to search.

An increase in search costs reduces the consumer surplus (the expected utility of the search process) for a given price, but in this case the firm may lower its price. To obtain the total effect, one can write the consumer surplus given the price charged by the firm as

$$V(v - p^*) = \frac{(v - g + \bar{U})^2}{16\bar{U}} = \frac{1}{2}\Pi(v).$$

Therefore, the comparative statistics on social welfare and consumer surplus are the same as those on the seller's profit, and therefore social welfare can potentially decrease with lower search costs and more search informativeness. Note that when the seller charges an optimal price, consumer surplus no longer always increase with search informativeness and decreases with search costs. This is because more search could lead to a higher price that the consumer has to pay.

PROPOSITION 5: *With endogenous pricing, consumer surplus decreases with lower search costs if $v \in (g + \frac{\sigma^2}{4c}, g + \frac{3\sigma^2}{4c})$.*

5. EXTENSIONS

5.1. Signals on the Value of a Product

An alternative way to model the search for information is to have the consumer receive, at each search stage, an independent signal of the value of the product to him. That is, one could have that

at any search opportunity i the consumer receives a signal S_i which is equal to the true valuation of the product for that consumer, which we can denote as U , plus an error ε_i , $S_i = U + \varepsilon_i$. The consumer does not know U but tries to infer U from the sequence of signals received plus his priors on U . Suppose that the consumer, before obtaining any signal, has a normal prior on U with mean v and variance e^2 , and that the errors ε_i are distributed normally with mean zero and variance s^2 .

We can then write the expected utility of buying the product, u , after observing t signals as,

$$u(t) = \frac{\frac{v}{e^2} + \sum_{i=1}^t \frac{S_i}{s^2}}{\frac{1}{e^2} + \frac{t}{s^2}}. \quad (16)$$

In comparison to the model above where $u(t) - u(t-1) = x_t$, we now have

$$u(t) - u(t-1) = x_t = (S_t - u(t-1)) \frac{z_{t-1}^2}{z_{t-1}^2 + s^2} \quad (17)$$

where z_{t-1}^2 is the variance of the posterior beliefs about U given the information obtained after checking $t-1$ signals. One can compute $z_t^2 = 1/(\frac{1}{e^2} + \frac{t}{s^2})$, and $z_t^2 = \frac{z_{t-1}^2 s^2}{z_{t-1}^2 + s^2}$ as the equation of motion for z_t^2 for all t . As above, the expected value of $u(t) - u(t-1)$ is zero, but the variance of $u(t) - u(t-1)$, given the information obtained until $t-1$, can now be shown to be $\sigma_t^2 = \frac{z_{t-1}^4}{z_{t-1}^2 + s^2}$, which is decreasing in the number of signals already checked, t .

In order to obtain a continuous version of this model, we can make the precision of the signals (the inverse of the variance of the signals) go to zero, such that one unit of signals allows the consumer to update his expected utility of buying the product, u , with the statistical properties of the discrete case described above. Suppose that a signal received of size dt has a variance s^2/dt , such that the update of the variance of the posterior beliefs over one unit of signals is equivalent to the update due one signal of variance s^2 . Therefore, the variance of posterior beliefs about U after obtaining t units of signals continues to be $z_t^2 = 1/(\frac{1}{e^2} + \frac{t}{s^2})$. Finally, we can obtain that the variance of the change in the expected utility of buying the product with dt signals, the variance of $u(t+dt) - u(t)$, is equal to $\sigma_t^2 dt = \frac{z_t^4}{z_t^2 + s^2/(dt)} = \frac{z_t^4}{s^2} dt$, for dt converging to zero.

One can then obtain in the continuous version that the change in expected utility of buying the product can be described as the Brownian motion

$$du = \sigma_t dw$$

where the variance of the Brownian motion decreases now with the number of signals received.

Applying the analysis above we can then obtain that the expected utility of the search process

given the expected utility of buying the product, u , and after observing t signals, has to satisfy (3). Note that now $V_t \neq 0$, that is, the expected utility of the search process depends on the number of signals checked, t , in addition to the current expected utility of buying the product, u .

Note also that now the boundary conditions also depend on the number of signals checked. That is, the upper and lower levels of u that trigger either a purchase or a no purchase decision now depend on the number of signals checked, $\bar{U}(t)$ and $\underline{U}(t)$. In particular, the boundary conditions are now represented by

$$V(\bar{U}(t), t) = \bar{U}(t) \quad (18)$$

$$V_u(\bar{U}(t), t) = 1 \quad (19)$$

$$V(\underline{U}(t), t) = 0 \quad (20)$$

$$V_u(\underline{U}(t), t) = 0. \quad (21)$$

For this case, one can show that, in general, the optimal search process is characterized by the consumer continuing to search as long as the expected utility of buying the product after observing t signals is in between $\underline{U}(t)$ and $\bar{U}(t)$. Furthermore, one can show that $\bar{U}(t)$ is decreasing, and $\underline{U}(t)$ is increasing, in the number of signals checked t . The intuition is that as the consumer checks more signals, the consumer knows that the new signal is going to be less informative on the expected utility of buying the product and, therefore, the consumer requires less in terms of the expected utility of buying the product in order to either stop searching and buying the product, or to stop searching and deciding not to buy the product. Signals become less informative as t increases. As the early signals are the most informative, the consumer starts out being hopeful, and loses hope quickly if the signals turn out to be unfavorable. As t approaches infinity, the signals are hardly informative, and the consumer makes a decision even if the expected utility of buying is relatively close to zero. Furthermore, as the variance of the posterior beliefs approaches zero at a decreasing rate, we can obtain that the purchase and exit thresholds also approach zero at a decreasing rate. That is, the purchase threshold, $\bar{U}(t)$, is convex, and the exit threshold, $\underline{U}(t)$, is concave, both approaching zero when $t \rightarrow \infty$. Figure 3 presents the boundaries $\bar{U}(t)$ and $\underline{U}(t)$ that determine the optimal search behavior for an example with $c = 0.01$ and $s = e = 1$.

5.2. Finite Mass of Attributes

In the base model above it was also considered that there was an infinite mass of attributes, such that the consumer, after observing t attributes, was in the same conditions in terms of attributes

that he could check as if he had observed t' attributes for any t and t' . This assumption is motivated by the fact that for many product categories there is a huge amount of information available online and a consumer never needs to worry about running out of attributes to check. For products that are relatively simple, it may be true that they have relatively fewer attributes. To incorporate such situations we extend our baseline model in this section to incorporate a finite number of attributes.¹⁵

The analysis of this situation is similar to the one considered above leading to the partial differential equation (3). However, in this case we do not have $V_t = 0$, as the situation of the consumer changes when the number of attributes checked is different. That is, after checking more attributes, there are fewer attributes that the consumer can check. Let T represent the mass of attributes that the consumer can check.

To solve this problem we have again that the boundary conditions on u are functions of the number of attributes checked, and satisfy (18)-(21), as in the subsection above. Furthermore, we have an additional boundary condition that

$$V(u, T) = \max[u, 0]. \tag{22}$$

One can obtain that $\bar{U}(t)$ is decreasing, and $\underline{U}(t)$ is increasing, in the number of attributes checked, t . That is, as the number of attributes checked increases, the consumer is less demanding on the expected utility of buying the product, u , to decide to either stop searching and buy the product, or stop searching and not buy the product. Furthermore, one can obtain that $\bar{U}(T) = \underline{U}(T) = 0$. That is, after checking all attributes the consumer decides to buy the product even if the utility of buying the product is barely above zero. Note also that each attribute checked affects the expected utility of buying the product in the same way, such that the purchase and exit thresholds do not change too much for t small, but start changing more as t approaches T . That is, one can show that in this case the purchase threshold, $\bar{U}(t)$, is concave, and the exit threshold, $\underline{U}(t)$, is convex. Figure 4 presents the boundary conditions that characterize the optimal consumer search behavior for an example with $\sigma = T = 1$ and with $c = .01$ and $c = 0.1$. The figure also illustrates that the search region enlarges as search costs decrease.

¹⁵See also the two-sided case in Roberts and Weitzman (1981) for the case with direct specific assumptions on how the posteriors change through search in a setting of research and development for a project.

5.3. Discounting

In the model above, search was considered to be done relatively quickly so that issues of discounting the future payoffs because of the time duration of search were not important. However, in some search problems the delay in obtaining the potential utility of purchasing the product could actually matter. For such problems it is important to discount the future potential payoff of purchasing the product.

To consider this case, equation (1) becomes now

$$V(u, t) = -c dt + e^{-r dt} EV(u + du, t + dt),$$

where r is the continuous time interest rate. Doing the Taylor expansion on this equation as above we have that (3) now becomes

$$rV = -c + V_t + \frac{\sigma^2}{2} V_{uu}. \quad (23)$$

Noting that the problem does not change with time, we again have $V_t = 0$. We can then obtain, by solving the differential equation (23), that $V(u)$ satisfies

$$V(u) = A_1 e^{\frac{\sqrt{2r}}{\sigma} u} + A_2 e^{-\frac{\sqrt{2r}}{\sigma} u} - \frac{c}{r}, \quad (24)$$

where A_1 and A_2 are two constants that are obtained together with \bar{U} and \underline{U} with the conditions (6) and (7). The derivation of A_1, A_2, \bar{U} , and \underline{U} is presented in the Appendix. In comparison with the no discounting case, one can also obtain the following result (proof is in the Appendix).

PROPOSITION 6: *With discounting, the purchase threshold (\bar{U}) is smaller than the absolute value of the exit threshold (\underline{U}).*

That is, with discounting we have $\bar{U} < -\underline{U}$. The intuition is that with discounting the future benefits of searching are smaller, and therefore the consumer wants to stop searching sooner: \bar{U} is lower, and \underline{U} is higher. This effect is more important when the consumer is more likely to buy the product, i.e., when $u > 0$, as there is a clear benefit of buying that is being discounted, and therefore the consumer does not need as much positive information about the product to decide to stop searching and buy the product. That is, with discounting \bar{U} goes down by more than \underline{U} goes up. While the consumer is always trading off search cost with the potential benefit from buying the product, the discount rate is applied to an immediate payoff \bar{U} when u is near the purchase threshold, and it is applied to more distant payoffs when u is near the exit threshold. An interesting

observation is that an ex-ante indifferent consumer with $v = 0$ buys the product with probability of .5 in the no discounting case, but buys the product with a higher probability if he discounts the future. As an example, consider the case where $\sigma^2 = .5$, $c = .1$, and $r = .1$. Without discounting we have $\bar{U} = -\underline{U} = 1.25$. With discounting with the interest rate $r = .1$, we have $\bar{U} = .87$ and $\underline{U} = -1.09$.

5.4. Choosing the Search Intensity

When the search for information takes time, a consumer may often be able to choose how many attributes to search per unit of time, with more attributes being searched per unit of time being more costly. For example, a consumer could decide not to search actively, and simply obtain information that comes to him without effort or, alternatively, to search intensively for information. In terms of the model above one can then endogenize σ^2 , the informativeness of the search process at each moment in time. Greater informativeness, greater σ^2 , would mean that the consumer incurs greater search costs, c , which we now write as $c(\sigma^2)$, as the search costs depend on σ^2 , with $c'(\sigma^2), c''(\sigma^2) > 0$. Consider first the case with no discounting.

From (2) and (3), maximizing $V(u)$ with respect to σ^2 leads to

$$c'(\sigma^2) = \frac{1}{2}V_{uu}. \quad (25)$$

Putting this equation together with (5) one obtains the condition that determines σ^2 as

$$\sigma^2 c'(\sigma^2) = c(\sigma^2). \quad (26)$$

Note that for this no discounting case we obtain that the optimal intensity of search, σ^2 , is independent of the expected valuation of buying the product, u . For example, for the case where $c(\sigma^2) = a_0 + \frac{a_1}{2}\sigma^4$, we obtain the optimal $\sigma^2 = \sqrt{\frac{2a_0}{a_1}}$.

With discounting, we have $\sigma^2 V_{uu}$ increasing in u , by (23), and given that $V(u)$ is increasing in u . Then, by (23) and (25) we have that the optimal σ^2 is increasing in u . That is, with discounting, a consumer invests more in search intensity the more likely he is to buy the product, which occurs when u is greater. The closer the consumer is to buying the product the more he invests in learning about the product, as the benefits of purchase are discounted.

6. CONCLUDING REMARKS

We have studied two central topics in this paper. First, how does a consumer optimally search for product information when his valuation of the product is uncertain? Second, given the consumer's optimal search strategy, what can the seller do to maximize his expected profit? We provide a parsimonious model of consumer search for gradual information that captures the benefits and costs of search, determining optimally the thresholds when the consumer decides to stop searching.

In summary, the consumer searches for further information whenever his valuation of the product is between two boundaries. When his valuation first hits the upper bound, he stops searching and buys the product. When his valuation first hits the lower bound, he stops searching without buying the product. The two bounds are determined by the informativeness and marginal cost of search. In particular, when search becomes more informative and less costly, the purchase threshold, or the upper bound, increases.

As for the seller, when the stand-alone utility is not too high, the optimal price that she should set takes into account both the stand-alone valuation (expected utility before search) of the product and the upper-bound valuation (expected utility after search given that the consumer purchases). For a high stand-alone expected utility, the seller chooses the price such that the consumer's net utility prior to search is exactly \bar{U} , hence the consumer buys the product immediately.

While providing an understanding of consumer behavior and implications for a firm's marketing strategies, the paper presents a tractable foundation to understand information-search behavior in a stochastic environment. From this perspective, one can extend the model in several directions. It would be interesting to investigate the seller's decision on how much information to provide. It would also be interesting to know what happens when the attributes can be ordered by their importance, and consider the order of the product attributes that the seller would choose to present, and how the consumer's purchase threshold may evolve over the search process.

APPENDIX

PROOF OF PROPOSITION 2: To obtain (10), we can see that

$$\begin{aligned}
 Pr(u) &= Pr(u \text{ hits } \bar{U} \text{ first} | u) \\
 &= Pr(u \text{ hits } \bar{U} \text{ first} | u + z') \cdot \frac{1}{2} + Pr(u \text{ hits } \bar{U} \text{ first} | u - z') \cdot \frac{1}{2} \\
 &= \frac{1}{2}Pr(u + z') + \frac{1}{2}Pr(u - z'), \tag{i}
 \end{aligned}$$

where $z' = E(du | du > 0) = \sigma \frac{\sqrt{dt}}{2\pi}$. From (i), $Pr(u + z') - Pr(u) = Pr(u) - Pr(u - z')$. Therefore, $Pr(u)$ is linear in u . Since $Pr(\underline{U}) = 0$ and $Pr(\bar{U}) = 1$, $Pr(u) = \frac{u - \underline{U}}{\bar{U} - \underline{U}}$.

DERIVATION OF EQUILIBRIUM CONSTANTS IN DISCOUNTING CASE: To obtain A_1, A_2, \bar{U} , and \underline{U} for the discounting case we can write conditions (6) and (7) as

$$A_1 e^{\frac{\sqrt{2r}\bar{U}}{\sigma}} + A_2 e^{-\frac{\sqrt{2r}\bar{U}}{\sigma}} - \frac{c}{r} = \bar{U}; \tag{ii}$$

$$A_1 \frac{\sqrt{2r}}{\sigma} e^{\frac{\sqrt{2r}\bar{U}}{\sigma}} - A_2 \frac{\sqrt{2r}}{\sigma} e^{-\frac{\sqrt{2r}\bar{U}}{\sigma}} = 1; \tag{iii}$$

$$A_1 e^{\frac{\sqrt{2r}\underline{U}}{\sigma}} + A_2 e^{-\frac{\sqrt{2r}\underline{U}}{\sigma}} - \frac{c}{r} = 0; \tag{iv}$$

$$A_1 \frac{\sqrt{2r}}{\sigma} e^{\frac{\sqrt{2r}\underline{U}}{\sigma}} - A_2 \frac{\sqrt{2r}}{\sigma} e^{-\frac{\sqrt{2r}\underline{U}}{\sigma}} = 0. \tag{v}$$

Solving (ii)-(v) together gives the solution to A_1, A_2, \bar{U} , and \underline{U} . Denoting

$$\bar{X} = e^{\frac{\sqrt{2r}\bar{U}}{\sigma}} \tag{vi}$$

and

$$\underline{X} = e^{\frac{\sqrt{2r}\underline{U}}{\sigma}}, \tag{vii}$$

from (iv) and (v) one can then obtain $A_1 = \frac{c}{2r\underline{X}}$ and $A_2 = \frac{c\underline{X}}{2r}$. Using these in (iii), and denoting

$$Z = \frac{\bar{X}}{\underline{X}}, \tag{viii}$$

one can then obtain

$$Z = \frac{\sigma\sqrt{2r}}{2c} + \sqrt{1 + \frac{\sigma^2 r}{4c^2}}. \tag{ix}$$

Using these in (ii) one can obtain

$$\log \bar{X} = (Z - 1) \frac{2c}{\sigma \sqrt{2r}} - 1, \quad (\text{x})$$

from which one can obtain immediately \bar{U} . From the expressions presented above it is then straightforward to obtain \underline{U} , A_1 , and A_2 .

PROOF OF PROPOSITION 6: To show that $\bar{U} < -\underline{U}$ we can just show that $Z > \bar{X}^2$, where $Z = \frac{\bar{X}}{\underline{X}}$ as defined in (viii), and \bar{X} and \underline{X} are defined by (vi) and (vii), respectively. Given that Z is computed to be (ix) and \bar{X} is obtained in equilibrium to satisfy (x), this inequality can be written as

$$\omega + \sqrt{1 + \omega^2} > e^{\frac{2\omega}{1 + \sqrt{1 + \omega^2}}} \quad (\text{xi})$$

where $\omega = \frac{\sigma \sqrt{2r}}{2c}$. Taking the logarithm (an increasing function) on both sides of (xi) one obtains

$$\log(\omega + \sqrt{1 + \omega^2}) > \frac{2\omega}{1 + \sqrt{1 + \omega^2}}. \quad (\text{xii})$$

Noting now that the left hand side of (xii) is equal to the right hand side of (xii) for $\omega = 0$ (also the result that $\bar{U} = -\underline{U}$ for $r = 0$), and that the derivative of the left hand is greater than the derivative of the right hand side for $\omega > 0$, we have that $\bar{U} < -\underline{U}$ for all $r > 0$.

REFERENCES

- ARMSTRONG M., J. VICKERS, AND J. ZHOU (2009), “Prominence and Consumer Search,” *RAND Journal of Economics*, **40**, 209-233.
- ANDERSON, S.P., AND R. RENAULT (1999), “Pricing, Product Diversity, and Search Costs: A Bertrand-Chamberlin-Diamond Model,” *RAND Journal of Economics*, **30**, 719-735.
- BAKOS, J. Y. (1997), “Reducing Buyer Search Costs: Implications for Electronic Marketplaces,” *Management Science*, **43**, 1676-1692.
- BERGEMANN, D., AND J. VÄLIMÄKI (1996), “Learning and Strategic Pricing,” *Econometrica*, **64**, 1125-1149.
- BERGEMANN, D., AND J. VÄLIMÄKI (2000), “Experimentation in Markets,” *Review of Economic Studies*, **67**, 213-234.
- BOLTON, P., AND C. HARRIS (1999), “Strategic Experimentation,” *Econometrica*, **67**, 349-374.
- CHEN, Y., AND K. SUDHIR (2004), “When Shopbots Meet Emails: Implications for Price Competition on the internet,” *Quantitative Marketing and Economics*, **2**, 233-255.
- COX, J.C., S.A. ROSS, AND M. RUBINSTEIN (1979), “Option Pricing: A Simplified Approach,” *Journal of Financial Economics*, **7**, 229-263.
- DE LOS SANTOS, B., A. HORTACSU, AND M.R. WILDENBEEST (2010), “Testing Models of Consumer Search using Data on Web Browsing and Purchasing Behavior,” *working paper*, University of Chicago.
- DIAMOND, P. (1971), “A Model of Price Adjustment,” *Journal of Economic Theory*, **3**, 156-168.
- DIXIT, A. (1993), *The Art of Smooth Pasting*, Harwood Academic Publishers: Chur, Switzerland.
- DUMAS, B. (1991), “Super Contact and Related Optimality Conditions,” *Journal of Economics, Dynamics, and Control*, **15**, 675-685.
- ELLISON, G., AND D. FUDENBERG (1995), “Word-of-mouth Communication and Social Learning,” *Quarterly Journal of Economics*, **110**, 93-125.
- GAL-OR, E., T. GEYLANI, AND A. DUKES (2008), “Information Sharing in a Channel with Partially Informed Retailers,” *Marketing Science*, **27**, 642-658.
- HAUSER, J.R., G.L. URBAN, AND B.D. WEINBERG (1993), “How Consumers Allocate Their Time When Searching for Information,” *Journal of Marketing Research*, **30**, 452-466.
- HAUSER, J.R., AND B. WERNERFELT (1990), “An Evaluation Model of Consideration Sets,” *Journal of Consumer Research*, **16**, 393-408.

- HONKA, E. (2010), “Quantifying Search and Switching Costs in the U.S. Auto Insurance Industry,” *working paper*, University of Chicago.
- JOHNSON, J.P., AND D.P. MYATT (2006), “On the Simple Economics of Advertising, Marketing, and Product Design,” *American Economic Review*, **96**, 756-783.
- KUKSOV, D. (2004), “Buyer Search Costs and Endogenous Product Design,” *Marketing Science*, **23**, 490-499.
- KUKSOV, D. (2006), “Search, Common Knowledge, and Competition,” *Journal of Economic Theory*, **130**, 95-108.
- LAL, R., AND C. MATUTES (1994), “Retail Pricing and Advertising Strategies,” *Journal of Business*, **67**, 345-370.
- MEHTA, N., S. RAJIV, AND K. SRINIVASAN (2003), “Price Uncertainty and Consumer Search: A Structural Model of Consideration Set Formation,” *Marketing Science*, **22**, 58-84.
- MERTON, R.C. (1990), *Continuous-Time Finance*, Basil Blackwell: Oxford, U.K.
- MOORTHY, S., B.T. RATCHFORD, AND D. TALUKDAR (1997), “Consumer Information Search Revisited: Theory and Empirical Analysis,” *Journal of Consumer Research*, **23**, 263-277.
- NARASIMHAN, C. (1988), “Competitive Promotional Strategies,” *Journal of Business*, **61**, 427-449.
- OFEK, E., AND V. SRINIVASAN (2002), “How Much Does the Market Value an Improvement in a Product Attribute?,” *Marketing Science*, **21**, 398-411.
- OTTAVIANI, M., AND A. PRAT (2001), “The Value of Public Information in Monopoly,” *Econometrica*, **69**, 1673-1683.
- RAO, R., AND N. SYAM (2001), “Equilibrium Price Communication and Unadvertised Specials by Competing Supermarkets,” *Marketing Science*, **20**, 61-81.
- ROBERTS, K., AND M.L. WEITZMAN (1981), “Funding Criteria for Research, Development, and Exploration Projects,” *Econometrica*, **49**, 1261-1288.
- SHUGAN, S.M. (1980), “The Cost of Thinking,” *Journal of Consumer Research*, **7**, 99-111.
- STAHL II, D. (1989), “Oligopolistic Pricing with Sequential Consumer Search,” *American Economic Review*, **79**, 700-712.
- SUN, M. (2010), “Disclosing Multiple Product Attributes,” *Journal of Economics and Management Strategy*, forthcoming.
- VARIAN, H. (1980), “A Model of Sales,” *American Economic Review*, **70**, 651-659.
- VILLAS-BOAS, J.M. (2004), “Consumer Learning, Brand Loyalty, and Competition,” *Marketing Science*, **23**, 134-145.

- VILLAS-BOAS, J.M. (2009), "Product Variety and Endogenous Prices with Evaluation Costs," *Management Science*, **55**, 1338-1346.
- WEITZMAN, M.L. (1979), "Optimal Search for the Best Alternative," *Econometrica*, **47**, 641-654.
- WERNERFELT, B. (1994), "Selling Formats for Search Goods," *Marketing Science*, **13**, 298-309.
- WOLINSKY, A. (1986), "True Monopolistic Competition as a Result of Imperfect Information," *The Quarterly Journal of Economics*, **101**, 493-512.
- ZHANG, J. (2010), "The Sound of Silence: Observational Learning in the U.S. Kidney Market," *Marketing Science*, **29**, 315-335.

Figure 1: An Example of $u(t)$

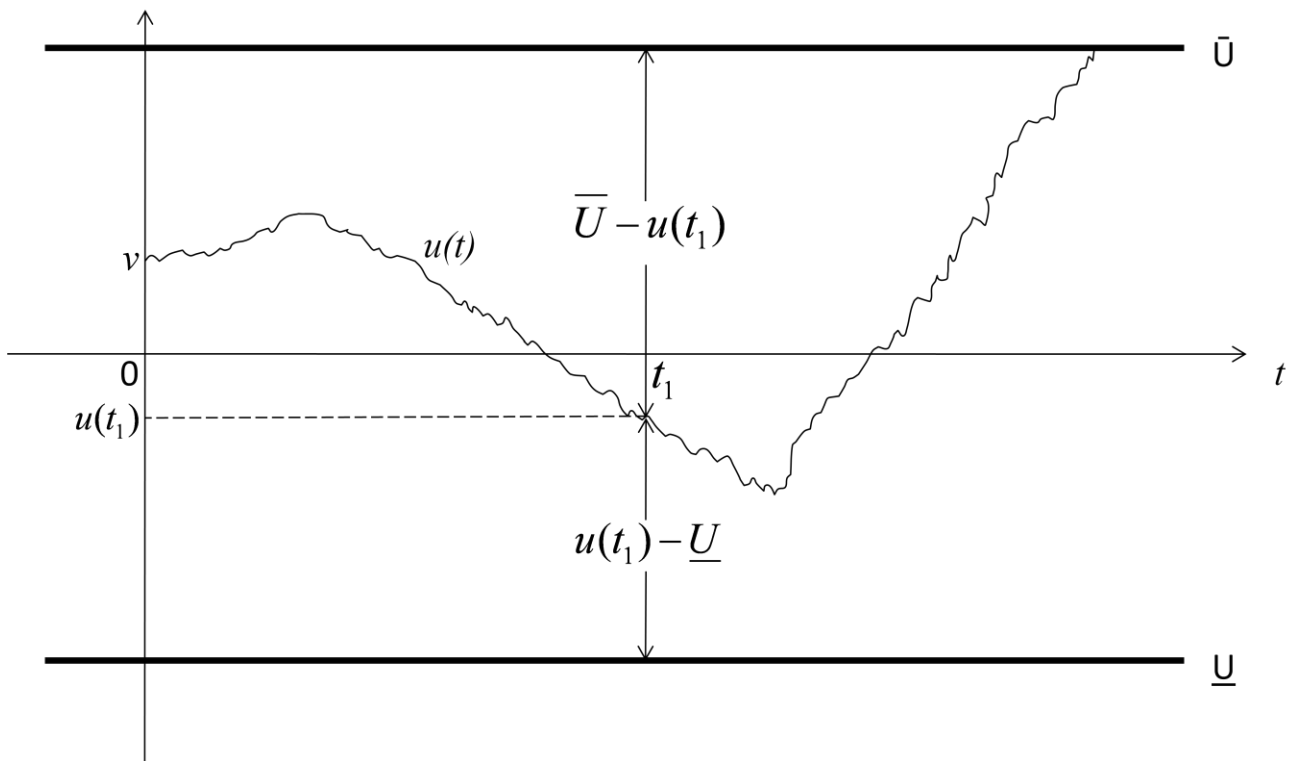


Figure 2: Graph of $V(u)$

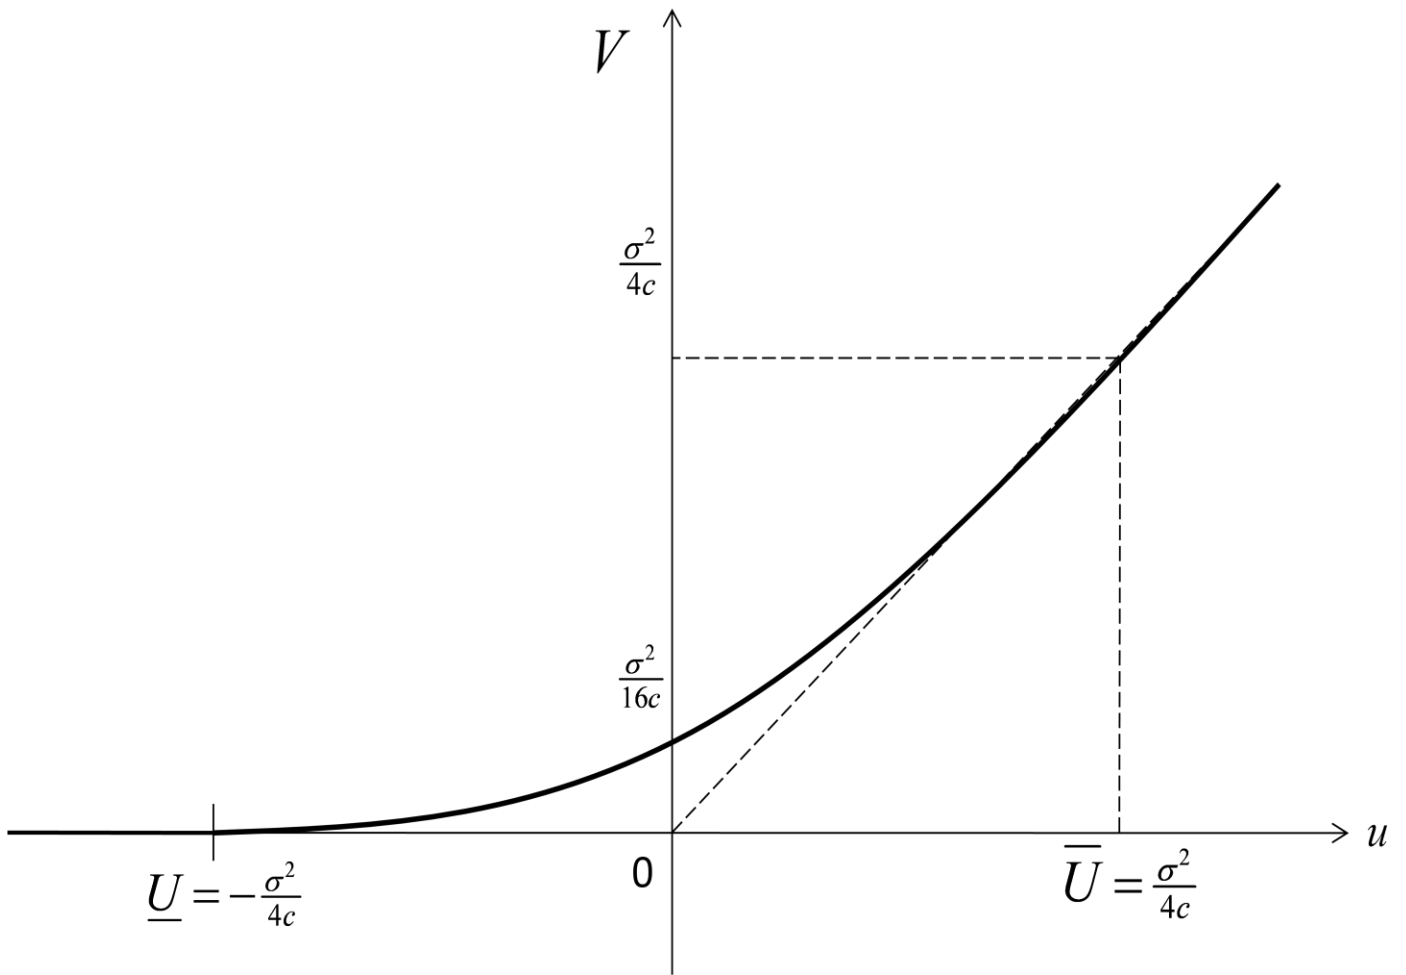


Figure 3: Stopping Boundaries with Independent Signals

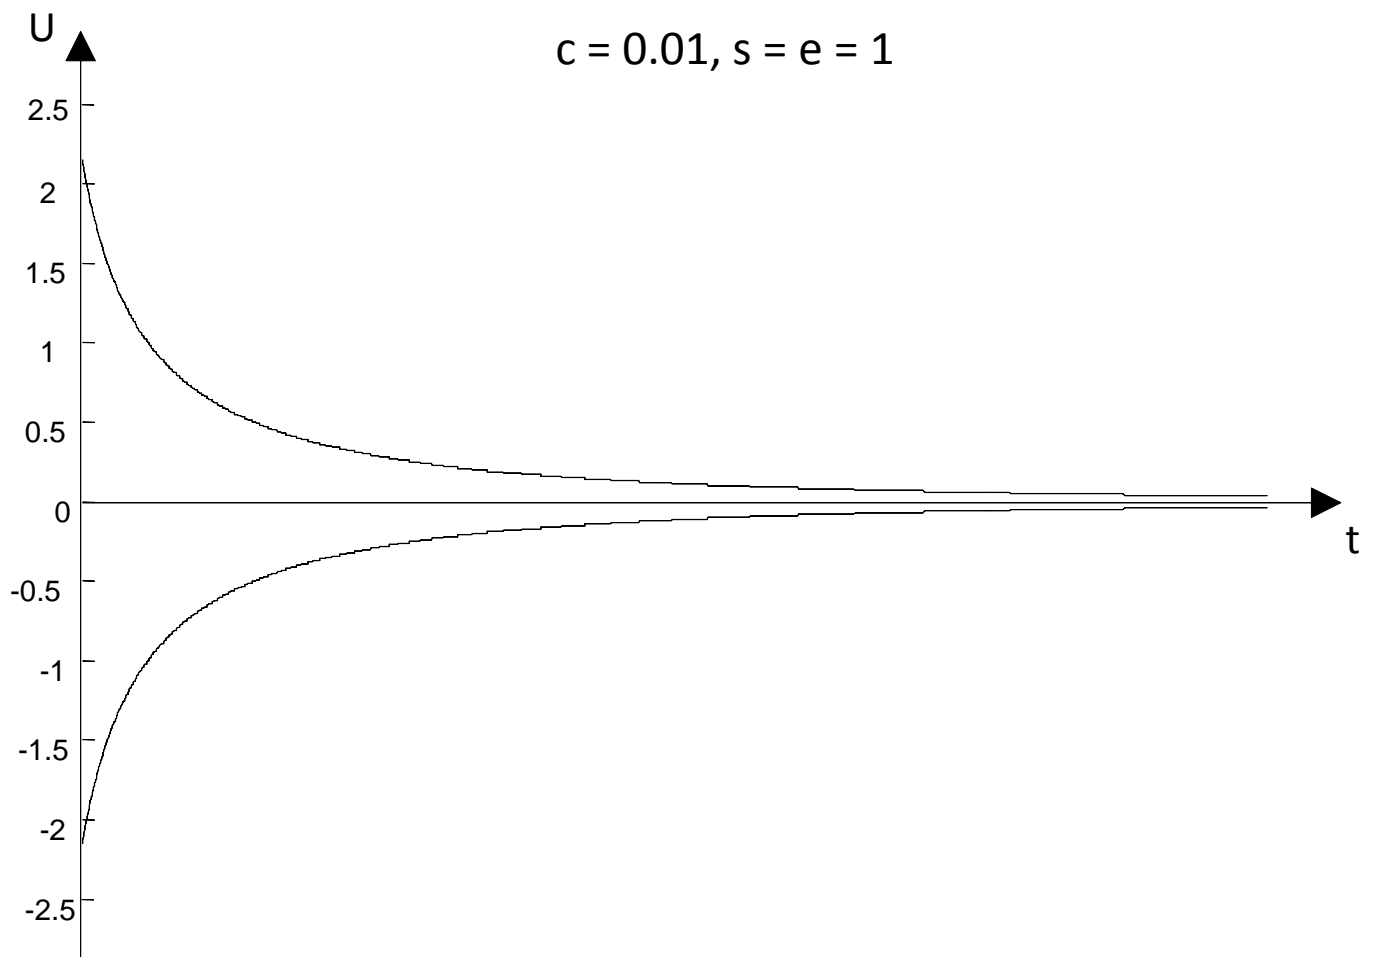


Figure 4: Stopping Boundaries with Finite Mass of Attributes

