

Price Wars and Niche Discovery in an Information Economy

Christopher H. Brooks

Edmund H. Durfee*

Rajarshi Das †

ABSTRACT

Electronic goods are flexible and have negligible marginal costs. These features allow a producer of electronic goods to explore pricing schemes, and in particular bundling, that would not be feasible with physical goods. However, they can also make it more difficult for a producer to differentiate itself from competitors offering identical goods. Previous research in this area indicates that in markets where producers compete over the sale of identical information goods, cyclical price wars often develop.

In this paper, we provide a characterization of the conditions that result in price wars and show analytically how the existence of niches within the consumer population can lead duopolist producers to each target separate niches and avoid price wars. In situations where producers have incomplete information about consumer preferences, and so must learn a strategy, producers will be concerned not only with the relative benefits of niche targeting as opposed to a price war, but also with the loss of revenue incurred in discovering a niche. We present experimental evidence describing how the difficulty of locating niches changes for an entrant and an incumbent producer as both niche size and the producers' learning strategies are varied. Our results provide support for the idea that information goods producers are not doomed to engage in price wars at all times, but instead can coexist without direct competition under certain market conditions.

1. INTRODUCTION

The continued growth and development of electronic com-

*Artificial Intelligence Laboratory, University of Michigan, Ann Arbor, MI 48109, {chbrooks,durfee}@umich.edu

†Institute For Advanced Commerce, IBM Research, PO Box 704, Yorktown Heights, NY 10598, rajarshi@watson.ibm.com

⁰This work was supported in part by an IBM University Partnership Grant and by the National Science Foundation under grant IIS-9872057.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'00, October 17-20, 2000, Minneapolis, Minnesota.

Copyright 2000 ACM 1-58113-272-7/00/0010 ..\$5.00

merce on the Internet is producing an environment in which automated agents will undoubtedly play a significant role, both in the sale and purchase of goods. Electronic marketplaces that allow agents to conduct commerce have been proposed by several researchers [7, 9, 17]. As autonomous computational agents command a greater presence within electronic markets, they will have a greater impact on the collective behavior exhibited by these markets. Autonomous agents have the ability to make rapid decisions over complex outcomes, but they have the disadvantage of not typically applying human-type heuristic decision-making. Therefore, it is important for us to understand both how to construct these agents and marketplaces and also how to make predictions regarding the aggregate behavior of the markets and the individual behavior of the agents within them. This will allow marketplace designers to construct mechanisms which promote desirable outcomes and allow agent designers to produce agents with desirable sets of behaviors.

In our previous research, we have studied the role of agents in information economies [3, 5]. We are explicitly interested in the ways in which information goods change the traditional producer-consumer problem. Information goods differ from physical goods in that producers of information goods have a large fixed cost to produce a single instance of a good, but negligible marginal cost, meaning that once the good exists, it may be replicated and resold virtually for free. This allows information goods producers to consider new sorts of pricing strategies, in particular selling goods in a bundle, that would not be feasible with physical goods. In our past work [5], we studied the effectiveness of different pricing schemes for a monopolist producer of information goods. We found that a producer that must learn about the preferences of the consumer population is better off choosing a moderately complex schedule, even though schedules with more parameters yield higher profit once learned. Ultimately, the cost of learning can outweigh the increase in profit. This underscores another of the points driving our research: it is important to understand the characteristics of the equilibrium behavior of an information economy in order to make long-term predictions, but knowledge about equilibrium behavior alone is not sufficient. If we wish to understand how information economies behave (and be able to predict or prescribe the behavior of agents within an economy), we also need to study the nonequilibrium dynamics.

A phenomenon that often occurs in the study of multi-producer information economies is the emergence of price wars, where producers sequentially undercut each other. The work of Kephart, et al. contains several examples of this

phenomenon [12, 13]. The addition of bundling makes price wars even more likely by inhibiting producers' ability to differentiate themselves from each other. If competing producers are able to bundle together additional goods without incurring extra marginal expense, then each producer has an incentive to bundle together all possible goods. Producers can then distinguish themselves from each other only through the prices they charge for their bundles of information goods.

The occurrence of price wars among computational producer agents is an unappealing result for at least two reasons. First, competing producers end up profiting less, since the price war drives prices down to marginal cost. Second, price wars of this nature, in which firms compete until all profit is gone, seem not to occur very often in the real world. Rather than engaging in a price war, firms appear to differentiate themselves from other firms so as to avoid this sort of destructive competition. In this paper, we describe system characteristics which lead to price wars and identify sufficient conditions for the formation of *niches*, or separate components of the consumer population. We consider an environment with two producers and two niches and describe the situations in which a producer is better off targeting a particular niche and avoiding competition, rather than competing directly for the entire population. We begin with an analytical approach that describes when duopolist producers with perfect information should pursue niches. Even in this very simple environment, there are cases in which producers are better off avoiding direct competition. Following this, we consider producers with imperfect knowledge of the preferences of the consumer population. We use simulation to examine the aggregate profits acquired by learning producers as they target separate niches. This allows us to explicitly study the nonequilibrium profits of these producers and determine how easy or difficult it is for learning producers to avoid a price war. In addition to the relative benefits of niche targeting as opposed to a price war, we show that producers that have incomplete information about the market will be concerned with the loss of revenue incurred in discovering a niche. The result of this work is a demonstration that there exist cases in information economies where price wars need not occur; in some instances, producers are better off specializing and targeting particular niches and avoiding the cost of competing as generalists, even if the niches may appear less profitable at first glance.

The remainder of the paper is organized as follows: Section 2 presents the details of our model, along with an analytical description of the conditions sufficient for an alternative to price wars. Section 3 presents experimental results showing how easy or difficult it is for producers with imperfect information to discover niches as niche size and learning behavior vary, and then discusses the features of the problem that produce these results. Finally, section 4 concludes and provides some future directions for our research.

2. NICHES AND PRICE WARS

In this section, we begin by discussing the phenomena which lead to price wars. We then present our model of consumer valuation and describe the introduction of clutter costs. Finally, we present conditions under which price wars will not occur if information producers are rational and have perfect knowledge.

2.1 Price Wars

It is well known that price wars occur in a duopoly when producers are selling identical products and are competing with each other only in terms of price. The underlying reason is that, since the goods offered are identical, a producer can capture the entire market by slightly undercutting its competitor's price. The Nash equilibrium strategy is for both producers to price their goods at marginal cost. If producers have perfect information and set prices simultaneously, then this will behave as a Bertrand game and both producers will immediately price at marginal cost. If producers are myopic or have incomplete information, they will iteratively lower their prices until one reaches marginal cost. At this point, each producer is indifferent as to whether to offer articles at all. If producers have different marginal costs, the producer with the higher marginal cost will drop out of the war rather than sell at a loss. Once one producer remains, it can then raise its price, since it is no longer facing competition. The losing producer can then re-enter at the next iteration, setting its price just below that of the incumbent producer and beginning another cycle of the price war. In general, there are two components of the problem that must be present in order for this to occur: the producers must be unable to differentiate themselves from each other (other than by merely charging a lower price) and the consumer population must be homogeneous in their preferences.

One question that has been raised within the context of an information economy is whether it is reasonable to assume that competing producers offer identical goods. Varian [18] argues that, in fact, producers typically offer differentiated products which are partially substitutable. In fact, this is the model adopted and developed by Bakos and Brynjolfs-son [2]. This is certainly the case for *content producers*, such as The New York Times or the Washington Post. These firms can offer similar content, but their specific articles are copyrighted and cannot be offered by competitors. However, there are also cases in which *content resellers* are able to offer identical goods. An example of this is newspapers that reprint articles purchased from the AP or UPI newswires; in this case, each paper has access to the same set of articles, and so they compete primarily on price. (Newspapers also derive a significant portion of their revenue from advertising, which leads them to target consumer niches that appeal to their advertisers. The same need to distinguish oneself from competing newspapers applies.)

Even if they are offering identical information goods, producers can still differentiate themselves through the use of different price schedules. For example, one producer might offer articles using a linear pricing schedule, where users pay a fixed price per article. This tends to appeal to consumers that are only interested in a small number of articles. Another producer might choose to offer its articles in a bundle, where a consumer pays a flat fee for access to the entire collection. This tends to appeal to consumers that value a large number of articles. If consumers must pay a cost to process each article purchased, then producers can use these pricing schedules to differentiate themselves. The details of this are discussed in section 2.2. For example, the New York Times sells the current day's edition as a bundle, under the assumption that a consumer may want to read a large number of articles, but no single article is highly valued. However, it sells access to archived articles on a per-article basis, un-

der the assumption that consumers of archived material are looking for a small number of particular articles that they value highly.

In addition, in order for price wars to be inhibited, there must be differentiation amongst the preferences of the consumer population. If all consumers make the same decision with regard to competing alternatives, then each producer's best strategy is adopt the relevant price schedule to undercut its opponents; only the producer with the best price will receive a positive profit. Moreover, the heterogeneous segments of the consumer population must be sufficiently large or valuable to appeal to producers. A single consumer with tastes that differ from the remainder of the population is unlikely to dissuade competing producers from a price war.

It should be noted that information goods producers can also avoid a price war through explicit collusion. Producers could agree to keep prices artificially high so as to avoid a price war. This requires a degree of trust on the part of each producer, as in an iterated Prisoner's Dilemma game [1]. Each producer cooperates so as to ensure the future cooperation of others. While we do not examine collusion specifically in this paper, a producer's decision as to whether to collude with another and share a niche rather than targeting a separate niche is analogous to the decision as to whether to engage in a price war (see section 2.5); in both cases, the producer chooses whether to share a consumer population or not based upon the expected profit.

A feature of information goods that increases the tendency toward price wars is their low marginal cost. As discussed previously, information goods are extremely cheap to reproduce, which makes pure bundling an effective strategy. If consumers are able to dispose of unwanted articles without cost, then pure bundling becomes at least as effective as per-article pricing (see, for example, Fay and MacKie-Mason [10]), making it more difficult for producers to differentiate themselves. However, the idea that a consumer can freely determine its valuation of a good or dispose of unwanted goods does not seem realistic. Consumers typically must spend some time or effort to determine if an information good is valuable, and must pay to download or store each good. To solve this problem, we introduce the idea of a *clutter cost*. This is a cost that is incurred by a consumer for each article it purchases. It can be thought of as a storage or bandwidth cost, or the opportunity cost of time paid by a consumer to assess an article's value. A consumer must spend some effort or time to assess the value of an information good. For example, if a consumer must spend two minutes reading the abstract of an article to determine its value, the article's worth is decreased by the value of those two minutes to the consumer. If this article turns out to be of little value, then the consumer has wasted the time spent evaluating this article. In addition to being realistic, clutter cost serves to weaken the advantage of bundling. If consumers do not pay some sort of clutter cost, then producers have a strong incentive to bundle together large numbers of articles. Since a consumer will never receive a negative valuation for an article, a producer always has an incentive to add articles to a bundle. Clutter cost emphasizes the distinction between consumers that prefer a large number of articles and consumers that prefer a small number of articles by tying the value of each article to the cost paid by the consumer to process a bundle of articles.

In summary, in order for price wars to be avoided, pro-

ducers must be able to differentiate themselves from each other other than through price. In addition, the consumer population must contain structures that make it rational for producers to differentiate themselves. In the following section, we describe these conditions in the context of our model.

2.2 Consumer Model

This work uses a model of consumer preferences based upon the one originally proposed by Chuang and Sirbu [8]. In this model, each producer offers the same N information goods. A consumer i has a positive valuation on a fraction k_i of these goods. Consumer i values its most-preferred article at w_i . If articles are then sorted in order of decreasing value (from indices 1 to N), the consumer valuation of the j th article (denoted $\mu(j)$) is given by:

$$\mu(j) = \begin{cases} w_i(1 - \frac{j-1}{k_i N}) & \text{if } j - 1 \leq k_i N \\ 0 & \text{if } j - 1 > k_i N \end{cases} \quad (1)$$

In addition, we introduce a clutter cost, denoted by the function $\alpha(|B|)$, which gives the cost incurred by a consumer for consuming a bundle of articles of size $|B|$. Therefore, the net value $V(B)$ of a bundle B is:

$$V(B) = \left(\sum_{j \in B} \mu(j) \right) - \alpha(|B|) \quad (2)$$

If a producer offers its goods as a bundle, the consumer will simply compute the function $V(B)$ for the bundle and purchase it if $V(B)$ is greater than or equal to the price. If the goods are offered on a per-article basis, then the consumer will calculate $V(B)$ for a B of size 0, 1, 2, and so on up to N . A consumer will then purchase the number of articles (0 through N) which maximizes its surplus, where surplus is $V(B)$ less the amount paid.

2.3 The Existence of Niches

Given the consumer valuation described above, we now turn to the question of when it is preferable for producers to target separate niches rather than engage in a price war.

We begin with an economy composed of two niches of consumers, n_1 and n_2 . Each niche is composed of identical consumers; consumers in n_1 all have the same w_1 and k_1 . Similarly, all consumers in niche n_2 have the same w_2 and k_2 . All consumers have the same clutter cost $\alpha(\cdot)$. The economy also contains two producers (or firms), f_1 and f_2 , which have perfect information about the consumer population. These producers set prices simultaneously. Their only strategic decision is whether to sell using a pure bundling strategy where consumers have access to all N articles for a flat fee or to charge a fixed price per article purchased. This leaves a producer with a single choice between well-understood strategies. In this section, we do not consider the additional producer decision as to what size bundle to offer. This will be treated in section 2.6. In short, most of the results here carry through to the case where producers choose bundle sizes and prices.

We next define the demand function $D_n(p_i, p_j)$. This is the number of articles bought from producer f_i by a consumer of niche n when prices p_i and p_j are offered. If f_i is offering a per-article schedule, then $D_n(p_i, p_j)$ will be in the range $(0, N)$, and if f_i is offering a bundle, $D_n(p_i, p_j)$ will

be either 0 or N . If firm f_1 sells per-article, it will receive a profit of:

$$\pi_1 = \sum_{n \in \text{niches}} |n| p_1 D_n(p_1, p_2) \quad (3)$$

If firm f_1 bundles, it will receive a profit of:

$$\pi_1 = \sum_{n \in \text{niches}} |n| p^* D_n(p_1, p_2) \quad (4)$$

where $p^* = \frac{p_1}{|B|}$.

Since the producers are offering identical articles, if they offer different prices, a consumer will buy from either f_1 or f_2 (or neither), but not both. Clearly, if both producers offer a per-article schedule or if both producers bundle, there will be a price war, since they will have no way of differentiating themselves from each other besides price. Therefore, we focus on the situation when producer f_1 sells per-article and f_2 bundles in order to determine conditions under which the two can each achieve positive profit.

Each consumer wants to maximize its surplus, which is the value of the articles bought less the price paid. For a bundle of goods, the potential surplus is $S(B, p) = V(B) - p$. For a per-article offering, the consumer can choose how many articles it wants to purchase, so the potential surplus is:

$$\begin{aligned} S(B, p) &= S(\text{nbought}(p), p) = \\ &V(\text{nbought}(p)) - p \times \text{nbought}(p) = \\ &\left(\sum_{i=1}^{\text{nbought}(p)} \mu(i) \right) - \alpha(\text{nbought}(p)) - p \times \text{nbought}(p) \end{aligned}$$

where $\text{nbought}(p)$ is the number of articles a consumer values positively at price p . We find nbought by finding the largest rank article (denoted with i) whose value less clutter cost is greater than p . (Note that nbought has a maximal value of $|B|$, the number of articles for sale.)

$$p \leq w \left(1 - \frac{i}{k|B|} \right) - \alpha(i)$$

$$p \leq w - \frac{wi}{k|B|} - \alpha(i)$$

$$\frac{wi}{k|B|} \leq w - p - \alpha(i)$$

$$i \leq \text{nbought}(p) \leq k|B| - \frac{kp|B|}{w} - \frac{\alpha(\text{nbought}(p))k|B|}{w} \quad (5)$$

So, in order for niche formation to occur, it must be the case that niche n_1 receives a greater surplus from f_1 than f_2 and niche n_2 receives a greater surplus from f_2 than f_1 (or vice versa). Also, each firm must receive more profit from targeting separate niches than from selling to the entire population. Writing out the consumer restrictions, we have the following four inequalities:

$$V_1(\text{nbought}_1(p_1)) - p_1 \times \text{nbought}_1(p_1) > V_1(B_2) - p_2 \quad (\text{i})$$

$$V_1(\text{nbought}_1(p_1)) - p_1 \times \text{nbought}_1(p_1) \geq 0 \quad (\text{ii})$$

$$V_2(\text{nbought}_2(p_2)) - p_2 \times \text{nbought}_2(p_2) < V_2(B_2) - p_2 \quad (\text{iii})$$

$$V_2(B_2) - p_2 \geq 0 \quad (\text{iv})$$

We begin with condition (iv). If producer f_2 (the bundler) has perfect information, it can set its price equal to one

niche's valuation for the entire bundle, making $V_2(B_2) = p_2$ and surplus equal to zero for that niche. The remaining three conditions needed for niche formation (from the consumer point of view) are that the per-article-purchasing niche not want to buy the bundle (condition (i)), that the per-article-purchasing niche's surplus for the per-article offering be nonnegative (condition (ii)), and that the bundle-purchasing niche's surplus for the per-article offering be less than the bundle surplus (condition (iii)). Condition (iii) is easy to define: Since we know that the bundle surplus is zero, we must find conditions under which the per-article surplus will always be negative. Niche n_2 will always receive negative surplus for a per-article price schedule if and only if the per-article price is greater than its net valuation for its most-preferred article. In other words, the price p_1 is greater than w_2 (the value of its most-preferred article) minus $\alpha(1)$, the cost of processing a single article.

To satisfy condition (i) (that niche n_1 not want to buy the bundle), it must be the case that there be some positive number of articles such that n_1 's net value for them less p_1 is greater than niche n_1 's value for the bundle less p_2 . This is certainly the case if $V_2(B)$ is larger than $V_1(B)$, or, simplified, $w_1 k_1 < w_2 k_2$. Since p_2 will be set equal to $V_2(B)$, niche n_1 would receive a negative surplus from purchasing the bundle in this case. It is also the case that n_1 will prefer the per-article schedule if the marginal value acquired by purchasing the additional $|B| - i$ articles (where i is the number of articles niche n_1 would purchase at price p_1) is less than $p_2 - p_1 i$, or the extra value is outweighed by the additional cost. Condition (ii) states that consumers in niche n_1 must not receive negative surplus from purchasing per-article. This is trivially true if no articles are purchased. In this case, however, firm f_1 will receive no profit, and will want to compete for niche n_2 . Consumers in niche n_1 also have a positive valuation on at least one article if clutter cost is less than $w_1 - p$. So, to satisfy condition (ii), clutter cost must be in the range $(0, w_1 - p)$.

So far, we have discussed the conditions necessary for the consumer population to potentially behave as niches exhibiting different purchasing behavior. In the next section, we discuss the producer's decision, and analyze the cases in which it is rational for a producer to actually target one niche and ignore the other, rather than attempting to target the entire population.

2.4 Producer Decision-making

The previous section described the conditions that must exist within the consumer population to allow for the formation of niches. An equally important concern is the appeal of these niches to producers. In order to avoid a price war, it must be the case that a producer will profit at least as much from targeting a single niche as from competing for the entire population.

Consider the problem first from the perspective of f_1 , the per-article firm. If p_1 is the optimal price to charge niche n_1 , and p'_1 is the optimal price to charge the entire population, then

$$p_1 |n_1| D_1(p_1, p_2) \geq \sum_{n \in \text{niches}} p'_1 (|n| D_1(p'_1, p_2)).$$

Similarly, from the perspective of f_2 , the bundling producer, if it is to target niche n_2 at price p_2 , rather than the popu-

lation at large at price p'_2 , then (if $p^* = \frac{p_2}{|B|}$ and $p^{*'} = \frac{p'_2}{|B|}$):

$$p^* |n_2| D_2(p^*, \cdot) \geq \sum_{n \in \text{niches}} p^{*'} (|n| D_2(p^{*'}, \cdot)).$$

In order to say much more about this, we must be slightly more concrete. In particular, we must specify a form for α . We will set $\alpha(n) = n^2$, where n is the number of articles purchased. The greater the degree of α is, the smaller that w , k , and $|B|$ must be in order for niche effects to occur. Figure 2.4 presents an illustration of this. In the first figure, a consumer's raw valuation for different numbers of articles is shown, along with a linear clutter cost and the net valuation. In the second plot, a quadratic clutter cost is shown. As clutter cost increases, consumers desire fewer articles, the advantages of bundling are lessened, and there is more potential for niche formation. If clutter cost becomes too large, then bundling is no longer an effective strategy and per-article pricing becomes more appealing. A quadratic clutter cost is a reasonable assumption if a consumer's opportunity cost for time is concave. In other words, if 10 minutes of time filtering through a bundle of articles to find those of value is more than twice as valuable to a consumer as five minutes spent filtering through articles, then the cost incurred by the consumer for losing those additional five minutes is superlinear. In the limit, one can imagine the difference in value of a single article delivered to a reader, and the value of that article when it is buried in a truckload of other articles.

Consider a situation where $\alpha(i) = i^2$, niche n_1 consists of a population with $w_1 = 95$ and $k = 0.1$ and niche n_2 consists of a population with $w_2 = 25$ and $k = 1$. A per-article producer could target niche n_1 at a price of 94. In order to compete with this, a bundling producer would have to charge a price of -5, meaning that this niche is not feasible for the bundler to target. The bundler could also choose to target niche n_2 , at a price of 25. The per-article producer could also choose to target niche n_2 ; it would have to charge 18.5, selling 2 articles to each consumer and extracting a profit of 37 from each consumer in n_2 . In addition, all the consumers in n_1 would buy at this price. This is a viable strategy for the per-article producer if $|n_1|18.5 + |n_2|37 > |n_1|94$, or n_2 is $\frac{75.5}{37} = 2.04$ times as large as n_1 . If n_2 is smaller than this, f_1 is better off targeting niche n_1 and leaving niche n_2 to f_2 .

In this case, the decision as to whether to target a niche is dependent upon the per-article producer. By reducing clutter cost, we can also construct scenarios in which the bundler has control over this situation. This leads us to another question. So far, we have assumed that producers f_1 and f_2 have decided their pricing strategies ahead of time and will not change them. What is to prevent firm f_2 in the previous scenario from adopting a per-article pricing strategy and starting a new price war?

For firm f_2 to decide to bundle and not engage in a price war, it must believe that it can gain at least as much profit by targeting a separate niche, even if that niche may appear to be smaller. In the following section, we describe a formulation for this game.

2.5 Selection of Pricing Strategies

Each firm has an initial choice of choosing either schedule 1 or schedule 2. Without loss of generality, let us assume that schedule 1 will target either niche n_1 or the entire pop-

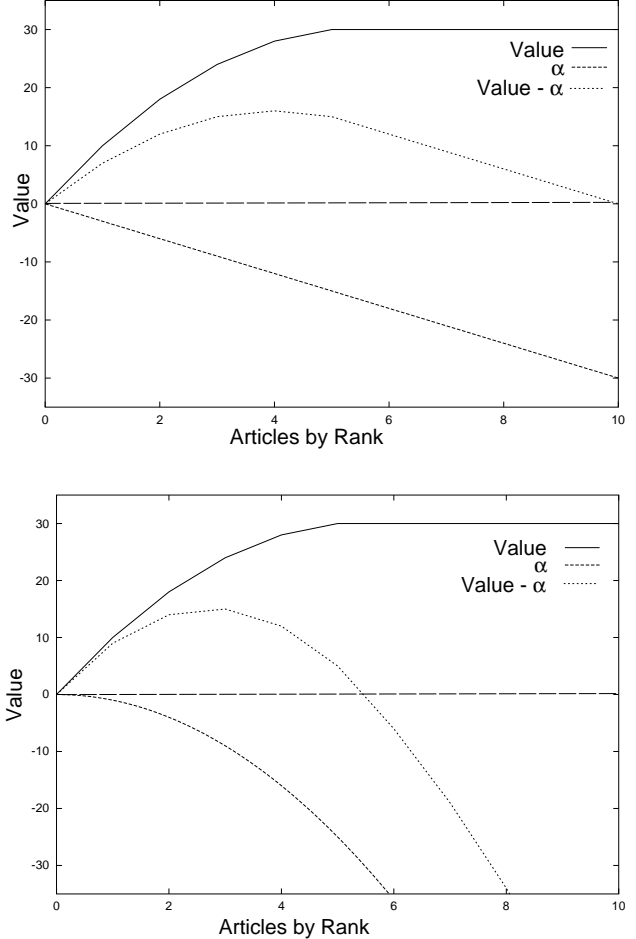


Figure 1: Valuation, clutter cost and net valuation for $w = 10, k = 0.5, N = 10$. The left figure shows $\alpha = -3i$, and the right shows $\alpha = -i^2$

	s1	s2
s1	$\frac{\pi(n_1)}{2}$ $\frac{\pi(n_1)}{2}$	$\pi(n_1)$ $\pi(n_2)$
s2	$\pi(n_2)$ $\pi(n_1)$	$\frac{\pi(n_2)}{2}$ $\frac{\pi(n_2)}{2}$

Figure 2: Game matrix for determining whether to engage in a price war or pursue separate niches

ulation, and schedule n_2 will target niche n_2 . If both firms choose schedule 1, a price war will occur. (In the previous scenario, schedule 1 was per-article pricing and schedule 2 was bundling.) Let us assume that, if each firm chooses a different schedule, it is more profitable to target niche n_1 than the entire population. (Otherwise, the dominant strategy is to engage in a price war for the entire population.) Let us also assume that the game happens in two stages: first, producers simultaneously choose schedules, observe each other’s choices, then simultaneously set prices. The choice we are particularly interested in is the choice of schedule.

We can then construct a game matrix describing the problem. This is shown in figure 2. $\pi(n_1)$ and $\pi(n_2)$ are the maximum profits that can be extracted from niche n_1 and n_2 , respectively. (These can be derived using equations 3, 4, and 5). This is, in fact, a version of the Hawk-Dove game [15], with schedule 1 being the “Hawk” strategy and schedule 2 being the “Dove” strategy. In the Hawk-Dove game, each participant can either take an aggressive strategy (the “Hawk”) or a passive strategy (the “Dove”). If both players are aggressive, they will compete for the larger reward and incur a cost for competing. Similarly, if both players are passive, they will compete for the smaller reward and pay a penalty for competing. If one player is aggressive and one passive, they will each receive a reward (the Hawk’s is larger) without paying a competition cost. If both players compete for the larger niche, then the largest expected payoff a producer can hope to receive is $\frac{\pi(n_1)}{2}$. If both producers set their price to extract the optimal profit, then each has a 50% chance of winning the niche. (Alternately, half the consumers buy from each producer.) An alternate line of reasoning on a Hawk producer’s part would be to charge slightly less than the optimal price and win the whole niche. However, since its opponent is perfectly rational, its opponent would charge slightly less than that. This leads to the price war being carried out within each producer’s reasoning process (as in a Bertrand game) and each producer pricing at marginal cost.

For this analysis, we will give each producer the benefit of the doubt and allow them to receive $\frac{\pi(n_1)}{2}$ for each playing schedule 1. If $\pi(n_1) > \pi(n_2) > \frac{\pi(n_1)}{2} > \frac{\pi(n_2)}{2}$, then each producer would prefer to be in a separate niche, rather than fighting. Of course, both would prefer to be in niche n_1 . We can apply standard game theoretic techniques (see, for example, Fudenberg and Tirole [11]) to determine that each producer should select schedule 1 with probability $\frac{2\pi(n_1) - \pi(n_2)}{\pi(n_1) + \pi(n_2)}$. As the niches become more equal in size, this approaches 0.5. With probability 0.5, each producer is trying to target a separate niche all of the time, but it must coordinate with its opponent. If the producers play an iterated game in which they are able to signal each other or

remember their past, they can then quickly settle on separate niches and play the optimal strategies with probability 1. Our previous work [4] contains an example of this with three niches.

This analysis assumes that each producer selects a schedule and has a single interaction with the consumer population. If, instead, a producer chooses a schedule and then has a series of interactions with the consumer population, the analysis changes.

Let us assume that, once a producer has chosen a schedule, it is locked in to that choice for the next c iterations. During this period, it can change its price, but not the schedule it uses. In that case, a producer which has a niche n all to itself receives a profit of $c\pi(n)$ over c iterations. A producer which engages in a price war will receive a positive profit on every other iteration. (Its opponent wins on the other iteration.) If each producer must lower its price by δ from the previous price, where $\pi(n) - c\delta = 0$, then over c iterations, a producer in a price war will receive a profit of:

$$\sum_{i=0}^{\frac{c}{2}} (\pi(n) - 2i\delta) = \frac{c}{2}\pi(n) - \frac{c^2\delta}{4} + \frac{c\delta}{2} \quad (6)$$

So, a producer should settle for the smaller niche n_2 rather than pursuing a price war over niche n_1 if $\frac{c}{2}\pi(n_1) - \frac{c^2\delta}{4} - \frac{c\delta}{2} < c\pi(n_2)$ or (simplifying), $\pi(n_1) < 2\pi(n_2) + (c + 1)\delta$. In general, following the same techniques as above, we find that a producer should choose schedule 1 with probability

$$q = \frac{2\pi(n_1) - \pi(n_2) + (\frac{c}{2} - 1)\delta}{\pi(n_1) + \pi(n_2) + (\frac{c}{2} - 1)\delta} \quad (7)$$

This additional $(\frac{c}{2} - 1)\delta$ serves to make schedule 1 more appealing. If it is small, then we are left with the same proportion we had previously. As it becomes larger, it pushes q closer and closer to 1. Recall that c is the number of iterations in the price war and δ is the price decrement. Together, they determine the gap between $\pi(n_1)$ and the marginal cost. The larger this is relative to the profit for the smaller niche, the more appealing a price war becomes.

2.6 General Bundling

In this section, we extend the results described above to the case in which producers are able to choose both a bundle size and a price. This strategy is sometimes referred to as generalized subscription [14]. In this paper, we will refer to it as general bundling. The producer offers a bundle of size b for a price p . The consumer is able to select the b articles it wants out of a larger set; we still are not requiring the producer to make this decision.

The results from the previous sections comparing per-article pricing and pure bundling extend in a straightforward manner. This is easy to see if one considers per-article pricing as offering bundles of size 1. Producers are then located at either end of the bundle-size line. By introducing general bundling, we allow producers to locate themselves anywhere on this line. This allows a producer to fit a niche’s demand more exactly and provides it with more ways of differentiating itself from potential competitors.

Note that the consumer demand function does not change at all if we allow producers to offer general bundles. Consumers don’t care whether a bundle’s size is determined exogenously or by a producer; they give it the same value in either case.

2.6.1 Niche Existence

The first question we ask is how the existence of niches changes if producers are allowed to offer general bundles. We begin by extending the demand function to include bundle size. Therefore, $D_n(p_i, b_i, p_j, b_j)$ is the number of articles bought from producer i by a consumer of niche n when i offers a bundle of size b_i for price p_i and producer j offers a bundle of size b_j for price p_j .

We can then define the profit received by a producer from a niche as:

$$\pi_n = |n|p^*D_n(p_i, b_i, p_j, b_j) \quad (8)$$

where $p^* = \frac{p_i}{b_i}$. (In other words, the price per article).

Recall that, for niche formation to occur, consumers must prefer different bundles, and it must also be rational for producers to target separate niches. The conditions for consumers (presented in section 2.3) generalize quite easily. It must be the case that consumers of niche 1 get more surplus from the bundle of producer i than they do from j (or vice versa) and that the surplus from producer i 's bundle be nonnegative. Likewise, these conditions must hold for niche 2 and producer j . However, these conditions are easy to ensure. Consider producer i and niche 1. With perfect information, producer i can offer a bundle of size equal to the number of articles positively valued after clutter cost is subtracted and price equal to the area under this curve. (See figure 1 for a graphical representation of this; the bundle size is the point where $V - \alpha$ intersects the x axis and the price is the area under this curve.) Since all potential profit is extracted, surplus will be zero. This means that a consumer of niche 1 will want to purchase a bundle from producer i rather than the bundle of producer j (targeted to niche 2) if and only if the bundle of producer j provides a negative surplus. This can happen if niche 2's w or k is sufficiently larger than niche 1's. (If both are larger, then producers will engage in a price war for the small w, k niche and the winner will get the larger niche for free.) By symmetry, the same arguments hold for niche 2 and producer j .

To summarize, if producers can set a bundle size and price, each producer can target a niche exactly and extract all possible surplus. A price war can be avoided if it is the case that the optimal bundle for one niche is not appealing to the other niche. It is still possible for a producer to choose an intermediate bundle that appeals to both niches; as in section 2.4, if a price war for the entire population is to be avoided, it must be the case that a producer can make more profit by specifically targeting a niche than by targeting the entire population.

2.6.2 Producer decision-making

The results derived previously in section 2.5 still hold true in this case. A producer that can offer a general bundle has three choices: target niche 1, target niche 2, or select an intermediate bundle and engage in a price war for the entire population. If we assume, as we did previously, that targeting the general population is not the best strategy (that is, there is some reason for producers to differentiate themselves), then we can return to the game matrix shown in figure 2. All that has changed in this is the values of $\pi(n_1)$ and $\pi(n_2)$. If each producer can set a bundle size and price with perfect information, they can extract all profit from a particular niche. This is an improvement from the per-article pricing, where a producer was best off setting a

price of $\frac{w}{2}$, selling $\frac{k_i|B|}{2}$ articles for a profit of $\frac{wk_i|B|}{4}$ and losing half the available profit. In short, this makes smaller niches more valuable. Recall that, the more valuable the smaller niche is, the less appealing a price war is, since the profit from not engaging in a price war increases relative to the profit from engaging in a price war. As producers increase the number of ways in which they can differentiate themselves from each other (in this case, by offering bundles of different sizes), it becomes easier for each producer to extract profit from a population containing niches and therefore avoid a price war.

3. NICHE DISCOVERY

In the previous sections, we have shown when it is reasonable for producers with perfect information to seek out separate niches rather than engaging in a price war. We also focused primarily on equilibrium solutions, showing optimal strategies to be played when all participants are perfectly rational. Of course, in real life, producers do not typically have perfect information, either about their opponent or about the consumer population. The previous analysis does not give us any indication of whether price wars will emerge when producers are learning, how easy or difficult it is to discover the existence of niches, or when the loss in profits incurred by discovering a niche is outweighed by the extra profit from avoiding a price war. In order to study this nonequilibrium behavior, we have constructed a set of experiments which show when a learning producer should focus on a separate niche and when it should compete with an opponent for the entire population.

3.1 Learning Producers

We use a rather naive model of a learning producer. A producer has no knowledge about the consumer preferences, and thus cannot take advantage of the underlying structure in the distribution of these preferences. Instead, it simply receives a profit signal from the environment as a result of the prices set by itself and its opponent. Note that a producer is able to observe its opponent's prices but not its profit. Since a producer is interested in maximizing its aggregate profit over time, it must trade off the cost of exploring further and potentially finding a better set of prices against the benefit of exploiting currently known information. In these experiments, we maintain a static exploration policy and demonstrate how its effectiveness changes as the consumer population varies.

Each producer uses the same learning algorithm. A producer maintains a model of the world in the form of a radial basis function network [6] which maps prices set by each producer into profit. The process through which a producer determines the next price to offer is depicted in figure 3. To determine a price to offer, a producer first constructs an estimate of its opponent's strategy (using the price the opponent offered on the last iteration) and determines a best response to this strategy. It does this by using Amoeba [16], a nonlinear hill-climbing algorithm, to search the landscape formed by the model it has learned. Once it has found a best response, it decides whether it should explore or exploit this knowledge. This decision is a biased random variable weighted by the agent's estimate of the fraction of the profit landscape it has observed so far. (Prices are discretized to integers for this calculation.) If the agent chooses to explore, it perturbs the value of the best response. For each

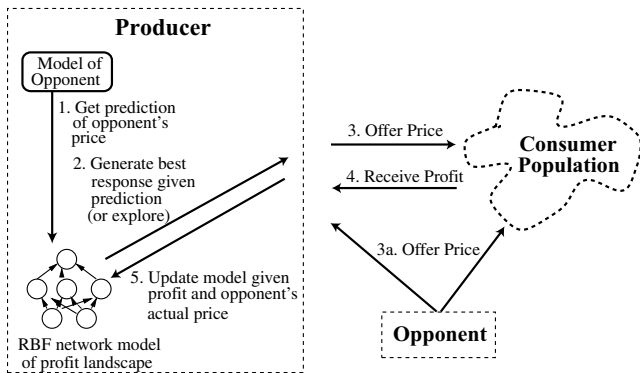


Figure 3: Process by which a learning producer decides upon a price to offer.

dimension in the schedule, it probes the model to find the direction which has been less explored locally and alters the best response in that direction by a fixed amount. Otherwise, the agent sets a price which maximizes its expected surplus in its model. We make no particular claim that this learning algorithm is optimal for this task, nor that this particular exploration policy is the best possible. That is not the question we are trying to answer. Instead, we are interested in how the performance of a learning producer changes as the characteristics of the problem change. Once an agent designer understands how the problem characteristics influence the difficulty of learning, he or she can then optimize a learning algorithm for this task.

3.2 Experimental Setup

We conducted two sets of experiments involving producers that needed to learn about consumer preferences. In the first experiment, the relative size of the niches was varied, and in the second, the number of iterations each producer had to learn before its opponent changed its price was varied. The purpose of this was to determine how easy it actually is for a producer to discover a niche. In the previous section, we described conditions under which a producer with perfect information should target a niche. However, in a real-life environment, a producer not only might not know the preferences of a niche, it may not even be aware of the existence of some niches at all.

In our experiments, we begin with two niches and an incumbent producer. One niche consists of consumers with a large w and small k . A producer targeting this niche is better off selling per-article. The second niche consists of consumers with a small w and large k . This niche can be effectively targeted by a producer offering a pure bundle. We assume the following scenario: an incumbent producer is selling information goods as a monopolist using a per-article schedule to a consumer population consisting of the two niches. This producer has enough time to learn an approximation of the optimal price. A second producer then enters and offers information goods in a bundle, with the intent of capturing one niche. Both producers must now adjust their prices, either to target separate niches or to compete for the entire population. We chose an entrant-incumbent scenario because we felt it was more realistic than one in which two producers entered simultaneously, and also

because it avoids the uninteresting result in which two producers with no knowledge completely inhibit each other's learning.

3.3 Niche Size

In section 2.5 we described a producer's decision as to whether to pursue a price war or to target a smaller niche in terms of the difference in size between the niches. This analysis assumed perfect information on the part of the producers. In this section we examine how a learning producer's cumulative profit changes as the relative size of each niche changes in imperfect-information situations.

We used an environment consisting of two niches of consumers. The first niche n_1 had a w of 95 and a k_i of 0.1. Consumers in the second niche had a w of 25 and a k_i of 1. All consumers had a quadratic clutter cost of $\alpha(i) = i^2$. The relative sizes of the niches were varied by changing the number of consumers in each niche. Recall from the example in section 2.4 that if producers have perfect information, producers should pursue separate niches if niche n_2 is more than 2.04 times as large as niche n_1 .

We began with a scenario in which each niche contained the same number of consumers. We then moved to a scenario in which niche n_2 , with low w and high k , had twice as many consumers, then three times as many, then four. As the size of niche n_2 increases, targeting separate niches becomes less attractive for the per-article producer. We assumed that the entering producer had already made the decision to bundle and not engage in a price war. The purpose of this experiment is to determine conditions under which the bundling producer can actually enter and achieve a positive profit, and how difficult it is for the per-article producer to select whether to target a niche or the population as a whole. Figure 4 contains graphs showing the four scenarios. The curves indicate the aggregate profit per iteration, error bars indicate standard deviation, and the horizontal lines show optimal profits for producers with perfect information. Data is averaged over 50 trials.

These figures present several interesting points. When the per-article producer sells as a monopolist, it reaps very high profits. These fall off quickly upon the entry of the bundler as the per-article producer must adjust its prices. In the first two experiments, the best strategy for the per-article producer is to target niche n_1 , and in the second two, the best strategy is to target the population as a whole, charging a price of \$6 and leaving the bundling entrant with no profit. Optimal profit for each niche, and for the per-article producer when it is best off targeting the entire population, is indicated with a horizontal line in each graph.

Even though the per-article producer would have been substantially better off in the last two experiments by targeting the population as a whole, the per-article producer often targeted niche n_1 , leaving the bundling producer with positive profit. Briefly, this is because the gap in price between the optimal price for niche n_1 and the optimal price for the entire population is so large. This is discussed in more detail in section 3.5.

Note that the standard deviation (σ) for the per-article producer is quite large, whereas it is smaller for the bundler. While some of this is due to the bundler's optimal profit being smaller, some of it is due to an indecision on the part of the per-article producer. Since it learns the profits associated with a price contingent upon the offerings of the

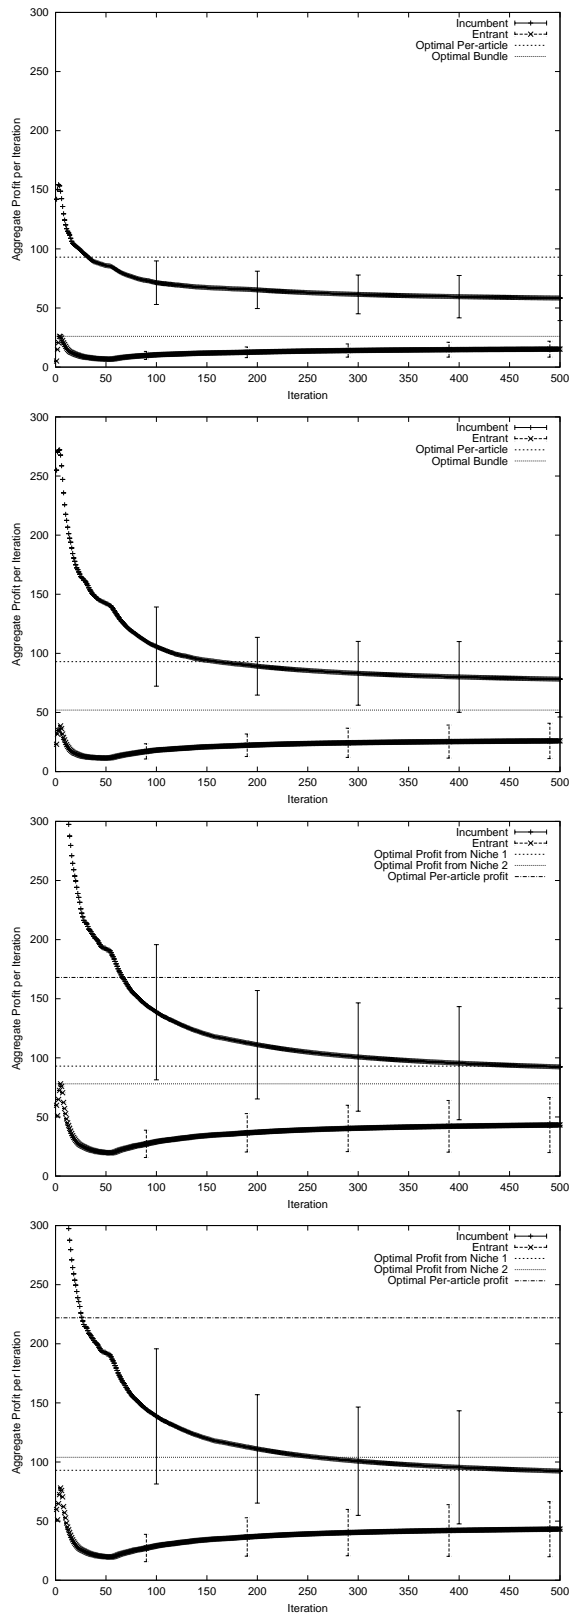


Figure 4: Cumulative profit per iteration for (top to bottom) equal size niches, Low w /high k niche twice as large, three times as large, and four times as large. The upper curve represents the incumbent, and the lower the entrant.

bundler, it was sometimes difficult for the per-article producer to systematically explore the pricing space. (More on this below.) Many times, the per-article producer settled into a sub-optimal state simply because, as it was exploring, the bundler set prices that led the per-article producer down a dead end. For example, if the bundler charged a very high price early on, the per-article producer never got any information about how the population would respond to its prices when there was serious competition with the bundler. The effect of this can also be seen in the bundler's standard deviation, which increases as the relative size of niche n_2 increases. When n_2 is small, the bundler's aggregate profit is very predictable, indicating that it almost always set the same prices. As the size of n_2 increased, the standard deviation of the bundler's aggregate profit also grew, indicating that it settled into a wider region of the bundling space. Since the bundler should receive no profit if the per-article producer acts optimally, this large standard deviation indicates that the bundler is sometimes successful at inhibiting the per-article producer's learning.

Also, notice that the error bars sometime extend above the optimal profit for the per-article producer. The reason for this is that this optimal profit is determined based upon a rational, perfect information best response by the bundler. If the bundler makes a particularly bad offer and extracts less profit than is possible, that leaves more profit for the per-article producer.

3.4 Varying Rate of Change

One phenomenon that was observed in some of the previous experiments was a sort of "chasing." Producers alternated the setting of prices; each would alternately set a best response to the others' last price. Since both producers were doing a great deal of exploration in the early stages, it was hard for either producer to push the system toward an optimum. Instead, it would wind up circling through a series of low-profit prices. This is referred to as the moving target function problem [19], since a producer's model of the opponent (and consequently, its payoff) is changing as its opponent learns. We hypothesized that if producers learned more serially, holding still for a time while their opponent learned, this problem could be avoided.

We used the population in which there were twice as many consumers in n_2 as in n_1 . We began with a baseline experiment in which producers alternated pricing, then allowed each producer to change prices for three iterations while its opponent held still, then five iterations. Results from these experiments (averaged over 50 trials), along with error bars indicating standard deviation and horizontal lines indicating optimal profit, are shown in figure 5.

The rate of convergence did not seem to change drastically as learning was serialized; the curve representing average aggregate profit per iteration seen in each of the three graphs in figure 5 is basically the same. What does differ between experiments is the standard deviation of these curves. In particular, the per-article producer's standard deviation improved as serialization was introduced. For example, after 300 iterations, σ for the per-article producer is 27 when the producers alternate every iteration and 21 when they alternate every 5 iterations. Similarly, after 400 iterations, the per-article producer's σ is 30 when producers alternate every iteration and 21 when they alternate every five. This would seem to indicate that serializing the learning provides

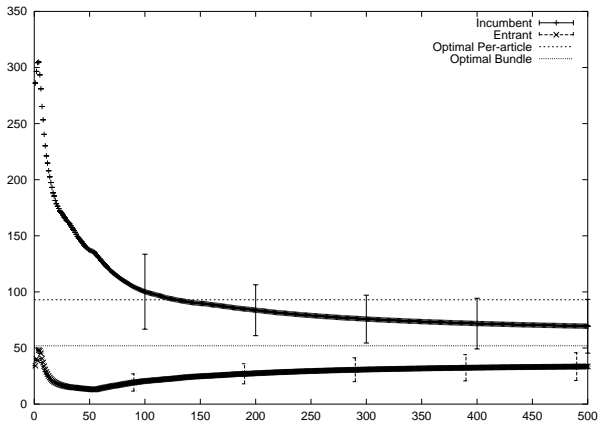




Figure 5: Cumulative profit per iteration when producers learn for one (upper left), three (upper right) and five (lower left) iterations, then remain still for an equal period. The upper curve represents the incumbent, and the lower the entrant.

the producers with an easier learning problem and so they converge to the optimal solution more frequently. This is a double-edged sword for the per-article producer, however. Serializing learning also reduces the chance that the bundler will make a large mistake and leave the per-article producer with an extra-large share of the market. Consequently, the average aggregate profit is slightly smaller (80 to 72 at 400 iterations) for the per-article producer when serial learning is introduced.

Both of these sets of experiments led us to wonder what it was about this problem that caused the per-article producer to get misled so often. For example, when niche n_2 has four times as many consumers as n_1 , it seems clear that the per-article producer should shut out the bundler and claim the entire market. Yet the bundler still managed to average a positive profit. Also, the “chasing” behavior mentioned above was still not completely clear. To answer these questions, we constructed profit landscapes for each producer.

3.5 Examining the Profit Landscape

In order to better understand what made this problem hard for the per-article producer, we constructed profit landscapes showing the profit for each producer as a function of prices for each producer. These landscapes are for a two niche consumer population, with the high w /low k niche composed of one consumer with $w = 95$ and $k = 0.1$ and the low w /high k niche composed of four consumers with $w = 21$ and $k = 10$. These landscapes are shown in figure 6.

There are several interesting features of these landscapes. First, the per-article producer’s landscape has two ridges, one where it is targeting the high w /low k niche and one where it is targeting both niches. As we discussed previously, with a large clutter cost, the per-article producer typically has control over whether or not to engage in a price war. In this case, it is more profitable for the per-article producer to avoid the price war and focus on the high w /low k niche.

The shape of the price war peak in figure 6 is a function of the low w /high k population. As their w increases, the slope decreases. In this case, consumers value articles more highly and are more tempted to switch over to the per-article producer when its price is low. The amplitude of this peak is controlled by this population’s k . Since a higher k leads to a greater area under their demand curve, there is more profit to extract.

These two peaks, along with the long sloping region, make the problem hard for the per-article producer. Initially, it starts out near one peak. If its explorations take it over to the other peak, it may realize that it is better off switching strategies. However, depending upon the makeup of the consumer population, it may need the assistance of its opponent to explore this peak. This helps to explain why, in cases where the per-article producer could win the entire population, it focused on one niche. It would have to descend through this chasm to find the other peak.

The bundler’s landscape is simpler; if the per-article producer is not trying to appeal to both markets (by pricing below 20 in this case), the bundler is insensitive to the per-article producer’s prices. This makes sense; both producers are essentially acting as monopolists at this point. The difficulty for the bundler comes in its large flat area, in which it receives no information about either the structure or the gradient of the landscape. The bundler can potentially spend a lot of time exploring that area and, consequently, misleading

the per-article producer as to its optimal price.

4. CONCLUSIONS AND FUTURE WORK

This paper has described some scenarios in which it is possible for duopolist information goods producers to avoid the sorts of price wars that have often seemed inevitable in other models studied in the literature. We have introduced the idea of niche populations and shown both conditions within a consumer population consisting of two niches that are necessary for the formation of niches and conditions which must hold in order for producers to decide to take on a strategy of targeting a niche population rather than competing in a price war. We have also presented experimental results which illustrate the practical difficulties learning producers can have in determining whether or not to engage in a price war and in finding the correct niche.

This work serves as an initial attack at an area which, until now, has been largely untapped. Our analytical results focus only on the case of two producers targeting a population consisting of two niches with identical consumers; our continuing research will extend these results to larger number of niches and producers and a greater degree of heterogeneity among the consumer preferences.

We are also extremely interested in endowing our learning agents with a degree of “economic intelligence.” Currently, they learn in a rather naive fashion, without exploiting economic knowledge, known properties of the consumer population, or regularities in their opponent’s behavior. A learning agent that also made some reasonable assumptions about (for example) the structure of consumer preferences or the sorts of strategies its opponent was likely to play would likely be able to learn more efficiently.

We would also like to extend this analysis to more complex pricing schedules. As we have discussed, both previously [5] and in this paper, there is a large space of pricing schedules beyond linear pricing and pure bundling; it may be that some of these are more effective at targeting niche populations. Additionally, we plan to consider cases in which producers can choose not only a schedule to offer, but the actual goods that they plan to bundle. This would give producers a greater ability to differentiate themselves from each other and provide a test for our conjecture that a producer’s ability to avoid a price war is dependent upon its ability to differentiate itself from its opponents in dimensions other than price.

5. REFERENCES

- [1] Robert Axelrod. *The Evolution of Cooperation*. Basic Books, New York, NY, 1984.
- [2] Yannis Bakos and Erik Brynjolfsson. Bundling and competition on the Internet: Aggregation strategies for information goods. *Marketing Science*, January 2000.
- [3] Christopher H. Brooks and Edmund H. Durfee. Congregation formation in information economies. In *Proceedings of the AAAI-99 Workshop on AI in Electronic Commerce*, 1999.
- [4] Christopher H. Brooks, Edmund H. Durfee, and Aaron Armstrong. An introduction to congregating in multiagent systems. In *Proceedings of the Fourth International Conference on Multiagent Systems*, pages 79–86, 2000.

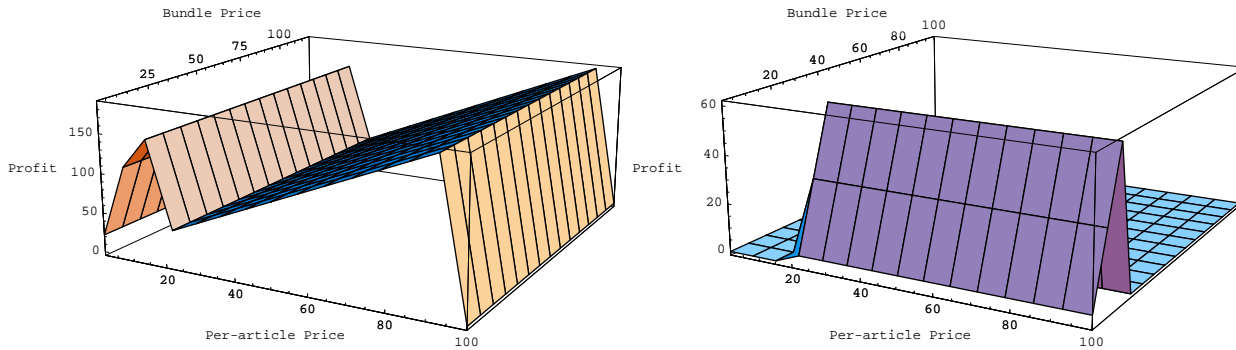


Figure 6: Per-article (left) profit and bundle (right) profit as a function of per-article price (x axis) and bundle price (y axis).

- [5] Christopher H. Brooks, Scott Fay, Rajarshi Das, Jeffrey K. MacKie-Mason, Jeffrey O. Kephart, and Edmund H. Durfee. Automated strategy searches in an electronic goods market: Learning and complex price schedules. In *Proceedings of ACM EC-99*, pages 31–40, 1999.
- [6] D. S. Broomhead and David Lowe. Multivariable functional interpolation and adaptive networks. *Complex Systems*, 2:321–355, 1988.
- [7] Anthony Chavez and Pattie Maes. Kasbah: An agent marketplace for buying and selling goods. In *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM'96)*, April 1996.
- [8] John Chung-I Chuang and Marvin A. Sirbu. Network delivery of information goods: Optimal pricing of articles and subscriptions. In *Proceedings of the Conference on Economics of Digital Information and Intellectual Property*, Cambridge, MA, January 1997.
- [9] Edmund H. Durfee, Tracy Mullen, Sunju Park, Peter Weinstein, and Jose Vidal. The dynamics of the UMDL service market society. In Matthias Klusch and Gerhard Weiß, editors, *Cooperative Agents II*, Lecture Notes in Artificial Intelligence, pages 55–78. Springer, 1998.
- [10] Scott A. Fay and Jeff MacKie-Mason. Competition between firms that bundle information goods. In *Proceedings of the 27th Annual Telecom Policy Research Conference*, Alexandria, VA, September 1999.
- [11] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, Cambridge, MA, 1991.
- [12] Jeffrey O. Kephart, James E. Hanson, David W. Levine, Benjamin N. Grosz, Jakka Sairamesh, Richard B. Segal, and Steve R. White. Dynamics of an information-filtering economy. In *Proceedings of the Second International Workshop on Cooperative Information Agents (CIA)*, Paris, July 1998.
- [13] Jeffrey O. Kephart, James E. Hanson, and Jakka Sairamesh. Price and niche wars in a free-market economy of software agents. *Artificial Life*, 4:1–23, 1998.
- [14] Jeffrey K. MacKie-Mason, Juan F. Riveros, and Robert S. Gazalle. Pricing and bundling electronic information goods: Experimental evidence. In B. Compaine and I. Vogelsang, editors, *Internet and Telecommunications Policy Research*. MIT Press, 2000.
- [15] John Maynard Smith. *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, UK, 1982.
- [16] William H. Press et al. *Numerical Recipes*. Cambridge University Press, 1992.
- [17] Maksim Tsvetovaty, Maria Gini, Bamshad Mobasher, and Zbigniew Wiecek. MAGMA: An agent-based virtual market for electronic commerce. *Journal of Applied Artificial Intelligence*, 6, 1997.
- [18] Hal R. Varian. Pricing information goods. In *Proceedings of the Research Libraries Group Symposium on Scholarship of the New Information Environment*. Harvard Law School, 1995. Available from <http://www.sims.berkeley.edu/hal/people/hal>.
- [19] Jose M. Vidal and Edmund H. Durfee. The moving target function problem in multi-agent learning. In *Proceedings of the Third International Conference on Multi-Agent Systems*, Paris, France, 1998.