

Improving Learning Performance by Applying Economic Knowledge

Christopher H. Brooks¹, Robert S. Gazzale², Jeffrey K. MacKie Mason², and Edmund H. Durfee^{3*}

¹ Computer Science Department, University of San Francisco, San Francisco, CA 94117-1080
cbrooks@usfca.edu

² Department of Economics, University of Michigan, Ann Arbor, MI 48109
{rgazzale, jmm}@umich.edu

³ EECS Department, University of Michigan, Ann Arbor, MI 48109
durfee@umich.edu

Abstract. Digital information economies require information goods producers to learn how to position themselves within a potentially vast product space. Further, the topography of this space is often nonstationary, due to the interactive dynamics of multiple producers changing their positions as they try to learn the distribution of consumer preferences and other features of the problem's economic structure. This presents a producer or its agent with a difficult learning problem: how to locate profitable niches in a very large space.

In this paper, we present a model of an information goods duopoly and show that, under complete information, producers would prefer not to compete, instead acting as local monopolists and targeting separate niches in the consumer population. However, when producers have no information about the problem they are solving, it can be quite difficult for them to converge on this solution. We show how a modest amount of economic knowledge about the problem can make it much easier, either by reducing the search space, starting in a useful area of the space, or introducing a gradient. These experiments support the hypothesis that a producer using some knowledge of a problem's (economic) structure can outperform a producer that is performing a naive, knowledge-free form of learning.

1 Introduction

Recent advances in networked information technology have led to the emergence of a digital economy. This is an economy in which digital goods, in particular information goods, are traded online. In these environments, a producer of information goods is faced with a potentially daunting problem: how to position itself within a large and dynamic product space. In addition, many of the participants in this information economy are computational agents, which leads us to examine automated techniques for solving this problem.

The large product space is in many ways a result of the characteristics of digital information goods. They can be easily unbundled and re-bundled to create different products. In addition to a pre-defined "newspaper" for example, a producer can sell the

* This work was supported in part by NSF grants IIS-9872057 and IIS-0112669.

whole news collection, categories of articles, individual articles, graphics, data tables, etc. Digital goods typically have low marginal cost: once the first copy is produced, subsequent copies can be reproduced essentially for free. This makes *bundling* of goods a particularly attractive strategy [7]. The effect of unbundling and re-bundling is to exponentially proliferate possible product configurations.

In addition, the economy is likely to be highly dynamic. First, the enormous size of the product configuration space implies that producers will need to *search* for good configurations. If there are competing producers also searching, then a second dynamic is introduced, since the path followed by competitors will affect profits and learning for a producer. Third, the composition of the consumer population or the preferences of individual consumers may be nonstationary. These factors imply that firms must engage in active learning in order to determine what to offer and what price to charge. It is thus important to understand the system's dynamics, in addition to its equilibria.

In this paper, we consider the problem of competing information producers positioning their offerings when faced with a heterogeneous consumer population. We focus on the role that economic knowledge can play in reducing the producers' search burden. When producers must simultaneously learn what to offer and how to price it, the naive learning problem becomes very difficult. However, if producers use even a limited amount of economic knowledge, they can sharply reduce the difficulty of each learning problem, locating a profitable niche and extracting a significant fraction of the equilibrium profit.

"Economic knowledge" is a broad term, and such knowledge can take on a variety of forms. We take the position that economic knowledge is information that allows a producer to more effectively or efficiently learn what good to sell or how to price a good. Much of our research (e.g. [3]) has examined producers at either extreme of the spectrum of knowledge. On one end, economic knowledge can mean complete information about consumer preferences and the strategies of all other producers. This is the classical analytic assumption taken in section 3. At the other extreme, a producer has no knowledge about the problem whatsoever. Even in this case, some basic economic knowledge (for example, that prices must be nonnegative) will likely be built in. In this paper, we explore the space between these two extremes, beginning with the relatively well-understood full-information solution and then progressively weakening the information these producers have. One of our goals in this paper is to study the way in which different sorts of economic knowledge alter the solution space, and the resulting problem difficulty.

Other researchers have considered the problem of producers attempting to locate suitable niches in a high-dimensional product space. Chang and Harrington [4] study conditions under which a chain ought to allow individual stores find profitable niches in 'idea space' and when it ought to centrally supply ideas. They find that when consumers are sufficiently heterogeneous, allowing each store to learn simultaneously provides higher profits than a centralized search.

Several researchers, including ourselves, have studied the problem of "zero-knowledge" agents, meaning agents that learn without taking advantage of the economic aspects of the learning problem. Our previous work examined zero-knowledge monopolists trading off the precision of a price schedule against the difficulty of learning that

schedule [3] and zero-knowledge duopolists attempting to discover whether to compete directly or target separate niches [2]. In both cases, the learning agents did not attempt to exploit any economic knowledge about the structure of the problem.

Kephart, et al [9], and Dasgupta and Das [6] have examined economies in which zero-knowledge agents adaptively adjust the prices they offer in an attempt to maximize aggregate profit. These sorts of systems often exhibit *price wars*, in which producers repeatedly undercut each other, driving prices down to marginal cost and profits to zero; a focus of these papers is on developing learning algorithms that can avoid price wars. We argue that a relatively small amount of economic knowledge can also serve this purpose, and that the paucity of information available in the zero-knowledge setting makes the learning problem overly difficult for each producer.

In this paper we analytically derive optimal producer behavior in a full-information static environment, then use agent simulations to study behavior in an incomplete-information dynamic environment with learning. In Section 2, we present the details of our model of an information market. In Section 3, we derive the static equilibrium conditions and profits, both for a monopolist and for multiple strategic firms. Next, in Section 4, we examine the problem of zero-knowledge duopolists learning what bundle to offer and show that naive learning performs quite poorly. Then, in Section 5, we show that a learning producer with even a limited amount of economic knowledge is able to extract a significant fraction of the available profit. Finally, we discuss our results and future directions in Section 6.

2 Modeling

In this section, we present our model of consumer preferences, describe the structure of the consumer population, and then characterize the resulting behavior of both consumers and producers.

2.1 Consumer Preferences

We assume that there is set of Γ disjoint information categories. Each consumer is willing to pay w for each article in her favorite category, γ^* . All consumers agree that the relationship between categories can be expressed by ordering them from $\gamma_1, \dots, \gamma_\Gamma$.⁴ We label categories by their index. Each consumer values a fraction $k < 1$ of the categories.⁵ We assume that values for articles in particular categories fall off linearly from γ^* . The value that consumer j places on category γ_i is thus:

$$\mu_j(\gamma) = \begin{cases} w_j(1 - \frac{2(|\gamma_i - \gamma_j^*|)}{k_j\Gamma}) & \text{if } |\gamma - \gamma_j^*| \leq \frac{k_j\Gamma}{2} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

Figure 1 illustrates the value that consumers place on articles in categories.

⁴ That all consumers agree on the relationship between categories is an admittedly large assumption.

⁵ Our formulation of consumer demand is similar to that introduced by Chuang and Sirbu [5].

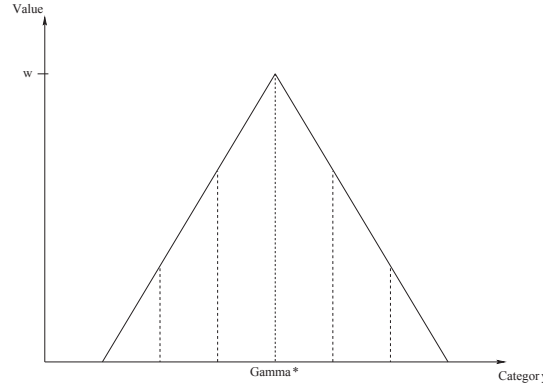


Fig. 1. Consumer valuations of articles in different categories.

A bundle is a collection of articles partitioned into categories. Bundles are useful for producers in this scenario, in that they allow a producer to extract greater profit from a heterogeneous consumer population. We define the size of a bundle as $|B| = \sum_{i=1}^{\Gamma} n_i$, where n_i is the number of articles in category j . While Equation 1 expresses the value of a single article, we wish to capture the notion that large bundles can place a burden upon the consumer due to the scarcity value of the consumer’s attention. We refer to this as clutter cost, $\alpha(|B|)$. We assume that this cost function is exponential (and thus convex): $\alpha(|B|) = e^{\lambda|B|}$. In practice producers may endogenously influence this cost function by providing search and indexing features, thereby changing λ . In this paper, we assume that λ is exogenous and common across all consumers.

Thus, a consumer j ’s valuation for a given bundle is

$$V_j(B) = \sum_{\gamma=1}^{\Gamma} n_{\gamma} \mu_j(\gamma) - e^{\lambda|B|}. \tag{2}$$

2.2 Consumer Population Structure

We assume that w and λ are common across consumers, so a consumer is fully specified by her values of γ^* and k . We also assume there is a finite number of k values, denoted by \bar{k} . We call each unique pair a “niche”, and label a pair by $c_l \equiv \{\gamma_l^*, k_l\}$, with $l = 1, \dots, \Gamma \times \bar{k}$. We assume that there is the same number of consumers in each niche.

2.3 Consumer Behavior

Consumers perfectly evaluate bundles before purchase. Each period, a consumer purchases a bundle which gives her the highest non-negative surplus, if at least one such bundle exists. All consumers in a niche purchase the same bundle except when multiple bundles offer the highest non-negative surplus, in which case each chooses randomly.

2.4 Producer Behavior

The producer's goal is to maximize cumulative (discounted) profits. In each period, each of M producers offers a single bundle, which consists of a number of articles in each category and a price for the bundle. We assume that consumers are anonymous; producers are unable to identify which consumers have purchased their bundles. Instead, each producer only receives its aggregate profit.⁶

We represent producer m 's offer as a vector of length $\Gamma + 1$, $\{n_{m1}, n_{m2}, \dots, n_{m\Gamma}, p_m\} = \{\vec{n}_m, p_m\}$, indicating the number of articles in each category and the bundle price. The producer chooses $n_i \in (0, \bar{n})$, where \bar{n} is the maximum number of articles a producer can offer in a single category,⁷ and chooses $p_i \in [0, \bar{p}]$.⁸

3 Full-Information Equilibria

In this section we characterize the solution of the firm's product configuration and pricing problem in a static, full-information environment. We assume that producers know the values of Γ , w , and λ , and know the distributions of k and γ^* . We first solve the problem for a monopolist, and then characterize the pure-strategy⁹ Nash equilibria when there are multiple complete-information firms.

3.1 One Firm Optimization

In order to provide insights into the effects of consumer demand on the producer's incentives, we first look at the case of a monopolist.¹⁰ The results derived for $M = 1$ are also useful for the analysis when $M \geq 2$.

Lemma 1 *If there is a single producing firm, then p is optimal for \vec{n} only if at least one niche receives zero surplus.*

⁶ This assumption makes offering multiple bundles simultaneously less appealing when a producer must learn, as it becomes more difficult to accurately assess the value of a bundle.

⁷ We set this to be a non-binding constraint for any profit-maximizing firm. For example, we can set \bar{n} to be the positive n such that $w * n - e^{\lambda * n} = 0$. For analytic convenience we assume that the number of articles offered in a category, n_i , has a continuous support. Imposing an integer constraint would greatly complicate the analysis and statement of results without changing their qualitative nature.

⁸ We set \bar{p} to be the maximal willingness to pay possible. As $k < 1$, for any bundle size, a consumer strictly prefers a bundle consisting only of articles in her most-preferred category to any other distribution of $|B|$. The willingness to pay of this consumer is $V = |B|w - e^{\lambda|B|}$. Maximizing this value with respect to $|B|$, and substituting back into the value function results in a willingness to pay equal to $\bar{p} = \frac{w(\log(\frac{w}{\lambda})-1)}{\lambda}$.

⁹ We have not analyzed the mixed-strategy solutions for this problem; this remains a topic for future work.

¹⁰ We assume, for analytical convenience, that optimal bundles are not constrained by the number of categories. That is, if consumers in a niche with $\gamma^* = \Gamma$ purchase, a producer would not change her offering to attract consumers with $\gamma^* = \Gamma + 1$, if they existed.

Proof: Assume not. Let c_l receive the smallest positive surplus from $\{\vec{n}, p\}$. By continuity of p , there exists a \hat{p} such that c_l receives zero surplus. With a higher price and equal demand, profits must increase, and so $\{\vec{n}, p\}$ cannot be optimal. ■

Lemma 2 *Assuming $\bar{k} \geq 2$, if $\{\vec{n}, p\}$ is the optimal offering for a monopolist, then at least 2 niches receive zero surplus.*

Proof: By Lemma 1, we know that at least one niche must receive zero surplus. We show that at least one other niche must receive zero surplus.

Assume that $n_j < |B|$ and that the only zero-surplus niche has $\gamma^* = \gamma_j$. By continuity of n , we can reallocate ϵ articles from γ_i to a category closer or equal to γ_j , for ϵ small. As this increases the surplus of $\gamma^* = \gamma_j$, p can be increased without changing total demand, and therefore $\{\vec{n}, p\}$ cannot not be optimal.

If $n_j = |B|$, all niches with $\gamma^* = \gamma_j$ receive the same surplus, regardless of k , and therefore receive zero surplus. ■

We can now characterize more completely the optimal offering by a monopolist. We assume that there are at least 2 k values, and from Lemma 2 we know that at least two niches receive zero net surplus. We call the two outermost niches, c_a and c_b , the *target niches*. We assume that the most-favored category of c_a is to the left of or the same as that of c_b , or $\gamma_a^* \leq \gamma_b^*$. We first show that given two target niches, it is optimal for the firm to restrict itself to choosing amongst a subset of bundle distributions, i.e. percentages of total articles in each category. Given that the monopolist chooses from this subset, the revenue that it receives is a function only of bundle size. The problem of choosing an optimal bundle and price, *conditional on a pair of target niches*, thus reduces to choosing the appropriate bundle size $|B|$.

At least one combination of target niches, when optimally targeted, must result in the highest profit. We call this combination of niches, c_a^* and c_b^* , the *optimal target niches*.

We start by characterizing optimal bundle distributions given a pair of target niches.

Lemma 3 *If $\{\vec{n}, p\}$ is optimal for a monopolist conditional on $V_{c_a} = V_{c_b} = 0$, then $n_i = 0 \forall \gamma_i < \gamma_a$ and $\gamma_i > \gamma_b$.*

Proof: Assume not, and that $n_i = 0$ for all $\gamma_i < \gamma_a$. Let γ_k be rightmost category with $n_k > 0$. By continuity of n_i , reallocating ϵ from γ_k to γ_{k-1} increases the surplus of c_a and c_b to \hat{p} . Only those with $\gamma^* \geq \gamma_k$ are worse off. All of these niches received positive surplus if they purchased the original bundle. Therefore, $p = \hat{p}$ strictly increases profits as all niches with positive surplus at $\{\vec{n}, p\}$ still have non-negative surplus. Therefore, $\{\vec{n}, p\}$ could not have been optimal. ■

Thus by Lemma 3, conditional on particular niches receiving zero surplus, only bundles where no articles are to the outside of either niche can be optimal. We now show that the set of all “interior” bundles which deliver zero surplus to the target niches are characterized by the mean category of the bundle, $\hat{\gamma} \equiv \sum_{i=1}^I \gamma_i \frac{n_i}{|B|}$, and show for which niches this characterization, together with the bundle size, is sufficient to determine gross surplus.

Lemma 4 *If $n_i = 0$ for all $i < \gamma^*$ ($i > \gamma^*$) then $V_j(B) + e^{\lambda|B|} = |B|w(1 - \frac{\hat{\gamma} - \gamma^*}{k_j\Gamma})$ ($= |B|w(1 - \frac{\gamma^* - \hat{\gamma}}{k_j\Gamma})$).*

Proof: We show that this is true for if $n_i = 0$ for all $i < \gamma^*$. The proof when $n_i = 0$ for all $i > \gamma^*$ is analogous.

As all articles are to the right of the consumer:

$$\begin{aligned}
 V_j(B) + e^{\lambda|B|} &= \sum_{i=\gamma_j^*}^{\Gamma} n_i w \left(1 - \frac{2(\gamma_i - \gamma_j^*)}{k_j\Gamma} \right) \\
 &= \sum_{i=\gamma_j^*}^{\Gamma} n_i w + \sum_{i=\gamma_j^*}^{\Gamma} \frac{2wn_i\gamma^*}{k_j\Gamma} - \sum_{i=\gamma_j^*}^{\Gamma} \frac{2wn_i\gamma_i}{k_j\Gamma} \\
 &= |B|w + \frac{2w|B|\gamma^*}{k_j\Gamma} - \frac{2w|B|\hat{\gamma}}{k_j\Gamma} \\
 &= |B|w \left(1 - \frac{2(\hat{\gamma} - \gamma^*)}{k_j\Gamma} \right), \tag{3}
 \end{aligned}$$

using the definition of $\hat{\gamma}$ and $\sum n_i = |B|$ on the third line. ■

Lemma 4 shows that assuming that all articles in a bundle are on one side of a niche's γ^* , that niche's value of that bundle depends *only* on the size of the bundle and the mean category. By Lemma 3, this result applies to the target niches. Our next result shows that for given target niches, there is a unique mean bundle category ($\hat{\gamma}$) that satisfies $V_{c_a}(B) = V_{c_b}(B)$.

Lemma 5 *If $n_i = 0$ for all $\gamma_i < \gamma_a$ and all $\gamma_i > \gamma_b$, then $V_{c_a}(B) = V_{c_b}(B)$ if and only if $\hat{\gamma} = \gamma_a + (\gamma_b - \gamma_a) \frac{1/(k_a\Gamma)}{1/(k_a\Gamma) + 1/(k_b\Gamma)}$.*

Proof: First, if $\gamma_a^* = \gamma_b^*$, then $n_{\gamma_a^*} = |B|$, and $\hat{\gamma} = \gamma_a^*$.

Now assume that $\gamma_a^* \neq \gamma_b^*$. By Lemma 4, $\hat{\gamma} \in [\gamma_a^*, \gamma_b^*]$. By Equation 3, the value of c_b is increasing on this interval, and that of c_a decreasing. Therefore, if there exists a $\hat{\gamma}$ such that $V_{c_a}(\hat{\gamma}) = V_{c_b}(\hat{\gamma})$, it will be unique. To find this $\hat{\gamma}$:

$$\begin{aligned}
 V_{c_a}(B) &= V_{c_b}(B) \\
 |B|w \left(1 - \frac{2(\hat{\gamma} - \gamma_a^*)}{k_a\Gamma} \right) &= |B|w \left(1 - \frac{2(\gamma_b^* - \hat{\gamma})}{k_a\Gamma} \right) \\
 \hat{\gamma} &= \gamma_a + (\gamma_b - \gamma_a) \left(\frac{1}{k_a\Gamma} \right) \left(1 / \left(\frac{1}{k_a\Gamma} + \frac{1}{k_b\Gamma} \right) \right),
 \end{aligned}$$

where the final row is the result of tedious algebra. ■

Therefore, given two target niches, $\hat{\gamma}$ characterizes any interior bundle satisfying $V_{c_a}(B) = V_{c_b}(B)$. By Equation 3, bundle size alone determines value of these bundles, and thus the price. We now show that, given two target niches receiving zero surplus, the surplus received by niches whose γ^* is exterior to a bundle is *independent* of bundle size.

Lemma 6 Given a bundle interior to target niches, the identity of the two target niches, c_a and c_b , is sufficient to determine the surplus of niche c_d if $n_i = 0$ for all $i < (>) \gamma_d^*$.

Proof: Assume that $\gamma_d^* > \hat{\gamma}$. As $surplus(c_b) = 0$, $surplus(c_d) = surplus(c_a) - surplus(c_b)$. As prices and clutter costs will be the same for both niches, we can simplify as follows:

$$\begin{aligned} surplus(c_d) &= |B|w \left(\left(1 - \frac{2(\gamma_d^* - \hat{\gamma})}{k_d \Gamma} \right) - \left(1 - \frac{2(\gamma_b^* - \hat{\gamma})}{k_b \Gamma} \right) \right) \\ &= \frac{2|B|w(\hat{\gamma}(k_b - k_d) + \gamma_b^* k_d - \gamma_d^* k_b)}{k_b k_d \Gamma}. \end{aligned} \quad (4)$$

All parameters positive by assumption, the sign of the surplus is the same as the sign of the outer parentheses of equation 4, the sign of which is determined solely by the niches. The proof for $\gamma_d^* < \hat{\gamma}$, using c_a , is analogous. ■

We can thus tell *a priori* which niches outside of the target niches will purchase the bundle. This is not the case for those niches for whom $\gamma^* \in (\gamma_a^*, \gamma_b^*)$. For example, if $\hat{\gamma} = 3$, then those niches for whom $\gamma^* = 3$ will certainly prefer a bundle of only category 3 to one of the same size with half in category 2 and half in category 4. Our next result shows that surplus of these consumers depends only on the distribution of articles, and not the size of the bundle.

Lemma 7 Given a bundle interior to target niches c_a and c_b , the share of the bundle in each γ_i is sufficient to determine the surplus of niche c_d .

Proof: We show for $\hat{\gamma}$ fixed and $n_i = 0$ for all $\gamma_i \notin [\gamma_a^*, \gamma_b^*]$, niche surplus does not depend on the size of the bundle.

The share in each category is equal to $s_i \equiv \frac{n_i}{|B|}$. Further define $\alpha = \sum_{i=\gamma_a^*}^{\gamma_d^*} s_i$, $\hat{\gamma}_l \equiv \sum_{i=\gamma_a^*}^{\gamma_d^*} s_i n_i$ and $\hat{\gamma}_r \equiv \sum_{i=\gamma_d^*+1}^{\gamma_b^*} s_i n_i$. The value of a bundle for any c_d is:

$$\begin{aligned} V_{c_d} + e^{\lambda|B|} &= |B|w\alpha \left(1 - \frac{2(\gamma_d^* - \hat{\gamma}_l)}{k_d \Gamma} \right) + \\ &\quad + |B|w(1 - \alpha) \left(1 - \frac{2(\hat{\gamma}_r - \gamma_d^*)}{k_d \Gamma} \right) \\ &= |B|w \left(1 + \frac{2((1 - 2\alpha)\gamma_d^* - \hat{\gamma}_r + \hat{\gamma})}{k_d \Gamma} \right) \end{aligned}$$

where the first equation is a result of algebra similar to that in the proof of Lemma 4 and the last line is the result of straightforward but tedious algebra and $\hat{\gamma} = \alpha \hat{\gamma}_l + (1 - \alpha) \hat{\gamma}_r$.

Subtracting $surplus(c_b) = 0$ from $surplus(c_d)$ gives us:

$$surplus(c_d) = \frac{|B|w((\hat{\gamma} + \gamma_d^* - 2\alpha\gamma_d^* - \hat{\gamma}_r)\gamma_b^* + (\gamma_b^* - \hat{\gamma})k_d)}{k_b k_d \Gamma},$$

the sign of which is entirely determined by the sign of the outer parentheses of the numerator, a function of the niches and the distribution of the bundle. ■

We can now characterize the optimal offering for a monopolist given two target niches.

Proposition 1 Given two target niches, c_a and c_b , the monopolist's optimal bundle configuration solves:

$$\max_{|B|} V_{c_a}(|B|) \text{ subject to } V_{c_a}(|B|) = V_{c_b}(|B|),$$

and the optimal price for the bundle is $p = V_{c_a}$.

Proof: Lemma 2 implies that $V_{c_a} = V_{c_b}$. Setting surplus equal to zero, the optimal price must be $p = V_{c_a} = V_{c_b}$.

Given c_a and c_b , the distribution of the optimal bundle must be of a certain form. For any offering, let the number of purchasing niches be $c^\#$. By Lemma 3, we know all articles are interior and are characterized by $\hat{\gamma}$ as defined by Lemma 5. We further restrict our subset to those distributions for which $c^\#$ achieves its maximal value on the set. This highest total demand subset is independent of bundle size and thus price by Lemma 7. As the set for which $c^\#$ is maximal is independent of bundle size, profit is maximized by choosing the bundle size $|B|$ that maximizes $p = V_{c_a}(|B|)$ and allocating this bundle by any of the distributions in our non-empty highest total demand subset. ■

bundle	1	2	3	4	5	6
k high	C	+	+	+	C	C
k med	-	0	+	0	-	-
k low	-	-	C,D	-	-	-

0: Target Niche, Will Purchase
 +: Will Purchase
 -: Will Not Purchase
 C: Purchase depends on niche characteristics
 D: Purchase depends on article distribution

Fig. 2. Consumer demand given two target niches. Each square is a niche defined by $\{\gamma^*, k\}$. All articles are in categories 2,3, and 4, with a mean category of 3 and a price equal to the willingness to pay of the targets.

Figure 2 demonstrates the link between characteristics of the target niche and the surplus that other niches get, regardless of bundle size. First, note that niches with most-favored categories outside of the target niches and smaller breadth of interest than the target niches do not purchase the bundle. Similarly, those whose breadth of interest is not smaller than the targets and whose ideal category is at least as close to 3 as the targets will purchase. Those consumers outside the targets but with higher k (denoted by C) might purchase. Finally, we look at the niche denoted by C, D . Note that only bundles of the form $\{0, \frac{\alpha}{2}|B|, (1 - \alpha)|B|, \frac{\alpha}{2}|B|, 0, 0\}$ are interior and satisfy equality of the targets. Regardless of α , the willingness to pay of the targets depends only upon $|B|$. The willingness to pay of C, D does depend on α . At $\alpha = 0$, this niche has positive surplus. At $\alpha = 1$, this niche receives negative surplus as the sum of differences between

most-preferred category and bundle category is the same as it is for the targets, yet this niche has a lower k .

Proposition 1 tells the firm how to find the optimal bundle conditional on any two target niches, c_a and c_b . By Lemma 2, the optimal offering for a monopolist will consist of a pair of target niches. To find the optimal target niches, the firm can find the optimal bundle and associated profit for each candidate pair of target niches, and then from this set choose the bundle and price that yield the highest overall profit.

3.2 Multi-firm Equilibrium

In this section, we analyze multiple interacting firms. Our goal is to characterize, to the extent possible, equilibrium outcomes.

Define the offerings of the M firms by \vec{N} , an $M \times \Gamma$ matrix with $N_{m,i}$ the number of articles from firm m in category i , and \vec{p} , a $M \times 1$ column vector where p_m is firm m 's price. Denote producer i 's profits as $\pi_i(\vec{N}, \vec{p})$. The next lemma gives a strong characterization of any equilibrium resulting from pure strategies.

Lemma 8 *In any pure-strategy Nash equilibrium, $\pi_i(\vec{N}, \vec{p}) = \pi_j(\vec{N}, \vec{p}) \quad \forall i, j \in \{1, 2, \dots, M\}$.*

Proof: If $\pi_i(\vec{N}, \vec{p}) > \pi_j(\vec{N}, \vec{p})$, then j could deviate by offering $\{\vec{n}_i, p_i - \epsilon\}$ for ϵ arbitrarily small. Therefore, $\{\vec{N}, \vec{p}\}$ cannot be an equilibrium. ■

Lemma 8 is a direct result of firms being unconstrained in the categories in which they can offer items. We shall show that such an equilibrium exists only if the category space is wide enough to accommodate M local monopolists. Letting Γ^m correspond to the number of γ^* -types that would be served by a monopolist, we define a local monopolist as one whose offering is optimal for $\Gamma = g < \Gamma^m$. The next proposition gives conditions sufficient for the existence of a multi-firm Nash equilibrium.

Proposition 2 *A pure strategy Nash equilibria exists if λ, w, Γ and \hat{k} are such that $\Gamma^m \leq \frac{\Gamma}{M}$.*

Proof: Assume a monopolist in a market characterized by the parameters would optimally serve no more than $\frac{\Gamma}{M}$ distinct γ^* -types. Let firm i configure the optimal monopolist bundle such the left-most niche served has $\gamma^* = (i - 1) * \Gamma^m + 1$. As all receive the monopolist profit, none has an incentive to deviate. ■

We now look at the case where all firms cannot act as monopolists. The next two lemmas greatly restrict the space of possible equilibria outcomes.

Lemma 9 *In no pure-strategy Nash equilibria are any consumers indifferent between the offerings of the two firms.*

Proof: An ϵ decrease in price by either firm leads to a discrete increase in demand, as all consumers in the previously indifferent niche now all purchase from the same firm. ■

Lemma 10 *In any pure-strategy Nash equilibrium, a niche gets positive surplus from at most one producer's offering.*

Proof: Assume not. By Lemma 8, we know that all firms earn equal profit. Let firms i and j provide positive surplus to the niche c_d . By Lemma 9, c_d is not indifferent. If c_d purchases from j , j offering $\{\vec{n}_i, p_{i-\epsilon}\}$ increases profits. ■

Therefore, all niches with the same γ^* that do purchase will purchase from the same firm, implying that each firm serves a unique subset of γ^* -types. We can now provide necessary conditions for the existence of a Nash equilibrium.

Proposition 3 *A pure-strategy Nash equilibrium exists only if each firm is a local monopolist such that $\pi_i = \pi_j \forall i, j$.*

Proof: By Lemma 10, each firm serves a unique subset of γ^* -types. Given this subset, the firm selects the optimal offering. By Lemma 8, all firms' profits must be equal. ■

Thus the addition of firms in our environment need not lead to firms directly competing over the same consumers. However, in searching for a profitable niche, even the monopolist producer needs to search over a very large product space to find an optimal bundle and price as detailed in Section 3. The addition of other firms aggravates this already difficult problem. When a producer uses profits as a guide to finding a good bundle, these profits will be affected by the actions of other learning producers. The ability to find a profitable niche and avoid competition requires a great deal of knowledge not only about both consumer preferences and the strategies of other firms.

4 Zero-Knowledge Learning

In the previous section, we showed that the optimal strategy for duopolist producers, when they know the distribution of consumer preferences and understand the economic structure of the problem, is for each to target a separate consumer niche, acting as a local monopolist. We expect that firms (and their computational agents) will not typically have complete information about the distribution of consumer preferences, nor about the economic structure of the problem. Therefore, we now examine what happens when the producers are learning what bundle to offer and how to price it. We are interested in examining whether learning producers can locate the equilibrium solution, and in characterizing the extent to which economic knowledge can help in locating profitable niches more easily, particularly in the face of competing producers.

We model the learning of each producer using a genetic algorithm (GA).¹¹ We chose to use GAs primarily because of their effectiveness at learning in high-dimensional, nonstationary problems. GAs are also commonly used [1,4] to study the adaptive behaviors of agents in a multi-agent systems. We cannot claim that a GA is the optimal algorithm for this problem, or even that the parameters or encoding we have chosen are ideal. Our goal is to identify the sorts of knowledge that will help a producer locate profitable niches, given that it has adopted a widely-used off-the-shelf learning algorithm, rather than determining and optimizing a particular algorithm to fit this particular problem.

We encode the number of articles in each category as a gene. A bundle is then represented as a chromosome: a list of the articles in each category, plus a price, encoded

¹¹ Goldberg [8] provides a comprehensive introduction to genetic algorithms.

as a binary string. The *fitness* of a bundle is the profit that that bundle receives when it is offered to the consumer population. We refer to the set of bundles that a producer's GA is currently evaluating as its *pool*. Note that each producer is learning simultaneously; this means that each producer's target (the function it is optimizing) is potentially changing. It is helpful in visualizing this learning process to realize that each producer is learning a mixed strategy, represented by the elements of its pool.

In the following experiments, we consider the following situation: a population of 100 consumers, with k (the fraction of categories valued) drawn uniformly from $[0.2, 0.8]$ and $\Gamma = 9$. Each consumer's γ^* is drawn from a uniform distribution, and all consumers have valuation $w = 10$ for their favorite category. This leads to an equilibrium in which each of two producers acts as a local monopolist, selling a bundle of 2 categories, with 113 articles in one category, 169 articles in the adjacent category, and a price of 1286. In expectation, if k is completely uniformly distributed, each producer will earn a profit of 36008.

We begin with the situation in which both producers are zero-knowledge producers. Each producer begins with a pool of 50 bundles in which the number of articles is drawn from $[0, \hat{n}]$ and the price is drawn from $[0, 2048]$. In this experiment, as well as the following ones, \hat{n} is set to 128.¹² In Section 5.5, we vary this parameter and discuss the implications of this choice. Once the experiment has begun, the GA allows producers to select any possible bundle and price; a producer's only feedback is the profit that a bundle earns. The results of this experiment are shown in Figure 3. The error bars indicate 95% confidence intervals for the data; that is, 95% of the data points for that iteration fall within the indicated range. As we see from Figure 3, when the producers are asked to locate equilibria without any other knowledge, they fail spectacularly. There are two potential reasons for this. First, the search space is extremely large. If we assume that a producer can offer 256 articles in a bundle, the maximum price is 2048, and prices are integers, then there are $256^9 \times 2048 \approx 9.6 \times 10^{22}$ bundles. Second, a large fraction of these bundles produce no profit. A random sample of one million bundles in this landscape reveals that only 55, or 0.005%, of the bundles have non-zero profit. This leaves a learning algorithm such as a GA, which depends upon a gradient, with little help in escaping zero-profit plateaus and locating optima. Even in those cases where producers are able to locate positively-valued bundles, the bundles they find are far from optimal. This indicates that the search space is pocked with plateaus and local optima.

It seems unlikely that information goods producers would have absolutely no knowledge of the goods they were selling or the preferences of the consumer population. More likely, they would have some (potentially uncertain) information that they could apply to their learning problem. In the following section, we examine the consequences of applying different sorts of economic knowledge on this learning problem.

5 Injecting Economic Knowledge

In real markets, even those in which consumers are anonymous and the consumer population is rapidly changing, producers typically have some knowledge about their learn-

¹² Note that \hat{n} limits the maximum number of articles in a category at the beginning of an experiment; the maximum possible number of articles in a category was set to 1024.

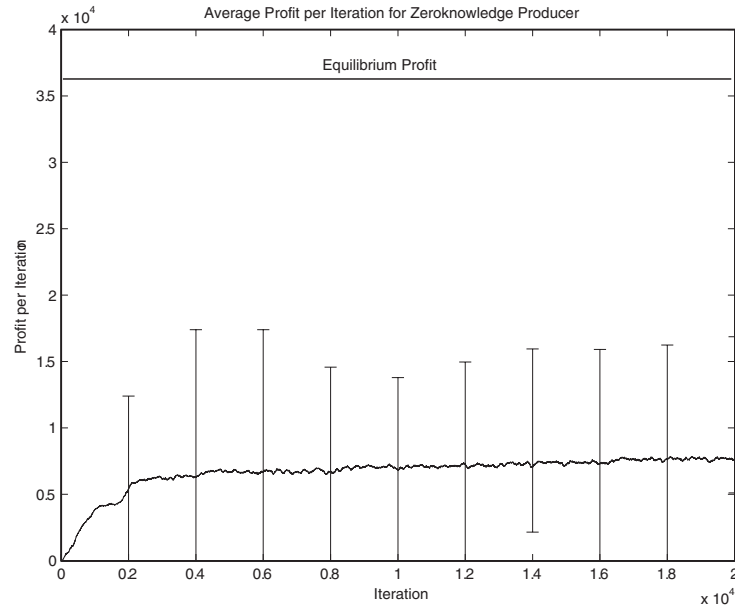


Fig. 3. Moving average of profit for zero knowledge Producers. Producers can only capture a small fraction of the surplus.

ing problem. At the very least, they know that their problem is an economic one, and so it will obey some simple criteria, such as that demand is non-increasing *ceteris paribus* as price increases.

This knowledge allows a producer to prune its search space in a way that a knowledge-free algorithm such as a GA, or optimization algorithms such as amoeba, which we have used in previous work [3], cannot. This leads us to suspect that results generated with zero-knowledge producers may be overly naive.

The experiments in this section map out some of the space between the full-information producers described in Section 3 and the zero-knowledge producers described in Section 4. In particular, we examine three ways in which economic knowledge can transform a learning producer's search problem: it can reduce the size of the search space; it can transform the landscape, introducing a gradient that aids adaptive algorithms such as GAs; or it can allow the producer to start in a more effective region of the search space.

5.1 Equilibrium Stability

We begin with producers that have near-perfect information. In this experiment, each producer's pool is seeded with all elements in the pool equal to the equilibrium solution. This provides the equivalent of the perfect information used to derive the analytic results in section 3. On every iteration, we increase the GA's mutation rate until it reaches 30%.




Fig. 4. Moving average of profit for perfect-information producers with mutation. Despite large amounts of mutation, producers are able to maintain the equilibrium solution.

This is meant to represent the situation where producers know what they “should” do, but occasionally make mistakes. It is also intended to help understand the nature of the equilibrium basin. It was not clear whether the difficulty in the zero-knowledge case was due to the size of the search space or the structure of the equilibrium. When producers are pushed off the equilibrium, are they able to easily return? Figure 4 shows the results of this experiment. Once again, error bars indicate a 95% confidence interval.¹³ As we see in Figure 4, the equilibrium solution is quite resistant to the introduction of mutation. The error bounds show that mutation can push the producer quite far from the equilibrium solution, yet it is able to return to the equilibrium. This implies that the basin of attraction around the solution is quite broad; even when producers are pushed away from it, there is a natural gradient they can follow to return back to the equilibrium.

Another way to understand this is to consider the change in each producer’s search space. The zero-knowledge agent searched over all possible bundles, which produced a vast space overwhelmingly filled with bundles that yielded zero profit. In contrast, the initial search space here consists of one bundle: the equilibrium bundle. As we increase the mutation rate, the GA considers more bundles, and so the search space will grow. At its highest, the mutation rate reaches 30%, meaning that there is a 0.3 probability

¹³ Since the experiments are conducted on a finite, randomly generated population, it is possible that, in a particular run, the equilibrium solution will differ slightly from the theoretical solution, allowing producers to earn slightly more than the theoretical profit.




Fig. 5. Moving average of profit for producers with uncertainty over consumer valuation for each category. Producers are still able to capture a large fraction of equilibrium profit.

that a bundle in a producer's pool will be mutated. In expectation, 15 new bundles will be added to the search space every time mutation is applied. At this rate, it would take approximately 6.6×10^{21} rounds of mutation, or 3×10^{23} iterations (since mutation is applied after all 50 bundles in a pool are evaluated) to sample all possible bundles.

5.2 Uncertainty over Consumer Valuation

In this section, we further weaken each producer's knowledge about the consumer population. In particular, we assume that they no longer know exactly how many articles to sell in each category. To achieve this, we seed each producer's pool with copies of the equilibrium bundle, and then apply Gaussian noise (mean 0, standard deviation 50) to the number of articles in each category for each of the bundles in the pool. This gives a producer a pool of approximate solutions. We continue to apply mutation as in the previous experiment. The results of this experiment are shown in Figure 5.

As we can see in Figure 5, applying significant amounts of Gaussian noise does not affect each producer's ability to return to the equilibrium solution. Although average profits decrease slightly, and we see the introduction of an initial transient in which producers are converging on a solution, the application of noise is not enough to push producers out of the basin of attraction of the equilibrium.

In this case, we have not reduced the total search space, since it is possible for Gaussian noise to produce any possible bundle. Instead, each producer starts its search in a profitable area of the search space, namely one that has a gradient. Random sampling of one million initial bundles generated according to this procedure shows that 27.5% of them have a non-zero profit. In a pool of 50 initial bundles, 13 of them will have non-zero profit, which provides the GA with a sufficient gradient for learning.

5.3 Adjacent Categories

In the previous experiment, we applied noise to each producer's solution, but we still seeded the producers with the proper coordination knowledge. That is, producer 1 was seeded with bundles that tended to have large numbers of articles in categories 2 and 3, and producer 2 was seeded with bundles that tended to have large numbers of articles in categories 7 and 8. In this experiment, we remove that knowledge. However, we allow producers to retain the knowledge that the consumer valuation of categories exist on a line, and so they are best off offering adjacent categories. In this sort of environment, 'general interest' bundles containing a wide variety of articles are not highly valued.

We implement this by seeding each producer's pool with the equilibrium bundle and applying noise, as above, and then rotating the categories in each bundle by a random amount. This produces a pool in which bundles tend to have large numbers of articles in adjacent categories, but the favorite category is distributed uniformly throughout the pool. Again, mutation is applied. The results of this experiment are shown in Figure 6. We see that the producers in this experiment suffer a loss in profit when compared to the producers shown in Figure 5. We have not changed the nonzero fraction of the search space relative to the previous experiment, but we have removed the initial coordination of solutions between producers. This can create inadvertent competition. Since the producers no longer start out in separate areas of the search space, there are cases in which they mis-coordinate and offer the same bundle. In fact, if each producer is choosing two adjacent categories at random, there is a $\frac{3}{8}$ chance that the producers will offer at least one category in common. Given this, it would seem that the producers still seem to do a relatively good job of coordinating, which implies that there is a strong 'learning gradient' toward acting as a local monopolist.

5.4 Bundle Diversity

In the following experiment, we further weaken each producer's knowledge by removing the assumption that producers know that categories are arranged on a line. Instead, we assume that they only know the approximate diversity of the bundle; that is, how many categories they should offer articles in. We implement this by seeding each pool with copies of the equilibrium bundle, applying Gaussian noise, and then permuting the articles offered in each category within each bundle. This produces a pool of bundles that tend to have two categories that contain a large number of articles, but which are not necessarily adjacent. The results of this experiment are shown in Figure 7. The profits in this experiment are very similar to those shown in Figure 6. This serves as an indicator that the knowledge that categories are arranged on a line may not be that essential to producer performance. Instead, knowing the number of categories to offer articles




Fig. 6. Moving average of profit for producers that know that category values are arranged on a line. Lack of coordination lowers profits slightly.

in may be a more useful piece of knowledge. Interestingly, by removing the adjacency knowledge, we reduce the probability that two random bundles will compete directly (to $\frac{17}{81}$, or 21%). However, random sampling of one million random bundles reveals that only 17% of them yield positive profit. So the reduction in potential competition is canceled out by the increasingly poor starting bundles.

An alternative explanation for why adjacency knowledge doesn't produce a significant improvement in performance might lie in the way in which GAs search for solutions. GAs build up solutions by composing contiguous sub-solutions (sometimes referred to as building blocks) that tend to perform well. This produces a selection bias toward solutions in which adjacent genes (or numbers of articles, in our case), both produce a high fitness. So the fact that the articles are encoded on a line on the gene may be enough for the producer to exploit that relationship within the model of consumer preferences, even without being given that knowledge in the form of initial conditions.

5.5 Total Bundle Size

As we have pointed out, the greatest challenge that these producers face in locating a suitable bundle to offer is the incredibly large search space. Many of the previous experiments have focused on ways in which a producer can apply economic knowledge to reduce the size of the search space. In this experiment, we allow each producer to exploit economic knowledge to instead start its search in useful areas of the search




Fig. 7. Moving average of profit for producers that know bundle diversity. Even though producers no longer know that categories are arranged on a line, profits remain at a level consistent with the previous experiment.

space. We do this by modifying the number of articles in each category that a producer will initially offer. Smaller initial bundles will be more likely to yield nonzero profit, but may potentially trap the producer in a local optimum.

In Section 4, we seeded the producers' pools with bundles containing a random number of articles in each category, where this number was drawn from $[0, \hat{n}]$, with $\hat{n} = 128$. In this experiment, we begin with the zero-knowledge producers from section 4. We then vary \hat{n} , thereby changing the region of the search space in which the producers can start. In Figure 8, we show the results of experiments in which the number of articles in each category of a producer's initial bundle were drawn from $[0, \hat{n}]$ with \hat{n} set to $\hat{n} = 25, 64, 128, 192,$ and 256 , respectively. As we can see from this experiment, starting with a smaller bundle definitely allows a producer to achieve better profits. In addition, when \hat{n} becomes too large (greater than 128), profits are either very low or zero, implying that the producers never find their way out of zero-profit regions of the search space. However, even a small \hat{n} is not enough to move the producers all the way to the equilibrium. This implies that there are local optima within the landscape in which producers can become stuck. In addition, we can see that the transient is different for the $\hat{n} = 25$ and $\hat{n} = 64$ cases. In the $\hat{n} = 25$ case, the initial bundles were simply too small, and so time (and a larger number of mutations) was needed to make the bundles large enough to extract significant profit. Beyond $\hat{n} = 128$, most of the initial bundles occur in zero-profit areas of the landscape, and so producers do quite poorly. This tells us that

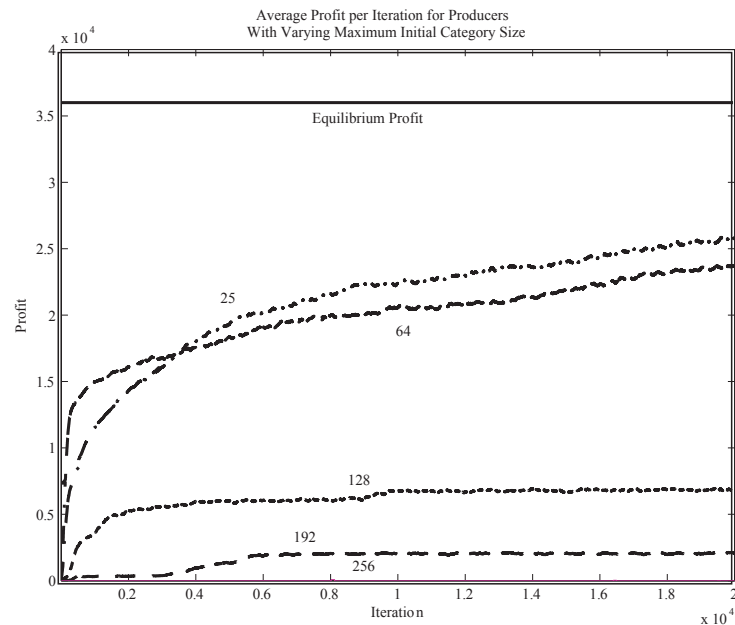


Fig. 8. Moving average of profit for producers with different initial numbers of articles in each category. Smaller initial numbers of articles allow producers to locate solutions, but overly small bundles introduce longer learning transients.

knowing something about where to start one's search is a useful piece of knowledge, but by itself it is not a replacement for more detailed economic knowledge about the behavior of either consumers or competing producers.

5.6 Introducing a Gradient

Another way in which producers can make their search problem easier is to transform the landscape so that it contains more of a gradient. One way to do this is to get feedback from the consumer population about how much they value a bundle, as opposed to only what they purchase. This allows a producer to distinguish between those bundles that receive zero profit because they're slightly undesirable and those bundles that receive zero profit because they're extremely undesirable.

In this experiment, we provide producers with a gradient by altering the way in which the GA assigns fitness to a bundle. Rather than using the profit that a bundle earns, we use the valuation that the five most-satisfied consumers assign to the bundle, even if some or all of these are negative.¹⁴ While the size of the search space does not

¹⁴ The assumption that producers can discover how negatively a bundle is valued is admittedly a bit artificial, although a clever producer might be able to estimate this quantity, either through offering rebates or cash to consumers, or packaging this bundle with another good whose value is known and observing the difference in demand.

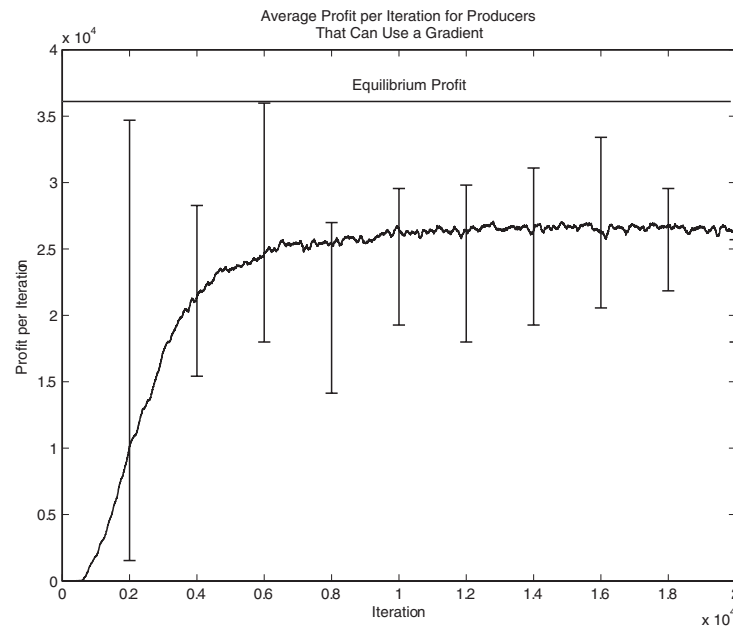


Fig. 9. Moving average of profit for producers with a gradient.

change at all, the underlying landscape is transformed, as there are (almost) no areas of zero profit. Instead, most of the large zero-profit plateaus will now be populated with valleys corresponding to negative valuations. This provides a gradient-using algorithm (such as a GA) with the feedback needed to learn effectively. By using only the valuations of the most-satisfied consumers, a producer is encouraged to adapt its bundle to better satisfy those consumers that are most inclined to buy its bundle, as opposed to consumers who would need drastic changes in order to purchase.

One potential complication that comes from using consumer surplus, rather than profit, as a measure of bundle fitness is that bundles that are priced at zero will be highly valued by the consumer (since they yield high surplus), even though they are not particularly useful solutions for the producer. We avoid this problem by fixing prices for this experiment and having producers learn only bundle contents. This solution is only partly satisfactory; a more comprehensive approach would combine both profit and consumer satisfaction in estimating bundle fitness.

Figure 9 shows the results of this experiment. There is still a significant transient period when producers are searching through low-profit areas of the landscape, but after about 400 iterations, a producer is able (on average) to extract about 70% of the equilibrium profit, which is a distinct improvement over the experiments in Section 4, where a producer was unable to extract any significant profit. As in the previous experiment, this tuning of the learning algorithm helps, but not as much as the application of specific domain knowledge, which can be used to drastically prune the search space.

6 Conclusions and Future Work

In this paper, we have examined the problem of producers determining where to locate in a high-dimensional information-goods product space. When the producers have perfect information, the pure-strategy equilibrium solution is for each producer to act as a local monopolist, serving a separate consumer niche. However, discovering those niches is quite difficult for naive learners, as the search space is incredibly large. We argue that learning agents should take advantage of the economic knowledge implicit (and explicit) in the problem. This can allow them to reduce the search space, transform the search space by adding a gradient, or start in promising locations.

We have not examined one particularly useful piece of economic knowledge: the fact that a producer is competing with other producers. Adding a layer of reflectivity, in which the GA's feedback is modulated by the distance between a producer's bundle and those of other producers, remains a topic for future work. In addition, the use of gradient information could be made more sophisticated through the introduction of a multi-objective fitness function that weights both profit and valuation of a bundle.

References

1. J. Andreoni and J. Miller. Auctions with artificial adaptive agents. *Games and Economic Behavior*, 10:39–64, 1995.
2. C. H. Brooks, E. H. Durfee, and R. Das. Price wars and niche discovery in an information economy. In *Proceedings of the 2nd ACM Conference on Electronic Commerce (EC-00)*, pages 95–106, Minneapolis, MN, October 2000.
3. C. H. Brooks, R. S. Gazzale, R. Das, J. O. Kephart, J. K. MacKie Mason, and E. H. Durfee. Model selection in an information economy: Choosing what to learn. *Computational Intelligence*, 18(4), 2002.
4. M. Chang and J. E. Harrington. Centralization vs. decentralization in a multi-unit organization: A computational model of a retail chain as a multi-agent adaptive system. *Management Science*, 46(11):1427–1440, November 2000.
5. J. C. Chuang and M. A. Sirbu. Network delivery of information goods: Optimal pricing of articles and subscriptions. Available at: <http://ksgwww.harvard.edu/iip/econ/chuang.html>, 1998.
6. P. Dasgupta and R. Das. Dynamic pricing with limited competitor information in a multi-agent economy. In *Proceedings of the Fifth International Conference on Cooperative Information Systems*, 2000.
7. S. A. Fay and J. MacKie Mason. Competition between firms that bundle information goods. In *Proceedings of the 27th Annual Telecom Policy Research Conference*, Alexandria, VA, September 1999.
8. D. E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison Wesley, Boston, 1989.
9. J. O. Kephart, J. E. Hanson, and J. Sairamesh. Price and niche wars in a free-market economy of software agents. *Artificial Life*, 4:1–23, 1998.