Information Bundling in a Dynamic Environment

Christopher H. Brooks* **Rajarshi Das†** **Jeffrey O. Kephart†**
Jeffrey K. MacKie-Mason‡ **Robert S. Gazzale§** **Edmund H. Durfee***

Abstract

Markets for digital information goods provide the possibility of exploring new and more complex pricing schemes, due to information goods' flexibility and negligible marginal cost. In this paper we compare the dynamic performance of price schedules of varying complexity under two different specifications of consumer demand shifts. A monopolist producer employs a simple direct-search method that seeks to maximize profits using various price schedules. We find that the complexity of the price schedule affects both the amount of exploration necessary and the aggregate profit received by a producer. The size of the bundle offered, the rate of population change, and the number of iterations a producer can expect to interact with a population in total all affect the choice of schedule. If the number of iterations is small, a producer is best off randomly choosing a high-dimensional schedule, particularly when the bundle size is large. As the number of interactions between the producer and a given consumer population increases, then two-parameter schedules begin to perform best, as their learnability allows the producer to find highly optimal prices quickly. Our results have implications for automated learning and strategic pricing in non-stationary environments arising from changes in the consumer population, in individuals' preferences, or in the strategies of competing firms.

1 Introduction

The distinguishing characteristic of information goods is their negligible incremental cost of reproduction and distribution once the good is initially produced. As a producer can almost costlessly package these goods in a wide variety of configurations, markets for digital information goods provide the possibility of exploring new and more complex pricing schemes. With multiple items (say, articles) to sell, a producer might offer a single price per article (linear pricing), sell the whole lot for a subscription price (bundling), or any number of other variants. Collectively, the family of schemes possible for multiple items (units of a single good, or combinations of different goods) are known as "nonlinear" pricing, as the revenue received is a nonlinear function of the number of items sold. Bundling is one example of nonlinear pricing that has received substantial recent attention. Conditions under which bundling information goods is desirable are well-recognized in the economics literature [Chuang and Sirbu, 1999; Bakos and Brynjolfsson, 1999].

Prior authors [Wilson, 1993] have calculated the benefit of bundling and other nonlinear pricing schemes based on the assumption that the producer knows the distribution of consumer preferences. Typically producers do not have complete knowledge of all relevant parameters of this distribution, and can only offer an estimate of the profit-maximizing price schedule. Thus, at any point in time, there are two returns from an offered price schedule: a current profit, and incremental information about the underlying consumer preferences which may be used to extract higher profits by changing the price scheme parameters — or even the price scheme itself — in the future.

The pricing problem is then inherently dynamic as information acquired in the current "period" affects actions in later periods. In a given period the producer generally does not offer the price scheme expected to maximize current or *myopic* profits given its current information about consumer preferences. Instead, it deviates from this strategy in order to set prices that constitute a better experiment by revealing more information about the true underlying preferences. While this experimentation clearly has long-run benefits, there are short-run opportunity costs in terms of the profits foregone by not acting myopically. The firm must therefore select pricing parameters that balance current profits, *exploitation*, and learning the nature of consumer demand, *exploration* [Thrun and Møller, 1992]. The pricing problem of the information goods producer is even more acute, as the relative advantage of a pricing schemes depends on its *potential profitability* as well

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

†Institute for Advanced Commerce, IBM Research, P.O. Box 704, Yorktown Heights, NY 10598, {rajarshi, kephart}@watson.ibm.com

‡School of Information and Department of Economics, University of Michigan, Ann Arbor, MI 48109, jmm@umich.edu, (734) 647-4856

§Department of Economics, University of Michigan, Ann Arbor, MI 48109, rgazzale@umich.edu

as its *learnability*.

The identified benefits of exploration for the firm rest on various assumptions. First, it is assumed that the underlying preferences to be learned are stationary. In addition, the standard approach assumes that a producer will have a long period of interaction with a population once the optimal prices are learned, so learning costs can be absorbed into long-run profit. This clearly need not be the case. For example, in many instances the set of information goods being offered by various producers is rapidly changing. This will result in an observed change in consumer preferences, represented as willingness-to-pay per item, will appear to change even if the underlying preference structure has not changed. Additionally, competing producers will regularly change their offerings, causing consumer preferences for a given producer's goods will shift. Finally, individual consumer tastes do evolve over time, as does the composition of heterogenous agents in the market.

Our work is motivated in part by the relative ease of changing price schemes for electronic goods: price listings are stored in electronic databases, and communicated through the data network. Powerful computational resources can be applied to analyzing consumer information in order to estimate improved pricing schemes ("data mining" is a prominent example of this). Likewise, purchasers — even retail customers, but certainly large business buyers — have cheap computational power to assist them in analyzing and responding to changing and possibly complex pricing schedules.

In this paper, we investigate the learning problem for the provider of information goods in a changing environment. In particular, we investigate whether or not there is a relationship between the complexity of a pricing schedule and the provider's ability to quickly learn consumer preferences. Further, we explore the relationship between the costs of experimentation and the complexity of the price schedule. In addition to the usual trade-off between exploitation and exploration for a price scheme (its *learnability*), we also are concerned with the robustness of a pricing scheme to changes in consumer preferences (its *adaptability*). We explore two models of consumer change, each of which correspond to different scenarios about how a consumer population might evolve.

In [Brooks *et al.*, 1999], we examined a version of this problem in which a monopolist learned a pricing schedule while facing a static population. We argued that the transient profits were important since the population could change. In this paper, we make that change explicit and examine the performance of price schedules of varying complexity.

In the very short run, learning a simpler schedule is not beneficial – choosing random parameters in a more complex schedule tends to yield more profit. However, in the medium and long run, the benefits of learning with a moderately complex two-parameter schedule outweigh the higher steady-state profit that can be extracted by an optimized complex schedule. In this case, the ease of learning the simple schedule wins out. The size of the bundle being offered also plays a role in selecting a schedule; fully nonlinear pricing schedules tend to perform better by comparison when the number of articles offered by a firm is large.

1.1 Economics of learning

The use of prices to learn about consumer preferences is not new in the economics literature. Grossman, Kihlstrom and Mirman [1977] were among the first to quantify the extent to which a firm may experiment with prices that are not myopically optimal, thus giving up short-run profits, in order to learn the slope of the demand curve it faces. Subsequent authors have attempted to generalize the problem and identify the conditions under which "complete" learning will occur in the long run in stationary environments. For example, Aghion, Bolton, Harris and Jullien [1991] determine such conditions, which include payoff functions that are smooth and quasi-concave, but note that neither adequate nor inadequate learning is generic.

The existing literature almost universally assumes that the preferences to be learned are stationary. In a rare exception, Keller and Rady [1999] identify the path of optimal actions for a monopolist facing an unknown demand curve that alternates between two states according to a continuous time Markov process. They find that for a given level of noise and a given discount rate, low rates of state changes will lead to "extreme" experimentation where the monopolist tracks the true state rather well, while higher rates of state changes will lead to "moderate" experimentation where the monopolist chooses actions near the action associated with the long-run average state and thus learns very little of the current true state of the world. They find that the transition from one experimentation regime to the other is discontinuous.

In our research we study provider price learning in environments that exhibit less restrictive forms of non-stationarity. Second, because we focus on information goods, we allow the provider to consider a variety of nonlinear pricing schemes, rathern than merely the linear scheme (a uniform price for each goods sold). Under these conditions, we are unable to obtain analytic solutions. We instead use simulation to answer our questions, guided by analytical solutions for simplified versions of the problem.

1.2 Machine Learning in Nonstationary Environments

Within the machine learning community, there has recently been an emphasis on learning nonstationary functions. Early examples of this include reinforcement learning algorithms such as Q-learning and TD-learning [Sutton and Barto, 1998], which can track a nonstationary function, and signal-processing algorithms such as Parzen windows [Duda and Hart, 1973], which can learn a changing probability distribution without the use of a parametric model. While these methods are able to successfully learn a nonstationary fuction in the limit, they require a large amount of data (a problem when data collection is costly), and neither explicitly addresses the exploration/exploitation tradeoff.

Strategies for exploring optimally have also been examined in machine learning [Thrun and Møller, 1992] and optimal control [Fe'ldbaum, 1965]. Typically these approaches either use a different definition of optimality (e.g., each data point is equally costly) or more restrictive assumptions about the nature of the problem to be optimized, such as stationarity.

There has also been work within the multiagent systems community on learning to interact with other agents. Vidal and Durfee [1998] show how agents can use nested models to more easily predict the actions of another agent. Wellman and Hu [1998] examine multiagent learning in a 2x2 game and shows how agents can simultaneously learn a Nash equilibrium strategy. Unlike our research, this work focus primarily on the long-run properties of the agent system and does not specifically address the costs of learning.

2 Model

We model a population of consumers served over time by one producer of information goods. The producer has N new and unique information items, called articles, to offer in each discrete time period. Consumer preferences (values) for articles differ across articles for a given consumer, and across consumers. Given these preferences, each consumer will buy some number of articles when the producer offers them according to some price schedule. Thus, the producer receives a revenue (actually, incremental profit because we assume incremental costs are zero) associated with a specific price schedule. However, the producer does not know the underlying distribution of consumer preferences, and thus must learn the mapping from price schedules to profits. In this section we describe consumer preferences and optimizing behavior, and the specification of the provider's problem.

2.1 Consumers

Consumers are risk-neutral expected utility maximizers. Each purchases the set of articles that yields the largest nonnegative expected surplus according to his or her individual preferences.

Each consumer wants at most one copy of any particular article. A consumer does not value all articles equally, and a particular article might be valued differently by different consumers. We model consumers using a simple two-parameter model introduced by Chuang and Sirbu [1999]. Preference heterogeneity across consumers is obtained by allowing them to have different parameter values. This formulation has the advantage of being analytically tractable while still providing significant nonlinearities in aggregate consumer demand when consumers are heterogeneous.

The two parameters are w, a consumer's value for its most-preferred article, and k, the fraction of the N articles available for which the consumer has a positive valuation. The valuation $V_j(i)$ of the ith most-preferred article by consumer j is a linear function of these variables, expressed by:

$$V_j(i) = \begin{cases} w_j(1 - \frac{i-1}{k_j N}) & \text{if } i-1 \leq k_j N \\ 0 & \text{if } i-1 > k_j N. \end{cases} \quad (1)$$

If the firm offers a price schedule in which the total payment depends only on the number of articles purchased by consumer j, $P(n_j)$, then consumer j's surplus from reading the n^* most preferred articles is the aggregate value less any payments made to the producer, $S_j = \sum_{i=1}^{n_j^*} V_j(i) - P(n_j^*)$. Each consumer chooses n_j^* to maximize S_j.

For the current version of this paper we limit consumer heterogeneity by assuming that w, the value of the most preferred article for each consumer, is the same for all consumers. Consumers differ in their values of k_j, which are distributed uniformly between 0 and $\overline{k}$. The probability density of article values is thus $f(k) = 1/\overline{k}$.

The consumer population varies over time, which is why the producer must be concerned about transitional performance; it cannot count on long-run profits to absorb any learning costs. We consider two separate models of consumer population change, each of which is plausible in different scenarios. In the first, each consumer's k is drawn from a fixed distribution. Periodically, the consumer population is removed and replaced with a new set of consumers drawn from the same distribution. The nonstationarity arises from differences in the particular realizations of each population. This would be an appropriate model for a setting in which subsets of a consumer population interact with a producer for a short period of time and then are sated and replaced by new consumers. The second model incorporates change in the underlying population. The distribution from which consumers are drawn drifts according to a random walk. This produces a model in which consumer preferences drift over time, and avoids the solution in which the producer simply learns the mean of the underlying distribution.

By exploring different of models of population change, we hope to both identify strategies that are effective for different sorts of nonstationarity and avoid generating solutions that are dependent upon a particular underlying model. In addition, we hope to gain further insight as to what sorts of models of consumer change actually merit further exploration.

2.2 Producer

A monopolist produces N articles in each period. Its objective is to maximize cumulative undiscounted profits.[1] Following the standard stylized assumptions for information goods, the marginal cost of duplicating and delivering any article is zero. Likewise there is no cost to bundling articles together to sell as a package.

The producer cannot track individual consumers across transactions, and is either unable to observe the number of articles read by a specific consumer, or to gain any advantage by using such information. This implies that a producer may freely choose from the family of price schedules described in Section 3, but must offer all consumers the same schedule. Therefore a price schedule can be represented by $P(n)$, which gives the total payment a customer makes if she purchases n articles.

We assume that the producer does not know the distribution of w and k across the consumer population, and in fact does not even know that preferences take the form given in (1), above. In a given period the producer announces a price schedule $P(n)$ and receives a profit.

We can evaluate a given price schedule by comparing its results with the maximal profit available. The maximum social surplus would be obtained if each consumer purchased

[1] Discounting can be incorporated without changing the nature of our results.

$k_j N$, as the marginal cost of production and distribution is assumed to be zero. Individual consumers would obtain a gross surplus of $\frac{wk_jN}{2}$, or $\frac{w\bar{k}N}{4}$ per consumer in expectation. A firm that could observe each consumer's k_j could extract this entire surplus. It would perfectly price discriminate by making a take-it-or-leave-it offer tailored to each individual for the full surplus. Since this is the maximal possible social surplus, this case serves is the upper bound on the maximum profit that a less-informed provider could earn.

In the next section we describe the price schedule families a producer may choose from. Each family has one or more parameters that are set by the provider. It is over these parameters that the provider explores in order to increase cumulative profits.

3 Price schedules

In this section we briefly characterize the pricing schedules which are available to the producer in our model, and analyze the performance of each when there is no value from learning. This allows us to present a hierarchy of complex price schedules and to separate their differing abilities to produce profits from the costs and benefits of learning. In Section 5 we introduce the opportunity to learn consumer preferences in a nonstationary environment, and study the interaction between learning and pricing choices.

We assume that a producer has access to a set of N articles that it can offer for sale to a consumer population. Producers then have the problem of first selecting a price schedule and then the prices within that schedule. As previously indicated, the producer must offer the same schedules to all consumers. Typically, the schedules considered for the sale of physical goods are limited to either pure bundling or linear pricing, both of which use just a single price parameter. However, the range of feasible choices is much greater for information goods, as they have high fixed or first copy cost, but negligible marginal production and distribution costs. For example, once a single digital copy of a good such as a newspaper article is produced, the marginal cost for additional copies is essentially zero. To explore this larger space of possible price schedules, we study a scenario in which a producer of information goods considers the following set of schedules, which range in complexity from 1 to N parameters:

- Pure Bundling. Consumers pay a fixed price b for access to all N articles.
- Linear Pricing. Consumers pay a fixed price p for each article purchased.
- Two Part Tariff. Consumers pay a subscription fee f , along with a fixed price p for each article.
- Mixed Bundling. Consumers have a choice between a per-article price p and a bundle price b.
- Block Pricing. Consumers pay a price p_1 for the first j articles, and a price p_2 for remaining articles.
- Nonlinear Pricing. Consumers pay a different price p_i for each article i.

The expected surplus for each schedule is presented in Table 1, as a function of the bundle size N and the population size n. In the limit of $N \to \infty$ and $n \to \infty$, an exact analysis of the schedules is possible [Brooks *et al.*, 1999]. These values, scaled by dividing by Nn, are given in column 3 of Table 1.

For finite size bundles and an infinite population size, the dependence of the profit upon the various parameters—the profit landscape—can be punctuated by discontinuous peaks or cliffs for small bundles, e.g. $N = 10$, but is smooth for sufficiently large bundles, e.g. $N = 100$. The ideal two-part tariff landscape for $N = 10$ and $N = 100$ are displayed in Figs. 1a and b. A combination of analysis and numerical analysis was used to derive the optimal profits in column 4, which are again scaled by dividing by Nn.

For finite size populations, the profit landscapes acquire numerous bumps, as illustrated in Fig. 1c. A cross-section of two realizations of the profit landscape for the $N = 100$ two-part tariff schedule, along with a cross section of the ideal $N = 100$ landscape, is shown in Fig. 2. Although the bumps decrease in height with population size, for a schedule such as two-part tariff for which the profit landscape has a broad central peak, for sufficiently small populations the bumps can be high enough to shift the peak location substantially from that of the ideal landscapes of columns 3 and 4. As a corollary, the landscape can shift substantially from one realization of the landscape to another (i.e. two different random selections of n consumers from the same distribution of w and k). The bumps can also be hazardous to algorithms that attempt to find the global maxima.

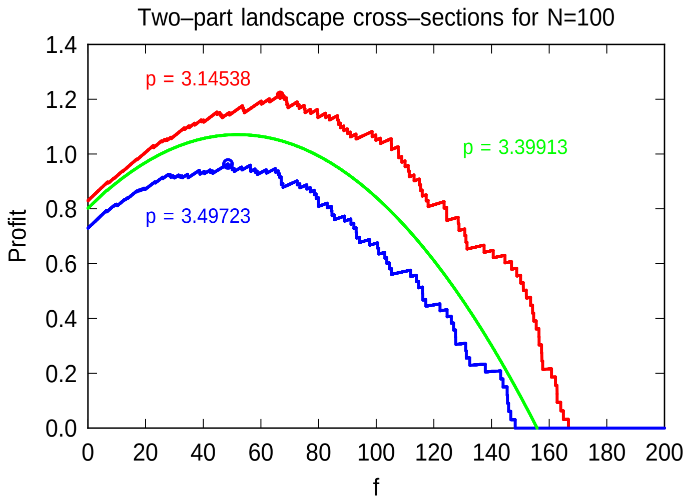

Figure 2: Cross sections of ideal $n \to \infty$, $N = 100$ landscape (green) and two realized $n = 100$, $N = 100$ landscapes (red and blue). Cross sections are for values of per-article parameter p fixed at the values indicated in the figure, which are chosen to be those at which the absolute peak in the landscape occurs.

Column 5 of Table 1 gives the peak profits for finite consumer populations, averaged over 10,000 landscapes, each of which was generated by choosing 100 random consumers from the given distribution. The peaks were determined by running the amoeba optimization algorithm 100 times on

Schedule	**Parameters**	$\Pi_{\textbf{opt}}(\infty, \infty)$	$\Pi_{\textbf{opt}}(100, \infty)$	$\Pi_{\textbf{opt}}(100, 100)$	$\Delta\Pi_{\textbf{opt}}(100, 100)$
Pure Bundling	b	0.8750	0.9010	0.9497 (0.0878)	0.0666
Linear Pricing	p	0.8750	0.8999	0.9044 (0.0515)	0.0042
Two Part Tariff	f, p	1.0370	1.0710	1.1043 (0.0777)	0.0464

Table 1: This table presents characteristics and optimal profits for a series of price schedules, when applied to the following population: $w = 10$, $N = 100$, and k drawn from $U[0, 0.7]$. Column 3 gives the optimal profit per good sold that can be earned by each schedule, with a bundle size $N \to \infty$ and a population size $n \to \infty$. Column 4 gives the corresponding results for a bundle size $N = 100$ and a population size $n \to \infty$. Column 5 gives the corresponding results for a bundle size $N = 100$ and a population size $n = 100$. The standard deviation across different realizations of the landscape are given in parentheses. Column 6 indicates the average additional profit that can be obtained by re-optimizing a landscape after redrawing a new population, assuming bundle size $N = 100$ and population size $n = 100$.

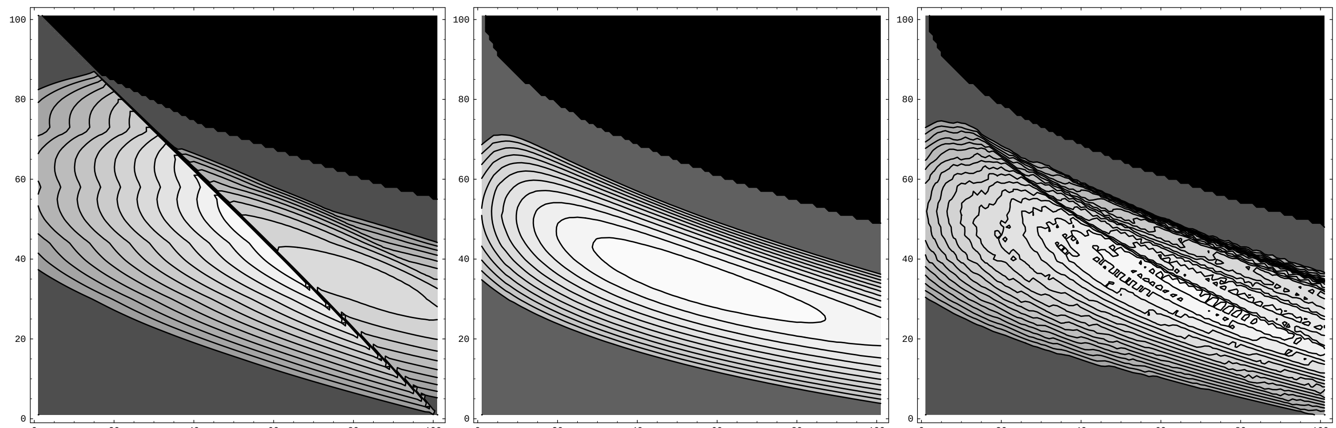

Figure 1: a) Ideal profit landscape with $N = 10$ and two-part tariff pricing. Contour plot of normalized profit, with contour lines starting at 0, 1.00, 1.05, 1.10, etc., ranging up to 1.75. Peak profit of 1.6500080 is attained at $(f = 5.345225, p = 4.654775)$. b) Ideal profit landscape with $N = 100$ and two-part tariff pricing. Contour plot of normalized profit, with contour lines starting at 0, 0.80, 0.825, etc., ranging up to 1.25. Peak profit of 1.07104834 is attained at $(f = 5.15788, p = 3.399134)$. c) Realized profit landscape with $N = 100$ and two-part tariff pricing, with 100 randomly generated consumers. Contour plot, with contour lines as in b).

each landscape and choosing the highest peak found over the 100 runs. Previous experimentation had established that this was much more than sufficient to find the global peak for the two-part tariff landscape. Taking the two-part tariff case as an example, the average peak profit of 1.1043 is measurably higher than the peak of the ideal landscape, which is 1.0710. This comes about because a realized landscape is bound to have portions which are both lower and higher than the ideal, and we are selecting the highest point on each landscape. The numbers in parentheses represent the standard deviation across different realizations of the landscape, which was much larger for pure bundling and two-part tariff than for linear pricing.

Since the landscapes can shift even when there is no underlying drift in the distribution according to which consumers are chosen, it is of interest to characterize how these shifts affect the overall profit. To do this, for each successive landscape we evaluated both the peak profit for that landscape and the profit that would have been attained at the price parameters for the previous landscape. The resultant difference is reported in column 6 of Table 1. This difference provides a sense of the amount of gain there is to be had from trying to track a changing but statistically stationary consumer population, as opposed to simply riding out the statistical fluctuations. The actual amount of gain may be less if the optimization algorithm fails to find the new global optimum in the landscape whenever the landscape shifts, either by getting stuck at a local optimum or having insufficient time before the landscape changes again.

For smaller consumer populations, we have verified that the amount of potential gain is larger. This is to be expected, as the relative size of the bumps in the landscape is larger.

4 Amoeba

Amoeba is a variant of the *simplex algorithm* [Nelder and Mead, 1965] for nonlinear unconstrained optimization problems. (This simplex algorithm should not be confused with the simplex algorithm for linear programming.) The amoeba algorithm maintains at each iteration a nondegenerate simplex, which is a geometric figure in n dimensions of nonzero volume that is the convex hull of $n+1$ vertices, $x_0, x_1, ..., x_n$, and their respective function values. Amoeba alternately reflects, contracts, and expands the simplex in an attempt to locate an optimum in a function's landscape. In each iteration, a new point (in this case, prices) is chosen and the corresponding value (profit) is determined.[2]

Particularly relevant for our work is the fact that amoeba does not make use of any gradient information when finding optima. This makes it a particularly suitable choice for this problem, as the profit landscape the producer must search is dotted with peaks and discontinuities (e.g., profit landscape in Figure 1c.)

Amoeba was designed for searching a fixed, deterministic landscape. In order to contend with a changing profit landscape, we have made some straightforward extensions. First, amoeba is essentially a "memoryless" algorithm - the simplex only contains the last $n+1$ observed points. In order to detect a change, a producer using amoeba must be able to observe that a previously visited point now yields a different profit. This is done by keeping a hash table of points and profits.

When amoeba determines that the landscape has changed, it discards the n lowest-profit points in the simplex and generates new, random points. All $n+1$ points in the simplex are then revisited. This introduces a bias that the previous 'best' point will be in the vicinity of the new optima; it will tend to pull amoeba in its direction. Empirically, this technique seems to perform well, thanks to amoeba's ability to converge quickly as long as the initial simplex is large enough.

5 Learning experiments

In order to determine which price schedules yield the greatest aggregate profit, we conducted a set of experiments. In each experiment, the consumer population was periodically replaced, creating a shock in the landscape.

As discussed in section 2, we consider two separate models of consumer population change. In the first, each consumer's k is drawn from a fixed distribution $[0, \overline{k}]$. At periodic intervals, the consumer population is removed and replaced with a new set of consumers drawn from the same distribution. The nonstationarity arises from differences in the particular realizations of each population. The second model incorporates change in the underlying population. Again, consumers are drawn from a distribution $[0, \overline{k}]$. At periodic intervals, they are removed and replaced. In this case, $\overline{k}$ is altered. A separate random variable $\alpha = U[\underline{\alpha}, \overline{\alpha}]$ is generated and added to the previous $\overline{k}$, producing a random walk by the bounds of the distribution. In these experiments, $\alpha = 0.3$. The new consumer population is then drawn from this new distribution.

In these experiments, w_0 was fixed at 10. Experiments were conducted for $N = 10$ and $N = 100$ goods. For each population change model, we conducted experiments where the shock occurred every 10, 50, and 250 iterations. This produces cases in which the population changes more quickly, at approximately the same rate as, and more slowly than amoeba is able to converge.

Figure 3 shows cumulative profit per iteration when the population changes every 10 iterations. The figures show $N = 10$ and $N = 100$ resampling from fixed and drifting populations. We see that in the very early stages, nonlinear pricing performs well, but after about 20 iterations it is overtaken by the two-parameter schedules, two-part tariff and mixed bundling. Even though these schedules start at worse points in the profit landscape, these schedules are more easily learned, and show the quickly dominate nonlinear pricing, despite the latter's higher steady-state profit.

Figure 4 shows cumulative profit when the population changes every 50 iterations. We see that the two-parameter schedules overtake nonlinear pricing even more quickly as the amount of nonstationarity is reduced.

Figure 5 shows cumulative profit when the population changes every 250 iterations, or once within the course of the experiment. In this experiment, bundle size makes a difference; by the end of the experiment, when $N = 10$, nonlinear pricing has had long enough to reach an optima and overtake the two-parameter schedules. When $N = 100$, it still lags

[2] For a detailed survey on amoeba refer to Walters [1991].

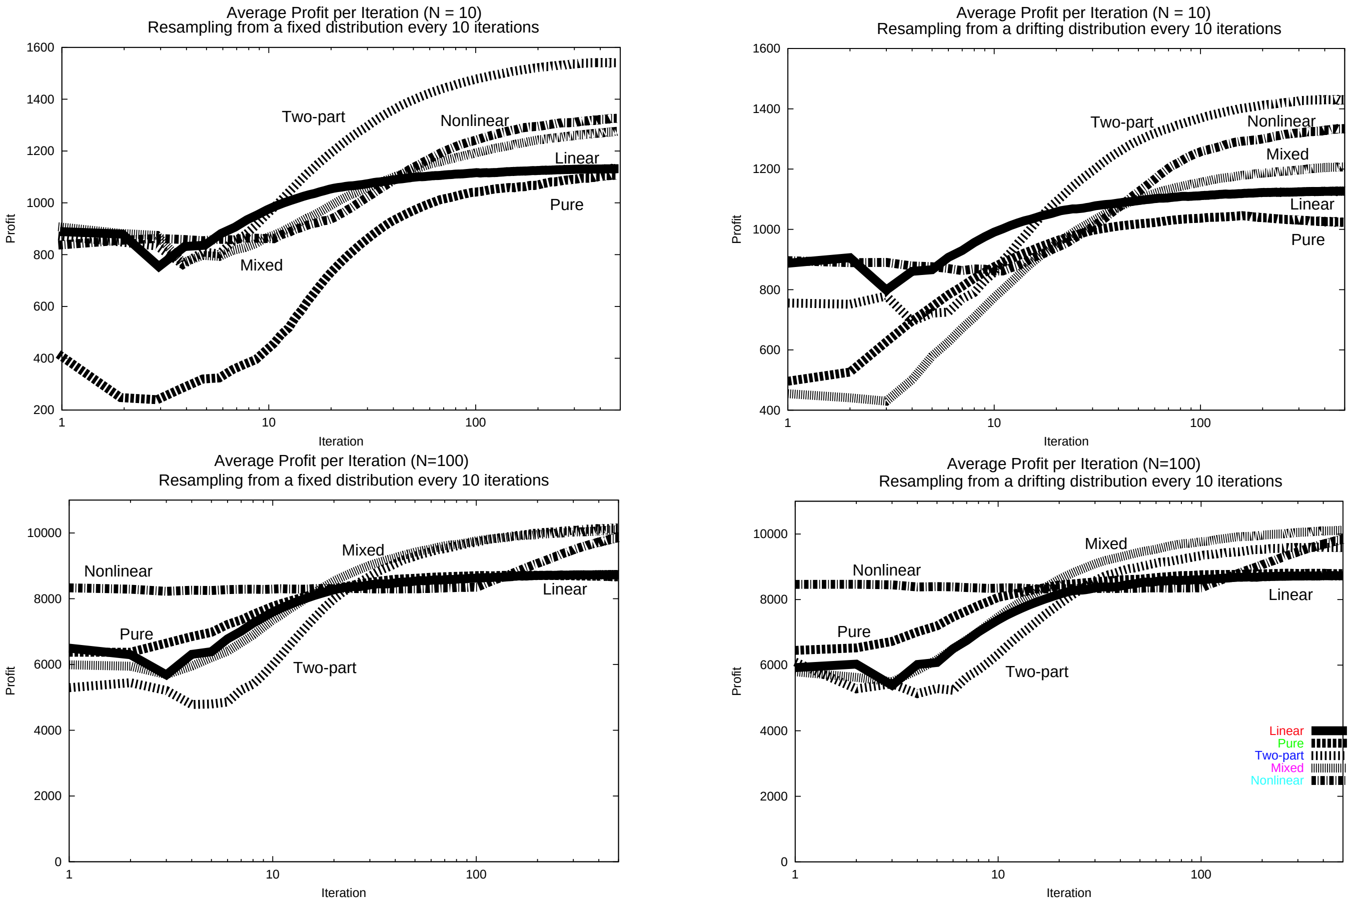

Figure 3: Cumulative profit per iteration when the population changes every 10 iterations

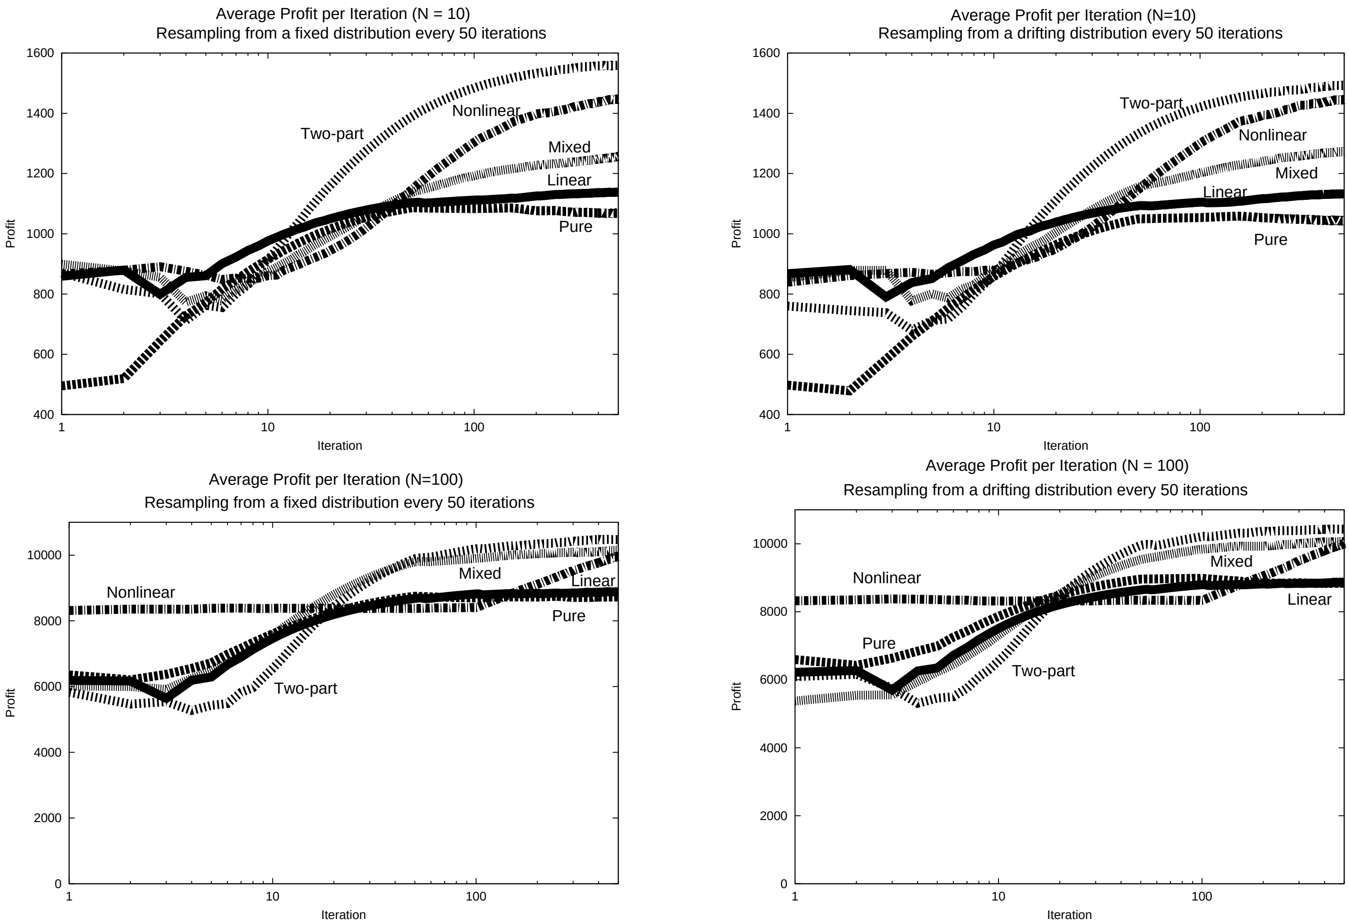

Figure 4: Cumulative profit per iteration when the population changes every 50 iterations

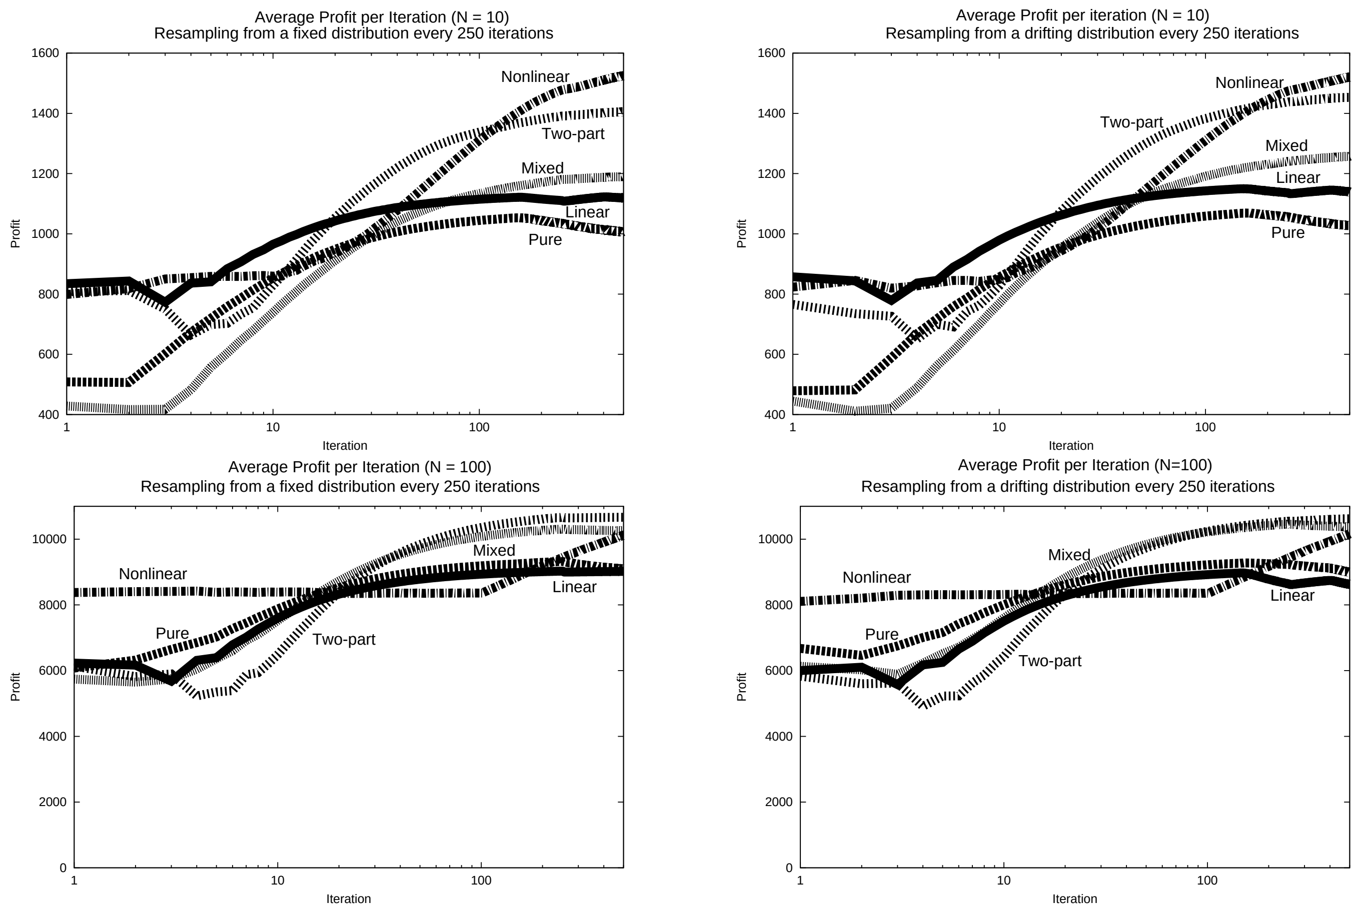

Figure 5: Cumulative profit per iteration when the population changes every 250 iterations

behind. (Recall that nonlinear pricing schedules have N dimensions.) Also worth noting is that after 250 iterations, the two-parameter schedules have converged on a solution; once the population changes, they are able to quickly shift to a new optima. In the case of two-part tariff, this is due to the large, flat mesa at the top of its hill. As discussed in section 3, realizations of two-part tariff tend to have local optima that lie within a plateau at the peak of the landscape. This makes it easy for a producer using two-part tariff to quickly locate a new optimum once the population changes. In addition, all solutions on this plateau tend to perform well, making the schedule resistant to consumer nonstationarity as modeled.

Another interesting detail is the different performance of mixed bundling as the bundle size increases. Recall that mixed bundling offers consumers a choice of a per-article price or a bundle price. This can be a disadvantage, as learning about one parameter does not appear to provide much assistance in learning the other. The two-part tariff, on the other hand, has more closely related parameters. Empirically, a producer using mixed bundling often winds up learning one parameter and abandoning the other, making it effectively a one-parameter schedule. This seems to happen more when $N = 10$. Perhaps the smaller nonzero region makes it more difficult to isolate and tune the effects of each parameter.

6 Conclusion

In this paper, we have explored some simple models of consumer nonstationarity and compared the performance of firms using different price schedules in this nonstationary environment. We have have seen that two-parameter schedules, and the two-part tariff in particular are attractive schedules as they are both robust to changing consumer population and learned rather quickly. This provides further evidence that when a producer is interested in maximixzing its aggregate profit in an uncertain environment, it must consider not only the steady-state profits yielded by a particular schedule, but also the ease with which this schedule can be learned as this greatly affects the profits earned in the interim periods.

This work has only scratched the surface of learning in a nonstationary environment. We plan to extend it in several ways. First, we would like to explore more complex models of nonstationarity, including continuous change, as well as evolutionary models in which consumer replacement is tied to their satisfaction. This also gives us an avenue to introduce learning on the part of the consumers.

Competition is an essential next step. Our previous work [Brooks *et al.*, 2000] examined a scenario in which producers learn to discover niches in an unknown but static population. Extending this to examine niche formation and discovery in a dynamic world is a promising extension.

Finally, we would like to extend the current model of consumer preferences. While the Chuang-Sirbu model is useful for its tractability, it does not have the richness one might like. For example, it doesn't consider substitutes or complements, and it assumes that the producer can show a set of N goods to the consumer, who then picks the ones she likes. If, instead, the producer must select these goods based on a prediction of the consumer tastes, the problem becomes much more complex.

Acknowledgments

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

References

[Aghion *et al.*, 1991] Philippe Aghion, Patrick Bolton, Christopher Harris, and Bruno Jullien. Optimal learning by experimentation. *Rev. Econ. Stud.*, 58(4):621–654, 1991.

[Bakos and Brynjolfsson, 1999] Y. Bakos and E. Brynjolfsson. Bundling information goods: Pricing, profits and efficiency. In D. Hurley, B. Kahin, and H. Varian, editors, *The Economics of Digital Information Goods*. MIT Press, Cambridge, Massachusetts, 1999.

[Brooks *et al.*, 1999] C. H. Brooks, S. Fay, R. Das, J. K. MacKie-Mason, J. O. Kephart, and E. 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.

[Brooks *et al.*, 2000] C. H. Brooks, E. H. Durfee, and R. Das. Price wars and niche discovery in an information economy. In *Proceedings of ACM Conference on Electronic Commerce (EC-00)*, Minneapolis, MN, October 2000.

[Chuang and Sirbu, 1999] J. C. Chuang and M. A. Sirbu. Network delivery of information goods: Optimal pricing of articles and subscriptions. In D. Hurley, B. Kahin, and H. Varian, editors, *The Economics of Digital Information Goods*. MIT Press, Cambridge, Massachusetts, 1999.

[Duda and Hart, 1973] Richard O. Duda and Peter E. Hart. *Pattern Classification and Scene Analysis*. John Wiley & Sons, 1973.

[Fe'ldbaum, 1965] A. A. Fe'ldbaum. *Optimal Control Systems*. Academic Press, New York, 1965.

[Grossman *et al.*, 1977] Sanford J. Grossman, Richard E. Kihlstrom, and Leonard J. Mirman. A Bayesian approach to the production of information and learning by doing. *Rev. Econ. Stud.*, 44(3):533–547, 1977.

[Hu and Wellman, 1998] Junling Hu and Michael P. Wellman. Online learning about other agents in dynamic multiagent systems. In *Proceedings of the Second International Conference on Autonomous Agents (Agents '98)*, 1998.

[Keller and Rady, 1999] Godfrey Keller and Sven Rady. Optimal experimentation in a changing environment. *Rev. Econ. Stud.*, 66:475–507, 1999.

[Nelder and Mead, 1965] J. A. Nelder and R. Mead. A simplex method for function minimization. *Computer Journal*, 7:308–313, 1965.

[Sutton and Barto, 1998] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1998.

[Thrun and Møller, 1992] Sebastian B. Thrun and Knut Møller. Active exploration in dynamic environments. In J.E. Moody, S. J. Hanson, and R. P. Lippmann, editors, *Advances in Neural Information Processing Systems 4*. Morgan Kaufmann, San Mateo, CA, 1992.

[Vidal and Durfee, 1998] Jose M. Vidal and Edmund H. Durfee. The moving target function problem in multi-agent learning. In *Proceedings of the Third Annual Conference on Multi-Agent Systems*, Paris, France, 1998.

[Walters *et al.*, 1991] F. H. Walters, L. R. Parker, S. L. Morgan, and S. N. Deming. *Sequential Simplex Optimization*. CRC Press, Boca Raton, Florida, 1991.

[Wilson, 1993] Robert B. Wilson. *Nonlinear Pricing*. Oxford University Press, 1993.