The Impact of Communication Costs and Limitations on Price Wars in an Information Economy

Jianhui Wu and Edmund H. Durfee
EECS Department, University of Michigan, Ann Arbor, MI 48109 USA
jianhuiw@umich.edu, durfee@umich.edu

Abstract

Price wars—the iterative undercutting of prices to the marginal cost by competitors—have frequently emerged in models of economic systems populated by computational agents. In this paper, we explore the prevalence and severity of price wars in models of multiagent ecommerce systems that include costs and limitations on interagent communication. The empirical results we describe in this paper indicate that, for a stationary consumer population, limiting the rate of penetration of price information can reduce the severity of price wars, and that charging producer agents for communication can in fact curtail price-undercutting before prices (and profits) bottom out. Furthermore, prices (and profits) do not bottom out for non-stationary consumer populations, where in fact cyclic price wars can arise.

1. Introduction

Using economic knowledge in decentralized resource allocation problems has been capturing increasing attention in recent work [3, 7, 11]. In domains like information economies, where the cost to replicate (information) goods (such as documents) is negligible, intense competition among myopic resource providers often occurs, triggering each competitor to try to cut its price below prices of others [1, 5, 14, 15, 16]. This is analogous to "price wars" in economic markets. A "price war" represents a period of intense competition in which competitors cut retail prices to gain business. If a producer can capture the entire market by slightly undercutting its competitors' prices, then the Nash equilibrium strategy is for each producer to sell its goods at marginal cost which, when the cost to replicate goods is negligible, approaches zero [5]. That is, when engaging in price wars, producers iteratively lower their prices until prices reach the marginal cost, or only one producer remains in the market.

Previous research [1, 5, 16] has studied the decision process of how to establish prices in a market having multiple producers. Based upon the assumption that all consumers can immediately know price shifts without any cost, the work [1, 5] has concluded that price wars are often inevitable among myopic producers who attempt to maximize immediate profits, if these producers cannot differentiate themselves from each other in ways other than through price.

In practice, however, iterative price-undercutting to the marginal cost seems to happen rarely. One reason for this is that producers often attempt to maximize their long-term profits rather than immediate profits. This is similar to an iterated prisoners' dilemma problem; a producer might encourage others to cooperate by adopting a tit-for-tat or a grim trigger strategy. In fact, producers could implicitly or explicitly collude to keep prices higher.

We hypothesize that another reason for the infrequency of price wars could be that, realistically, propagating awareness in a distributed system is generally neither instantaneous nor free. In this paper, we explore this hypothesis by factoring the costs and limitations of communication into a model of an information economy having myopic producers to reveal the effects on price wars. In the following discussion, we say that price wars become less severe if the rate of price-undercutting is slowed down (but prices would still eventually approach the marginal cost) and price wars are eliminated if, despite competition, producers charge prices above the marginal cost (and thus earn positive profit) indefinitely.

To simplify our investigation, we use a duopoly market, in which two producers sell identical information goods, as our model. Section 2 describes this model, and defines three types of agents in an information economy—consumer, advertiser, and producer—and their interactions. Section 3 presents analytical and simulation results showing whether price wars can be less severe or can even be eliminated when limitations in communication cause incomplete advertising penetration. We examine whether producers would in fact choose limited advertising penetra-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
AAMAS'04, July 19-23, 2004, New York, New York, USA.
Copyright 2004 ACM 1-58113-864-4/04/0007...$5.00

tion among themselves to avoid costly price wars. Of course, a producer that could have higher advertising penetration than its competitor could have an advantage; in Section 4 we study this, but assume that a higher advertising penetration will cost a producer more. Section 5 explores the impact of communication costs and limitations on price-setting decisions by producers when consumers are transient. Finally, Section 6 summarizes this work and outlines our future directions.

2. System Model and Agent Behaviors

Our model consists of two producers named *producer*$_1$ and *producer*$_2$, an advertiser, and a consumer population. In order to familiarize readers with some fundamental economic knowledge used in the study, this section starts with an introduction to the behaviors of the three types of agents modeled and then describes interactions between those agents.

2.1 Consumer

This work uses a simple consumer model to characterize the preferences of a consumer. In this model, a consumer k values the information good being sold at a value w_k. Each consumer attempts to maximize its surplus by deciding to buy from *producer*$_1$ or *producer*$_2$ (or neither) based upon their advertised pricing information, where surplus is w_k less the amount paid. Note that a producer might not be able to afford to pay the advertiser for sending its pricing information to consumers every iteration, and so the pricing information known by some consumers is perhaps stale. In order to distinguish between the price a producer actually charges and the pricing information most recently received by a consumer, in the following discussion, we use $price_k^p$ to represent the pricing information currently known by a consumer k about *producer*$_p$, and $actualprice^p$ is the actual price *producer*$_p$ currently charges for its good.

2.2 Advertiser

In our model, a producer cannot directly send pricing information to the consumer population; it needs to pay the advertiser to do this. The probability that a consumer receives the advertised pricing information is determined by the advertising fee paid to the advertiser by the producer. Viewed from the producer's perspective, this fee paid is the advertising cost. In other words, the probability of a consumer receiving the price announcement, denoted as *adpenetration*, can be represented as a function of advertising cost (*adcost*) to the producer:

$$adpenetration = f(adcost) \qquad (2\text{-}1)$$

The type of function that f is is determined by the method of advertising. In this paper, we use linear functions, although we have also studied exponential functions.

2.3 Producer

In our work, all producers use the same learning algorithm, and it is assumed that each producer has no prior knowledge about consumer preferences when it initially enters the market. Each producer maintains a model of the world in the form of a radial basis function (RBF) network [8].

At each iteration, based upon the myopic assumption that its competitor will maintain the same decisions as in the previous iteration, each producer attempts to maximize its immediate net profit by reconsidering its pricing and advertising decisions. In detail, each producer p is learning a decision function $\delta_i^p(w): W \rightarrow A$ that predicts the right actions to take in the world state w at iteration i. Recall that pricing information of some consumers is perhaps stale. We use S_i, which is a set of prices, to represent the state of knowledge about pricing information for the whole consumer population. That is, for a market having a consumer population of size n and two producers, S_i can be represented as {($price_1^1$, $price_1^2$), ($price_2^1$, $price_2^2$), …, ($price_n^1$, $price_n^2$)}. The subscript i in S_i indicates that it is the pricing information at the beginning of the i^{th} iteration. If S_i is fully observable, then the decision function δ_i^p can be represented as

$$\delta_i^p = argmax_{a_i^p}\ rbf\_net(S_i, a_i^p, a_i^{\bar{p}}) \qquad (2\text{-}2)$$

where $\bar{p}$ represents the competitor of producer p. a_i^p is the action taken by producer p at iteration i, which includes price-setting and advertising decisions. Each producer p is learning consumer preferences by training a radial basis function network (rbf_net), which maps inputs ($S_i \times a_i^p \times a_i^{\bar{p}}$) to net profit ($netprofit_i^p$).

Note that in a real-world market, it is usually impossible for a producer to accurately track S_i because of the large size of the consumer population. Viewed from the perspective of producer p, S_i can be approximately represented as a function of p's price and its market share in the previous iteration. As a result, the decision function δ_i^p can be rewritten as:

$$\delta_i^p = argmax_{a_i^p}\ rbf\_net(ms_{i-1}^p, actualprice_{i-1}^p, a_i^p, a_i^{\bar{p}}) \qquad (2-3)$$

where ms_{i-1}^p and $actualprice_{i-1}^p$ are the market share and the price, respectively, of producer p in iteration *i-1*.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
AAMAS'04, July 19-23, 2004, New York, New York, USA.
Copyright 2004 ACM 1-58113-864-4/04/0007...$5.00

As we have stated, a producer lacks knowledge of consumer preferences initially, and so the producer has to balance the tradeoff between exploration and exploitation. That is, it needs to decide whether it should use the best strategy so far to gain a good profit in the current iteration or explore further to search for a potentially better strategy for the future. A parameter θ , often called "*temperature",* is used in our work. In detail, we apply a Boltzmann distribution [19], which utilizes current utility estimates to make action selections. The probability for an action, which includes price-setting and advertising decisions, to be selected is a nonlinear function of its estimated net profit:

$$probability(a) = \frac{e^{f(a)\times\theta^{-1}}}{\sum_{a'} e^{f(a')\times\theta^{-1}}} \qquad (2-4)$$

where the function *f(a)* represents the normalized estimated net profit for the given action *a,* which includes price-setting and advertising decisions. (The estimated net profit is computed by using the trained RBF network.) When temperature θ is large, then even actions that are expected to incur low profits have some probability of being selected. Contrarily, only the best action, which produces the highest expected profit based upon the current knowledge, can be selected when temperature θ is close to zero. In this work, temperature θ is initially set to a large positive value and then gradually reduced if the estimated profits by using the RBF network are close to the actual profits. Otherwise, temperature θ is increased again in order to do more exploration for completing the RBF network.

It should be noted that we do not claim that the RBF network and Boltzmann exploration policy are optimal for this task, nor that the approximation of S_i is the best possible. In this work we are interested in the impact of communication limitations and costs on price-setting and advertising decisions made by a producer instead of the optimality of its learning algorithm. Although market dynamics might be different with different learning algorithms initially, producers will make similar decisions in the long term independently of which learning algorithm they use.

2.4 Interactions Between Agents

We now turn to describe interactions between the agents modeled, depicted in Figure 2-1. At the beginning of each iteration, each producer attempts to collect market information, including its competitor's pricing and advertising decisions and its own profit in the previous iteration, and then updates its RBF network. After that, according to the decision function δ_i^p (equation 2-3) and Boltzmann exploration policy

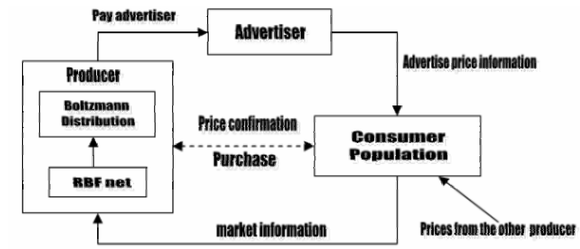

Figure 2-1: Interactions Between Agents

(equation 2-4) along with an estimate of $a_i^{\bar{p}}$ [1], each producer reconsiders its price and decides to do advertising or not in order to maximize its net profit.

If a producer decides to do advertising, then it pays the advertiser to send its new pricing information to consumers. Once advertising is done by both producers, each consumer *k* chooses to buy from the producer charging a lower price (or neither if both prices are higher than the consumer's valuation of the information good) based on its local pricing information $\{ price_k^1, price_k^2 \}$.

2.5 Experimental Setup

In this work, each consumer's valuation of the information good, w_k, is independently chosen at random with uniform probability from a given range. The results reported in this paper are based upon the range [20 80]. For this consumer population, the monopolist profit by charging price *x* is:

$$\begin{aligned} Profit(x) &= (x \times (\frac{80-x}{60}) + 0 \times (\frac{x-20}{60})) \times size \\ &= (-\frac{x^2}{60} + \frac{4}{3}x) \times size \qquad (2\text{-}5) \end{aligned}$$

where *size* is the size of the consumer population. In our simulations, the size of population is set to be 200. From equation 2-5, we can find that the optimal price is 40 and the monopolist's expected optimal profit in an iteration is 5340. In the following sections, we can see that producers usually receive much lower profits in a duopoly market because of competition.

3. Negligible Advertising Cost

Previous research [1, 5, 16] has illustrated that price wars are inevitable when advertising cost is negligible and advertising penetration is unlimited. In

[1] In this work, a producer myopically assumes that its competitor will maintain the same action as in the previous iteration, i.e, $a_i^{\bar{p}} = a_{i-1}^{\bar{p}}$

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
AAMAS'04, July 19-23, 2004, New York, New York, USA.
Copyright 2004 ACM 1-58113-864-4/04/0007...$5.00

this section, we attempt to answer whether price wars can become less severe or even be eliminated when advertising penetration is restricted (while advertising cost is still negligible). We begin with the assumption that producers are given the same advertising penetration when entering the market; a study of how a producer might select its own advertising penetration is postponed to Section 3.2.

3.1 Limited Advertising Penetration

To start our study of how advertising penetrations affect producer decision-making on price, Figure 3-1 shows a series of price decisions of $producer_1$ on two advertising penetrations. We can observe that, when there is no limitation in the advertising penetration (*adpenetration* =100%), a price war develops and so forces $producer_1$ to quickly lower its price to the marginal cost. However, when the advertising penetration is restricted to 7.5% for both producers, the price falls much more slowly. This is because the low advertising penetration (*adpenetration* =7.5%) makes it difficult for a producer to catch the attention of its competitor's consumers even when its price is lower than its competitor's price. In other words, it takes longer for new pricing information to be received by potential consumers, and so producers undercut each other much less frequently when the advertising penetration is low.

Figure 3-2 shows the aggregate profits over 1500 iterations for each producer for various advertising penetrations. It is assumed that initially consumers don't know at all about the existence of producers. For a very small advertising penetration, such as 0.01%, the aggregate profit is close to zero because it is almost impossible for a producer to capture the attention of any potential consumers given such a small advertising penetration. The aggregate profit rapidly increases as the advertising penetration in-

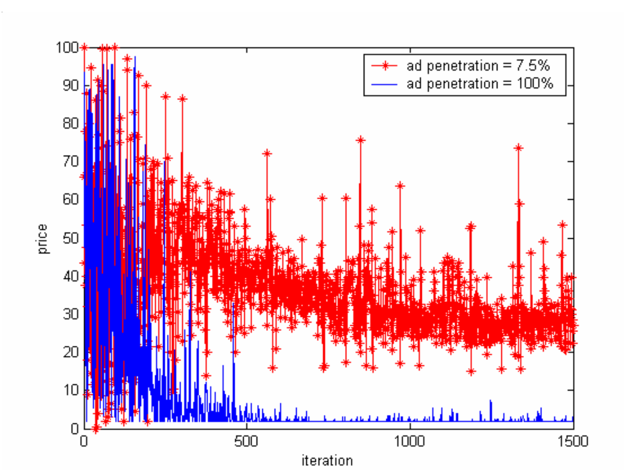

Figure 3-1: Effects of penetrations on prices

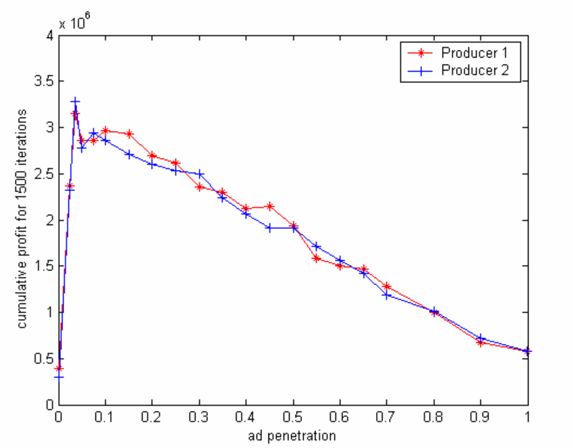

Figure 3-2: Effects of penetrations on aggregate profits

creases. It reaches its peak when the advertising penetration is near 4%. This is because, for an advertising penetration close to 4%, it is not too difficult for a producer to start to gain profits. On the other hand, small advertising penetrations make prices fall slowly and so producers can have good profits over a much longer period. As the advertising penetration is increased further, the aggregate profit starts to decrease. This is because the larger the advertising penetration is, the more easily the profit of a producer is affected by its competitor's price. As a result, more intense competition occurs and forces producers to lower their prices more frequently. When prices approach the marginal cost, the profits are close to zero and so the aggregate profits are naturally smaller.

Note that realistically, a producer cannot undercut its competitor by an infinitesimal amount, so that a producer might not lower its price because its expected loss due to a price decrement might exceed its expected gain from attracting a (small) fraction of its competitor's consumers. Our simulation results show that a very small advertising penetration makes it possible for producers to charge the same price (since they almost do not inference with each other when attracting the consumers without market information), and so the price might stabilize at a value significantly above the marginal cost.

In sum, when the advertising penetration is low, price wars become less severe. That is, for a long time, producers charge prices much higher than the marginal cost. Moreover, when advertising penetration is very small, price wars might even be eliminated.

3.2 Decisions on Advertising Penetration

In the previous discussion, it is assumed that the duopolist producers are given the same advertising

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
AAMAS'04, July 19-23, 2004, New York, New York, USA.
Copyright 2004 ACM 1-58113-864-4/04/0007...$5.00

penetration when entering the market. In a more realistic environment, it is more natural to think that each producer can individually choose its advertising penetration upon entry into the market.

Changing advertising penetration generally incurs some costs. In this study, we assume that a producer can decide whether to do advertising or not at each iteration of the market lifetime, but that advertising penetration cannot be changed once it is initially set by the producer.

For a producer, choosing a zero advertising penetration makes no sense and choosing a 100% advertising penetration could cause severe price wars. Figure 3-2 shows that when producers are restricted to choose the same advertising penetration, the optimal advertising penetration for each producer is near 4%. However, realistically, choosing a penetration is generally more complex since each producer makes its decision separately. Problems, like the well-known prisoner's dilemma, might arise. For example, Figure 3-3 shows a fraction of the payoff matrix, in which the utility represents the aggregate profit over 1500 iterations for each producer, for the experimental setup discussed in Section 2.5. Clearly, the strategy of choosing a 20% advertising penetration strongly dominates the strategy of choosing a 4% advertising penetration.

	$Producer_1$: ad penetration =4%	$Producer_1$: ad penetration =20%
$Producer_2$: ad penetration = 4%	$P1=3.2\times10^6$ $P2=3.2\times10^6$	$P1=3.8\times10^6$ $P2=1.9\times10^6$
$Producer_2$: ad penetration =20%	$P1=1.9\times10^6$ $P2=3.8\times10^6$	$P1=2.7\times10^6$ $P2=2.7\times10^6$

Figure 3-3: A fraction of the payoff matrix

A producer can choose any advertising penetration within the range of [0% 100%]; thus, it is not easy to find Nash equilibrium strategies by theoretical analysis. In the following discussion, we run simulations to illustrate the decision process for the optimal advertising penetration for each producer.

As described above, each producer decides upon its advertising penetration when entering the market. It has the information of its aggregate profit over all iterations when exiting the market, so a producer can analyze this aggregate profit feedback signal and improve its decision upon its advertising penetration when it enters the market the next time. For the purpose of simplicity, we assume that each producer has perfect information about the aggregate profit for any combination of penetrations, and the capability of collecting information about the advertising strategy of its competitor.[2] We run the simulations with random initial advertising penetrations and various length simulations. One of the results is shown in Figure 3-4. These results illustrate that, even given a random initial penetration, each producer eventually chooses an advertising penetration that falls in a moderate range of 5% to 35%. For such penetrations, it takes a long time for prices to be reduced to the marginal cost.

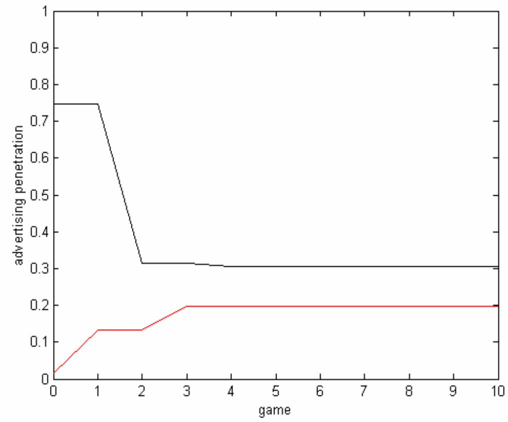

Figure 3-4: The process of decision making on the advertising penetration

It should be noted that the optimal penetration is related to the length of the simulation. The simulation length represents how long the information good offered by producers is desired by consumers. If consumers lose desire for information goods quickly, then a producer should prefer a higher penetration because it wants to enter the market quickly and attract consumers from its competitor quickly. On the other hand, if consumer preferences can be stable for a long time, then producers should choose low penetrations since neither of them wants the occurrence of intense competition, which can quickly drive prices to the marginal cost.

In general, we conclude that, when advertising cost is negligible, producers would choose advertising penetrations in a moderate range in most cases (except when consumers lose desire for information goods quickly). Such penetrations, although still causing prices to approach the marginal cost, force prices to fall much more slowly than an advertising penetration of 100%. That is, price wars become less severe when limitations of communication are presented and even voluntarily self-imposed.

[2] If a producer does not have complete information, it needs to do exploration to learn this.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
AAMAS'04, July 19-23, 2004, New York, New York, USA.
Copyright 2004 ACM 1-58113-864-4/04/0007...$5.00

4. Non-negligible Advertising Cost

In Section 3, we concluded that price wars can be slowed down by restricting advertising penetration and might even be eliminated for a very small penetration. However, when each producer can choose its own advertising penetration, they generally choose much larger advertising penetrations and so prices would eventually approach the marginal cost. In this section, we attempt to answer a question similar to that proposed in Section 3, but concentrate on the case where higher advertising penetration will cost a producer more. The structure of this section is similar to Section 3: we start with the assumption that producers have the same advertising penetrations established when they enter the market, and then we study how a producer might select its own advertising penetration.

4.1 Limited Advertising Penetration

To begin with, we assume a linear advertising function: *adpenetration=k*adcost*, and let k be 1/2000, which means that the advertising cost is 2000 for a 100% advertising penetration. Recall that the optimal profit in an iteration for our experimental setup is 5340 (and its corresponding optimal price is 40); that is, it is still possible for a producer to have a good profit, even when it spends 2000 on advertising each iteration. The results are shown in Figure 4-1.

We set k to 1/4000 and run the simulation again; the results are shown in Figure 4-2.

From Figures 4-1 and 4-2, we observe that price curves stop decreasing at some positive values rather than converging to the marginal (zero) cost, and furthermore, as advertising cost changes, the price curves end at different values. This is because the increment of profit by lowering the price diminishes as the price decreases. For example, suppose a producer can attract fifty more consumers by lowering its price, and so the producer can earn an extra profit of about 2000 when it reduces its price from 40 to 39.99, but it can only get an extra profit of about 1500 when the price is reduced from 30 to 29.99. Once the advertising expense counteracts the increment in profit, a producer will not be willing to lower its price again. More specifically, if a market state is stable, it must satisfy the following inequality for any positive δ_i for each producer:

$$(p_i - \delta_i) \times ((PC(p_i - \delta_i) - c_i) \times adpenetration_i + c_i) - p_i \times c_i < adcost_i \quad (4-1)$$

where p_i is the price that $producer_i$ charges for an information good, c_i is the number of consumers buying goods from $producer_i$, p_i-δ_i is a new price, $PC(p_i$-$\delta_i)$ represents the expected number of consumers who are willing to pay at least p_i-δ_i for a good, and $(p_i - \delta_i) \times ((PC(p_i - \delta_i) - c_i) \times adpenetration_i + c_i)$ is the expected profit for a given price p_i-δ_i after doing advertising with $adpenetration_i$. If the expected profit for any price lower than the old price is less than the old profit plus advertising cost, then a producer has no incentive to reduce its price. If neither of the producers is willing to lower its price, then a state is stable.

From the discussion so far, we can conclude that when advertising cost is not negligible (and also not too high since too high advertising cost will prevent producers from entering the market), although producers may undercut each other in cases where prices are high, producers would not undercut each other after prices drop down to some values.

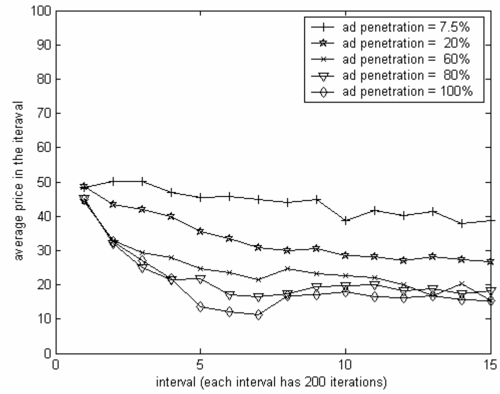

Figure 4-1: Average prices in each interval for various penetrations when *k*=1/2000

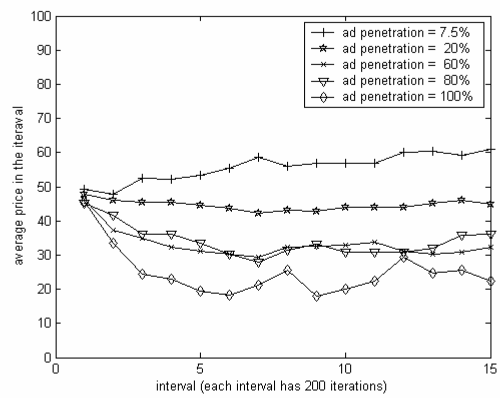

Figure 4-2: Average prices in each interval for various penetrations when *k*=1/4000

4.2 Decisions on Advertising Penetration

Now we come back to the question: when a producer has freedom to choose its own advertising penetration, what is its choice? We use a similar method as in Section 3 to answer this.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
AAMAS'04, July 19-23, 2004, New York, New York, USA.
Copyright 2004 ACM 1-58113-864-4/04/0007...$5.00

We run simulations with random initial advertising penetrations to illustrate the process of decision making for each producer. One of the results is shown in Figure 4-3. Analyzing the results over the various experimental setups, we notice that prisoners' dilemma problems generally arise. However, the price can usually be stabilized near the best stable price that produces the highest profit among all possible stable prices. The underlying reason is that, after reaching the stable price, a producer can stay in the market without competition. The better the price it can charge, the more profit it can earn.

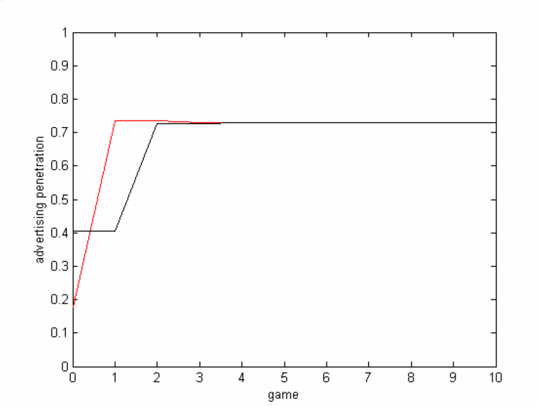

Figure 4-3: The process of decision making on the advertising penetration for a linear ad cost function (k=1/4000)

In summary, when advertising cost is not negligible (and also not too high), price wars can be eliminated. Moreover, prices would stabilize at the values significantly above the marginal cost when advertising is expensive.

5. Non-stationary Consumer Population

In the previous sections, we assume that the consumer population is stationary. We now study the impact of communication costs and limitations on price wars in a more general market in which each consumer has a probability, denoted as consumer change rate α, of leaving the market at each iteration, and being replaced by a new consumer (who does not have knowledge about the market).

Figure 5-1 shows the average price in each interval (100 iterations) for various advertising penetrations on the experimental setup where consumer change rate $\alpha = 5\%$ and advertising cost is negligible. We can see the occurrence of cyclic price wars in a wide range of advertising penetrations. That is, as prices approach the marginal cost, they suddenly jump up and a new cycle of price wars initiates.

It should be noted that, when advertising penetration is unlimited and advertising is free of charge, no cyclic price wars happen. The underlying reason is that, when advertising penetration is close to one, the producer charging a high (optimal) price can only attract a negligible number of the new consumers since most of the consumers hear about the lowest price immediately upon entry into the market. We also notice that, although 1% advertising penetration leads to prices being set close to the optimal price (and thus would appear to be a good choice), producers only receive 63% of the profit they would have earned by using 60% advertising penetration in this simulation (where the simulation length is set to 1500 iterations). A small advertising penetration means that producers can only serve a small fraction of the consumer population since it is not easy to propagate pricing information to the newly-arrived consumers.

Cyclic price wars can also be seen when advertising cost is not negligible. Figure 5-2 shows the simulation result for a consumer change rate $\alpha = 5\%$ and a

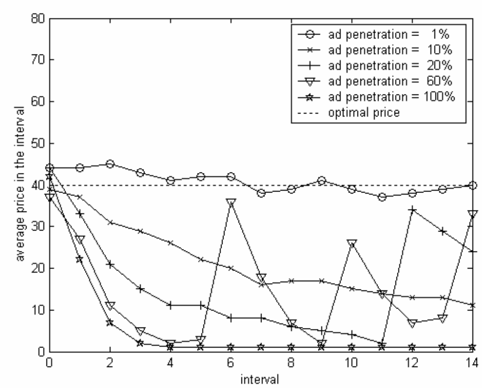

Figure 5-1 Average price in each interval (100 iterations) for various penetrations when con $\alpha = 5\%$ and advertising cost is negligible

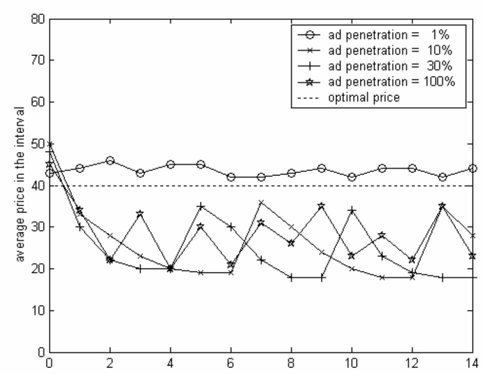

Figure 5-2 Average price in each interval (100 iterations) for various when $\alpha = 5\%$ and k=1/2670 (linear cost function)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
AAMAS'04, July 19-23, 2004, New York, New York, USA.
Copyright 2004 ACM 1-58113-864-4/04/0007...$5.00

linear advertising cost function with k = 1/2670 (i.e. advertising with 100% penetration would cost 50% of the monopolist profit). We observe that prices jump up at a value far above the marginal cost when advertising is expensive. Moreover, cyclic price wars occur even for an unlimited advertising penetration. When advertising has a cost, producers would stop price-undercutting and advertising once prices are below some values. As time passes, consumers are gradually replaced; a producer thus resets its price to the optimal price and does advertising again to attract new consumers – a new cycle of price wars starts.

6. Conclusions and Future Work

In this work, we use a duopoly market for information goods as a model to investigate our hypothesis that communication costs and limitations can give one possible explanation for why we seldom see producers selling goods at prices close to the marginal cost (even when producers are myopic). We concluded that, when advertising cost is negligible, each producer would choose its advertising penetration in a moderate range in most cases (except when consumers lose desire for the goods quickly). For such penetrations, although price wars cannot be eliminated, it takes a relatively long time for prices to approach the marginal cost. When advertising cost is not negligible, price wars can be eliminated. Moreover, when advertising is expensive, although price-undercutting initially occurs, prices would stabilize at the values far above the marginal cost and producers can have good, although not optimal, profits. In a more general market having transient consumers, cyclic price wars usually occur, i.e. the price gradually decreases and then suddenly jumps up to the optimal price.

Our work suggests that judiciously constraining communication in an information economy can improve system performance. One direction of our future work is to extend our duopoly market to a market with a larger producer population. Moreover, advertising in our work simply means propagating new pricing information to the consumer population. In the future, we are interested in exploring the effect of advertising in more general ways, such as as a means to change a consumer's expected valuation of a product.

Acknowledgements

This work was supported, in part, by NSF under ITR grant IIS-0112669. The authors thank three anonymous reviewers for their helpful suggestions and comments toward improving this paper.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
AAMAS'04, July 19-23, 2004, New York, New York, USA.
Copyright 2004 ACM 1-58113-864-4/04/0007...$5.00

References

[1] A. R. Greenwald, J. O. Kephart and G. J. Tesauro. "Strategic Pricebot Dynamics." In *Proc. ACM Conf. on Electronic Commerce*, 1999.

[2] A. Schaerf, Y. Shoham and M. Tennenholtz. "Adaptive Load Balancing: A Study in Multi-agent Learning." *Journal of Artificial Intelligence Research*, (2): 475–500, 1995.

[3] A. A. Lazar and N. Semret. "Auctions for Network Resource Sharing." CTR Tech. Report. Columbia University New York, February 1997.

[4] B. Arthur. "Inductive Reasoning and Bounded Rationality." *AEA Papers and Proceedings*, 84(2):406–411, 1994.

[5] C. H. Brooks, E. H. Durfee and R. Das. "Price Wars and Niche Discovery in an Information Economy." *Proc. of ACM Conf. on Electronic Commerce,* 2000.

[6] 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 a Complex Price Schedule." In *Proceedings of ACM Conf. on Electronic Commerce*, 31-40, 1999.

[7] D. Ferguson, Y. Yemimi and C. Nikolaou. "Microeconomic Algorithms for Load Balancing in Distributed Computer Systems." In *Int. Conf. on Distributed Computer Systems*, 1988.

[8] D. S. Broomhead and D. Lowe. "Multivariable Functional Interpolation and Adaptive Networks". *Complex Systems*, 2:321-355, 1988.

[9] E. A. Billard and J. C. Pasquale, "Dynamic Scope of Control in Decentralized Job Scheduling," *Proc. Int. Symp. on Autonomous Decentralized Systems*, Kawasaki, Japan, March 1993.

[10] G. Keller and S. Rady. "Optimal Experimentation in a Changing Environment." *Review of Economic Studies*, July 1999, 66 (3): 475-507.

[11] K. Kuwabara and T. Ishida. "Equilibratory Approach to Distributed Resource Allocation: Toward Coordinated Balancing." In Castelfranchi and Werner eds., *Artificial Sociality*, Lecture Notes in AI 830, Springer-Verlag, pp. 133-146, 1994.

[12] K. Bagwell "The Economics of Advertising." *The Int. Library of Critical Writings in Economics*, Vol. 136, Edward Elgar Pub. Ltd., 669 pp., 2001.

[13] K. Bagwell and G. Ramey, "Coordination Economies, Advertising and Search Behavior in Retail Markets," *AER* 84.3, 1994, 498-517.

[14] J. O. Kephart and A. R. Greenwald. "Shopbot Economics." In *Proceedings of the Third International Conference on Autonomous Agents (Agents '99)*, Seattle, WA, pp. 378–379. 1999.

[15] J. O. Kephart, J. E. Hanson and A. R. Greenwald. "Dynamic Pricing by Software Agents." *Computer Networks* 32(6): 731–752, 2000.

[16] J. M. DiMicco, A. Greenwald and P. Maes. "Learning Curve: A Simulation-Based Approach to Dynamic Pricing." *Electronic Commerce Research.*, 3:245–276, 2003.

[17] S. B.Thrun. "Efficient Exploration in Reinforcement Learning" Tech. Report CMU-CS-92-102, School of Computer Science, CMU, 1992.