Competitive Bundling of Categorized Information Goods

Jeffrey O. Kephart
Institute for Advanced Commerce
IBM Thomas J Watson Research Center
Yorktown Heights, New York 10598
kephart@us.ibm.com

Scott A. Fay
University of Florida
Warrington College of Business
Gainesville, FL 32611
faysa@notes.cba.ufl.edu

ABSTRACT

We introduce an information bundling model that addresses two important but relatively unstudied issues in real markets for information goods: automated customization of content based on categories, and competition among content providers. Using this model, we explore the strategies that sellers (or automated agents acting on their behalf) might use to set both price and bundle composition, and the market dynamics that might ensue from such strategy choices. The model incorporates different categories of information, explicitly accounts for finite production and consumption costs, and allows for possibly heterogeneous valuations by consumers. First, we determine the optimal bundle composition and price for a monopolist as a function of the seller's production costs and the consumers' preferences and consumption costs. For finite costs, finite-sized bundles are optimal. Then, we use game-theoretic analysis and simulation to explore the behavior of the market when there are multiple content providers. We find that, if consumer preferences are homogeneous, sellers choose to offer the same bundle that a monopolist would choose, but that competition forces sellers to offer the bundles at cost. For heterogeneous preferences, positive profits are possible, but there appears not to be a pure strategy Nash equilibrium. This is manifested as a never-ending cycle of prices and bundle choices when sellers employ a myopic best-response algorithm.

1. INTRODUCTION

The extremely low marginal cost of replicating and distributing information goods on the Internet has led to a resurgent interest in the study of product bundling. Most of the literature has focused on the question of whether a seller ought to sell items individually or as a fixed bundle, depending on the structure of consumer preferences, production costs, and a variety of other conditions. However, several practically important issues in electronic markets for information goods have received very little attention. One such issue is the development of automated (and hence low-priced) technologies that permit consumers to pick and choose items to compose their own customized bundles based on information categories. Another is that information providers such as online journals will compete with one another both on price and on the categorical composition of their bundles. It is of great practical interest to explore the strategies that sellers (or automated agents acting on their behalf) might use to set both price and bundle composition, and the market dynamics that ensue from such strategy choices.

This paper presents and analyzes a model in which multiple sellers compete to offer bundles of categorized information goods. It explicitly considers the consumers as (human or automated) agents that actively choose sets of bundles that will best satisfy their individual needs. First, in section 2, we review some of the most relevant bundling literature and discuss why it fails to address many of the issues that we feel are relevant for markets of bundled information goods. Then, in section 3, we introduce a novel information bundling model that incorporates different categories of information, explicitly accounts for finite production and consumption (or clutter) costs, and allows for possibly heterogeneous valuations by consumers. In section 4, we determine the optimal bundle composition and price for a monopolist as a function of various seller and consumer parameters, finding that finite-sized bundles are optimal when costs are finite. Then, in section 5, we present a game-theoretic analysis of an oligopoly with homogeneous consumer preferences. In a sequential game in which content choices precede price competition, we show that an oligopoly will achieve tacit collusion, producing a total output matching that of a monopolist. The profits earned by each firm will be positive, but will sum to less than that of a monopolist. If firms can instantaneously adjust their bundle composition, and thus can make content and pricing choices simultaneously, then each seller will independently set its bundle to that of the monopolist, and profits will be driven to zero. In section 6, we simulate an oligopoly in which the sellers employ a myopic best-response algorithm, showing that it reproduces the game-theoretic behavior. We then use simulation to investigate more complex scenarios that include heterogeneous preferences and more than two sellers, finding that these can exhibit more complex behavior in which both the prices and the bundle compositions can cycle, but the sellers can make positive profits. Finally, we summarize our findings and indicate some plans for future work in section 7.

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

2. RELATED WORK

Selling heterogeneous content as a bundle is not a new concept in the economics literature. Early papers considered a single provider bundling two goods [1, 31, 36]. More recent papers have extended the analysis to a monopolist bundling N goods [3, 8]. Analysis of a complete N-good bundling model with 2^N bundle combinations and N-dimensional consumer preferences quickly becomes intractable, even when competition is not considered. As a result, these papers restrict themselves to two simple alternatives: unbundling (selling each article in the collection separately) and pure bundling (offering the entire bundle for a single price). Some recent work does consider multiple content providers [16, 4, 15]. Each of these papers considers a rather specific construction of consumer preferences. None allow firms to control the degree of heterogeneity of content in their offerings. The characteristics that differentiate articles from each other are either unobserved or not manipulable by the firms. [1]

The model presented in this paper allows firms to choose product composition as well as price. There does exist a vast economics literature on endogenous product differentiation, but the form of differentiation typically studied is inappropriate for categorized information goods. The basic model originated by Hotelling [24] allows firms to choose which type of product they wish to offer. This choice is represented by a location on a line. Each consumer has an ideal product location, so that differences in consumer preferences are indicated by a distribution in the density of demand along the line. Consumers buy the product that is closest to their ideal. The primary issue considered by papers using this Hotelling model is the degree of differentiation firms choose. If firms choose the same location on the line, then they are minimally differentiated. Maximal differentiation is represented by firms choosing to locate at opposite endpoints. This occurs if transportation costs are quadratic, *i.e.* the penalty associated with not obtaining one's ideal product is a non-linear function of the distance between the ideal and the actual product consumed. Various references [2, 28, 39, 33, 13, 11] consider variations of models in which firms choose location (holding price choices fixed). Others broaden the analysis to the situation in which firms simultaneously control both location and price [10, 32, 43], or extend to firms the option to open multiple outlets [18, 29, 5].

All of these papers assume that firms sell each product separately. This is a reasonable restriction for the Hotelling model, since consumers are assumed to buy at most one product. However, in a market for information goods, consumers are likely to desire to read numerous different articles, and sellers are likely to find bundling advantageous.

3. MODEL

Our information bundling model assumes a population of B buyers and S sellers. First, consider the sellers. Each seller offers a single bundle[2] consisting of a selection of articles in some combination of choices from C categories. Specifically, at any given moment, each seller s will charge price p_s for a bundle consisting of a mix of articles in different categories: β_{sc} articles for category c, for each $0 \leq c < C$ [3]. Each seller experiences a cost r_s for producing and distributing each article that it sells, so the cost of producing a bundle is $r_s \sum_c \beta_{sc}$. [4]

Now, consider the buyers. We suppose that buyer b's valuation of n articles from category c is $v_{bc} f_{bc}(n)$, where v_{bc} can be thought of as an *intrinsic valuation* and $f_{bc}(n)$ can be thought of as a *saturation function*—a concave function of n satisfying $f_{bc}(0) = 0$, $f_{bc}(1) = 1$ and $f_{bc}(n) \leq n$ for $n \geq 2$. The intrinsic valuations v_{bc} are chosen independently from some distribution $g_c(v)$. Furthermore, articles in different categories are not substitutable for or complementary to one another, i.e. a buyer's valuation of a set of articles drawn from different categories is simply equal to the sum of its valuations for the articles in each separate category. Finally, the buyer b experiences a clutter cost ρ_b for sifting through each article that it receives, regardless of its valuation for that article.

Each buyer b decides to purchase q_{bs} bundles from each seller s, and in so doing it receives $q_{bs}\beta_{sc}$ articles in category c from seller s. Buyer b's valuation of this set of articles is $\sum_c v_{bc} f_{bc} \left(\sum_s q_{bs}\beta_{sc}\right) - \rho_b \sum_{sc} q_{bs}\beta_{sc}$, and it pays a total of $\sum_s p_s q_{bs}$ to obtain them. Therefore b's surplus is

$$\sigma_b = \sum_c v_{bc} f_{bc} \left(\sum_s q_{bs}\beta_{sc}\right) - \rho_b \sum_{sc} q_{bs}\beta_{sc} - \sum_s p_s q_{bs}. \tag{1}$$

Each buyer b attempts to set its purchase vector q_{bs} so as to maximize this quantity.

Returning to the seller, we see that seller s's revenue will be $p_s \sum_b q_{bs}$. Subtracting the bundle production cost $r_s \sum_c \beta_{sc}$ and normalizing by the number of buyers B, we find that seller s's profit per buyer is

$$\pi_s = \frac{1}{B} \sum_b q_{bs} \left[p_s - r_s \sum_c \beta_{sc}\right]. \tag{2}$$

Each seller s attempts to set p_s and β_{sc} so as to maximize this quantity.

For simplicity, we shall restrict q_{bs} to be either 0 or 1, i.e. a buyer purchases at most one bundle from each seller.

[1]Some papers do allow for distinct groups of products [30, 14]. But, in these models, the two firms produce variations of components that combine to form a system. A component (without its counterpart) is valueless. This is not a reasonable description for how content is valued by readers.

[2]A more sophisticated version of the model would allow s to set a price schedule based on the number of articles consumed by the buyer, rather than only permitting a single bundle to be offered. Even more generally, the price could depend explicitly on the categories and number of articles within each category in arbitrarily complex and nonlinear ways. The restriction to a single bundle here is primarily motivated by a desire to make the analysis tractable for multiple sellers, but that same fear of intractability could also practically limit sellers from availing themselves of more flexible price schedules. A broader range of price schedules is examined by Brooks et al. [7] for the case of a monopolist.

[3]We anticipate that, in markets in which specialization occurs, β_{sc} will be zero for many of the categories c.

[4]Any production costs associated with the creation of the article are assumed to be amortized over a very large number of buyers, and therefore negligible.

4. MONOPOLY ANALYSIS

As a first step in our analysis, consider a market with just a single seller. In this case, Eqs. 1 and 2 reduce to:

$$\pi = \frac{1}{B}\sum_b q_b \left[p - r\sum_c \beta_c\right] \tag{3}$$

and

$$\sigma_b = q_b\left[\sum_c v_{bc} f_{bc}(\beta_c) - \rho_b \sum_c \beta_c - p\right]. \tag{4}$$

Buyer b must decide whether to purchase the seller's bundle or not. Clearly, if the term in square brackets in Eq. 4 is positive, then it ought to purchase the bundle, i.e., it should set $q_b = 1$; otherwise, it should not purchase the bundle ($q_b = 0$). Substituting this condition into Eq. 3, we obtain:

$$\pi = \frac{1}{B}\sum_b \Theta\left(\sum_c v_{bc} f_{bc}(\beta_c) - \rho_b \sum_c \beta_c - p\right)\left[p - r\sum_c \beta_c\right] \tag{5}$$

where $\Theta(x)$ represents the step function, equal to 1 if $x > 0$ and 0 otherwise.

In the remainder of this section, we shall determine the price p and bundle $\vec{\beta}$ that maximizes a monopolist's profit, given two different assumptions about the normalized valuations v_{bc}. In subsection 4.1, we assume that they are equal for all buyers, i.e., $v_{bc} = v_c$. In subsection 4.2, we assume that the v_{bc} are all independent and distributed uniformly between zero and one.

4.1 Equal valuations

Assume that the valuations v_{bc}, the nonlinear saturation functions f_{bc}, and the clutter costs ρ_b are the same for all buyers. To determine the optimal price p^* and the optimal bundle $\vec{\beta}^*$, we first compute the optimal price for *any* bundle, and then use this result to compute the optimal bundle.

To compute the optimal price for any given setting of $\vec{\beta}$, note that, in Eq. 5, the monopolist's profit is maximized when p is maximized subject to the constraint given by the Θ function. Thus the monopolist's profit is maximized at the price

$$p^* = \sum_c v_{bc} f_c(\beta_c) - \rho \sum_c \beta_c. \tag{6}$$

Substituting this price into Eq. 5 yields a profit per buyer of

$$\pi = \sum_c v_c f_c(\beta_c) - (r + \rho)\beta_c \tag{7}$$

To compute the profit-maximizing bundle $\vec{\beta}^*$, note that each term in the summation appearing in Eq. 7 is independent. Therefore, we can solve independently for each component β_c. The optimal value β_c^* is simply the one that maximizes $v_c f_c(\beta_c) - (r + \rho)\beta_c$. Thus β_c^* depends simply on the ratio $\gamma \equiv \frac{r+\rho}{v_c}$. Since β_c^* is an integer, there will be regions of γ that share the same value of β_c^*. The boundary between a region in which $\beta_c^* = \beta_0$ and one in which $\beta_c^* = \beta_0 + 1$ occurs when

$$f_c(\beta_0 + 1) - f_c(\beta_0) = \gamma = \frac{r+\rho}{v_c}. \tag{8}$$

The boundaries between regions with different $\vec{\beta}^*$ are a superposition of the boundaries of the individual components. Figure 1 illustrates these boundaries as a function of ρ and r for a specific scenario. The number of categories is $C = 2$, and the saturation function for each category is $f_c(x) = \sqrt{x}$. [5] The valuations v_1 and v_2 are 1 and 1/2, respectively, *i.e.*, all of the buyers value a single article in category 1 twice as highly as they value one in category 2.

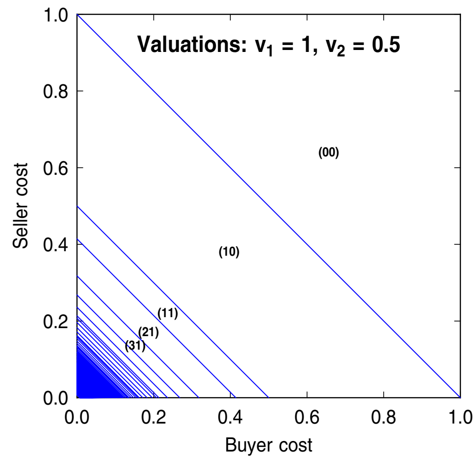

Figure 1: Optimal $\vec{\beta}^*$, indicated in diagram as (β_1^* β_2^*), as a function of the buyer clutter cost ρ and the seller production cost r. The number of categories is $C = 2$. Each buyer has equal valuations for articles in the two categories: specifically, $v_1 = 1$ and $v_2 = 0.5$.

The regions in Fig. 1 are labeled according to their optimal ($\beta_1^*\beta_2^*$). Several trends are readily apparent. First, the boundaries are linear, of the form $r + \rho = const$, *i.e.*, the optimal β^* depends solely on the sum of the buyers' clutter cost and the seller's production cost. This follows naturally from the fact that β_c^* depends only on the ratio $\gamma \equiv \frac{r+\rho}{v_c}$.

A second observation is that the individual components of the bundle vector, β_1^* and β_2^*, increase as $r + \rho$ decreases. Third, the rate at which the β_c^* decrease increases as $r + \rho$ is reduced, resulting in severe crowding of the boundaries as $r + \rho \to 0$. In this particular scenario, for example, β_1^* shifts from β_0 to the next higher value when $r + \rho$ is reduced below the threshold defined by $\sqrt{\beta_0 + 1} - \sqrt{\beta_0}$. Thus the transition from $\beta_1^* = 0$ to $\beta_1^* = 1$ occurs at $r + \rho = 1$, and successive transitions occur at $r + \rho = \sqrt{2} - 1 = 0.4142$, $\sqrt{3} - \sqrt{2} = 0.3178$, at $2 - \sqrt{3} = 0.2679$, etc. For β_2^*, the corresponding thresholds occur at half of these values. Thus the boundaries displayed in Fig. 1 can be viewed as a superposition of two infinite sets of boundaries, one of which is half the scale of the other. Using Eq. 8, one can show that the observed inverse relationship of β_1^* and β_2^* to $r + \rho$ holds for any concave saturation function f.

[5] This can be taken as an instance of the somewhat more general class of saturation functions given by $f_c(x) = x^{\alpha_c}$, with $0 < \alpha_c < 1$. A saturation function of this form with $\alpha_c \to 0$ implies that consumers have no interest in obtaining more than one article in category c (extreme subadditivity), while $\alpha_c \to 1$ implies that consumers have an insatiable appetite for articles of that category, and are willing to consume them in essentially infinite quantity.

The third effect—the severe crowding together of the boundaries when costs are low—can be understand by noting that, for sufficiently small $r + \rho$, Eq. 8 can be approximated as

$$f_c'(\beta_0) \approx \gamma = \frac{r+\rho}{v_c}, \quad (9)$$

which yields

$$\beta_c^* \approx \left(\frac{v_c}{2\,(r+\rho)}\right)^2. \quad (10)$$

Taking the derivative with respect to $r + \rho$, we find that the distance between successive boundaries diminishes approximately as $(r+\rho)^3$, which explains the severe crowding of boundaries at low costs. More generally, if the saturation function is of the form $f(x) = x^\alpha$, then a bit of algebra shows that this distance scales as $(r+\rho)^{\frac{2-\alpha}{1-\alpha}}$. Thus boundary crowding occurs even for extreme saturation ($\alpha \to 0$), and becomes increasingly severe for weaker saturation ($\alpha \to 1$).

4.2 Uniform distribution of valuations

Here we determine the optimal price and bundle for a monopolist under the assumption that the buyers' preferences are heterogeneous. Specifically, we suppose that the valuations v_{bc} are independent and uniformly distributed between 0 and 1. As before, we assume that the nonlinear saturation functions f_{bc} and the clutter costs ρ_b are the same for all buyers.

It is helpful in this case to visualize B individual valuation vectors (one for each buyer) as points lying within a C-dimensional hypercube. At any given price p and bundle $\vec{\beta}$, some of the buyers may purchase the bundle while others will not. From the step function in Eq. 5, it follows that the purchasers are those with valuation vectors v_{bc} satisfying:

$$\sum_c v_{bc} f_c(\beta_c) > p + \rho \sum_c \beta_c. \quad (11)$$

For any fixed p and $\vec{\beta}$, the boundary between purchasers and non-purchasers is a hyperplane that cuts the unit hypercube to form a simplex of non-purchasers. (Any component c for which $\beta_c = 0$ does not contribute to either side of Eq. 11. Therefore the problem reduces to a hypercube and hyperplane in a subspace consisting of just those components c for which $\beta_c > 0$.) In the limit as the number of buyers $B \to \infty$, the points v_{bc} are distributed uniformly, and the fraction of non-purchasers is just the volume of the simplex. Thus the fraction of purchasers is that of the hypercube (1) minus this volume. Substituting this for all terms in Eq. 5 except for the term in square brackets, we obtain the seller's profit per buyer:

$$\pi = \left[1 - \frac{(p + \rho(\sum_c \beta_c))^\nu}{\nu! \prod_{c \in \beta_c > 0} f_c(\beta_c)}\right] \left[p - r \sum_c \beta_c\right] \quad (12)$$

where

$$\nu \equiv \sum_{c \in (\beta_c > 0)} 1 \quad (13)$$

is the number of non-zero components of $\vec{\beta}$.

The optimal values of p and $\vec{\beta}$ can now be obtained by numerically optimizing Eq. 12. Figure 2 displays the resultant regions of r and ρ for which a given $\vec{\beta}$ is optimal, under exactly the same conditions as in Fig. 1, except that the v_{bc} are distributed uniformly rather than being homogeneous.

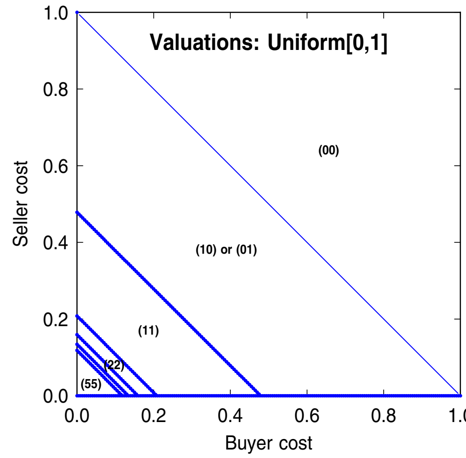

Figure 2: Optimal $\vec{\beta}^*$, indicated in diagram as (β_1^* β_2^*), as a function of the buyer clutter cost ρ and the seller production cost r. The number of categories is $C = 2$. Each buyer's valuation for each category is chosen uniformly from the unit interval. Regions in which the optimal number of articles in a category exceed 5 are not shown.

Figure 2 is qualitatively similar to Fig. 1 in a number of respects. The regions are again separated by boundaries of the form $r + \rho = const$, although this property is not immediately obvious from Eq. 12. Again, the optimal bundle size is inversely related to the buyer and seller costs, and the boundaries become increasingly crowded as $r + \rho \to 0$.

The results for both uniform and equal valuations are strongly reminiscent of those obtained for a related information filtering model that has been investigated previously [27]. For very high costs, such that the combined buyer and seller costs per article exceed 1, the seller does not have a viable business, and the optimal action is to offer no articles. For combined costs that are relatively high but do not exceed 1, the optimal action is to offer a bundle consisting of one article in one category. When the combined costs are somewhat lower, it becomes worthwhile for the seller to offer a bundle that includes both categories, and the number of articles in each category becomes larger as the costs are decreased, tending to infinity as the costs go to zero.

The main difference between the information filtering and information bundling models is the interpretation of the product parameter vector $\vec{\beta}$. In the bundling model, information is delivered in discrete bundles, and the β parameters specify a definite number of articles appearing in each category in every bundle. In the information filtering model, the information is delivered article by article, and the seller's β parameters represent real number probabilities for articles in particular categories to be let through by the filter.

5. OLIGOPOLY ANALYSIS

In this section, we perform a game-theoretic analysis of an oligopoly. In order to make the analysis tractable, we only consider the case in which consumer preferences are homogeneous. In the next two subsections, we consider two distinct scenarios. First, in subsection 5.1, we suppose that sellers make their content and price decisions sequentially. This assumption is appropriate when content choices are much less easily adjustable than prices, which could be due to the lead times required to develop or advertise a new product. In the first stage, the sellers simultaneously set their content parameters $\vec{\beta}$. Then, in the second stage, they observe the content parameters of all the other sellers and then simultaneously set their prices based on these settings. In the first stage, sellers make their content decisions with full awareness that, in the second stage, they and their competitors will know one another's content decisions before they set their prices. In subsection 5.2, we consider an alternate scenario in which content and prices are set simultaneously. In each case, the analysis is simplified by taking the content parameters β_{sc} to be continuous variables.

5.1 Sequential Content and Price Choices

In this subsection, we analyze the case in which the sellers first choose content simultaneously and then, having observed one another's content choices, they simultaneously set their prices.

First, we derive the equilibrium prices that will result for any given content decisions. Since consumer preferences are homogeneous, all consumers will behave alike, i.e., $q_{bs} = q_s$ and $v_{bc} = v_c$ for all b. To make a positive profit, each seller s must behave in such a manner that $q_s = 1$. Substituting this into Eqs. 1 and 2, we obtain:

$$\sigma = \sum_c v_c f(\sum_s \beta_{sc}) - \rho \sum_{sc} \beta_{sc} - \sum_s \rho_s \qquad (14)$$

and

$$\pi_s = p_s - r \sum_c \beta_{sc} \qquad (15)$$

According to Eq. 15, a seller would like to price its bundle as high as possible, subject to the constraint that $q_s = 1$. In other words, the seller must price the bundle just low enough to convince the buyer to purchase it. This is the point at which the buyer is indifferent between purchasing the bundle and not purchasing it, given any other purchases that the buyer intends to make. In other words, for seller s', the optimal price satisfies the condition:

$$\begin{aligned} \sum_c v_c f(\sum_s \beta_{sc}) - \rho \sum_{sc} \beta_{sc} - \sum_s p_s &= \\ \sum_c v_c f(\sum_{s \neq s'} \beta_{sc}) - \rho \sum_{s \neq s', c} \beta_{sc} - \sum_{s \neq s'} p_s & \end{aligned} \qquad (16)$$

or

$$p^*_{s'} = \sum_c v_c \left[f(\sum_s \beta_{sc}) - f(\sum_{s \neq s'} \beta_{sc}) \right] - \rho \sum_c \beta_{s'c} \qquad (17)$$

Substituting Eq. 17 into Eq. 15, we find that, at this equilibrium, the sellers' profits are given by:

$$\pi^*_{s'} = \sum_c v_c \left[f(\sum_s \beta_{sc}) - f(\sum_{s \neq s'} \beta_{sc}) \right] - (r + \rho) \sum_c \beta_{s'c} \qquad (18)$$

Treating the β_{sc} as continuous variables, we can solve for their optimal value by taking the appropriate partial differentials of Eq. 18 and setting them to zero:

$$\frac{\partial \pi_s}{\partial \beta_{sc}} = v_c \frac{\partial f(\sum_s \beta_{sc})}{\partial \beta_{sc}} - (r + \rho) = 0 \qquad (19)$$

Defining the total content vector $\beta_c^{\text{tot}} \equiv \sum_s \beta_{sc}$, this can be rewritten as

$$f'(\beta_c^{\text{tot}}) = \frac{(r + \rho)}{v_c} \qquad (20)$$

The interpretation is that there is a continuum of Nash equilibria—any combination of β_{sc} that sums to a value β_c^{tot} satisfying Eq. 20 is a Nash equilibrium. Comparison with Eq. 9 yields a further interpretation: at the Nash equilibrium, the total number of articles in each category produced by all sellers together is equivalent to what would be offered by a monopolist.

After settling upon one of the Nash equilibria defined by Eq. 20 in the first stage, the sellers set their prices in the second stage according to Eq. 17. The prices depend on the particular Nash equilibrium that is realized, as do the sellers' profits. This is illustrated in Figure 3, which plots the profit for a seller (and for both sellers combined) in a two-seller market in which the buyers' valuations are $v_1 = 1$ and $v_i = 0$ for $i > 1$, and the saturation function is $f(x) = \sqrt{x}$. The costs are chosen to be $r = 0.2$ and $\rho = 0.2$. Under these conditions, the optimal value of β_1^{tot} is $(2(r+\rho))^{-2} = 1.5625$, and so there is a continuum of Nash equilibria in which $\beta_{11} + \beta_{21} = 1.5625$. Seller 1's profits increase monotonically with the amount of content it provides in category 1, β_{11}. The combined duopoly profit (the thick line) is greatest when either seller 1 or seller 2 provide all of the content. When both sellers produce a finite amount of content, the total duopoly profit is less than would be earned by a monopolist. All equilibria are efficient, since the consumer ends up purchasing the efficient number of articles, but how the surplus is shared between the consumers and the two sellers is greatly affected by the choice of equilibrium.

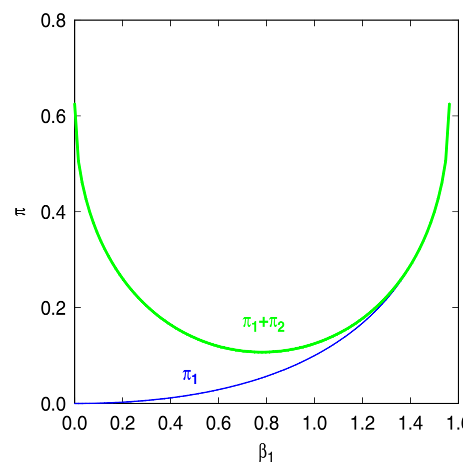

Figure 3: Duopoly profits. Lower curve: profit for seller 1 as a function of β_1. Upper curve: combined profit for sellers 1 and 2. Other parameters are as described in text.

5.2 Simultaneous Content and Price Choices

The above results depend critically on the assumption that sellers' decisions occur sequentially: first content is chosen, then pricing decisions are made. However, we show in this subsection that, if content and pricing decisions are made simultaneously, then none of the outcomes found in the previous subsection constitute a Nash equilibrium. For analytical tractability, we focus here on the duopoly case.

For an outcome to be a Nash equilibrium, each seller must be responding optimally to the other sellers' product and pricing choices. Consider one of the sequential equilibria derived in subsection 5.1. For this to be an equilibrium in the simultaneous game, Seller 2 must not have an incentive to alter its bundle and price (given Seller 1's bundle composition and price). Suppose Seller 1 is producing $\vec{\beta}_1^*$ and Seller 2 is producing bundle $\vec{\beta}_2^*$. Seller 1's price, p_1^*, is given by Eq. 17. Now, Seller 2 considers increasing its amount of type c' content by ϵ. Denote this new proposed bundle as $\vec{\beta}_2^{dev}$, with $\beta_{2c}^{dev} = \beta_{2c}^*$ for all $c \neq c'$ and $\beta_{2c'}^{dev} = \beta_{2c'}^* + \epsilon$. If Seller 2 can simultaneously choose a new price p_2^{dev} such that its profit is increased, then the sequential Nash equilibrium will not be a Nash equilibrium in the simultaneous game.

Note that, by increasing its content, Seller 2 has upset the balance in which the buyer was indifferent between buying one of the bundles or both. Thus the buyer will choose between Seller 1's and Seller 2's bundles. Seller 2 must set its new price, p_2^{dev}, just low enough so that the buyer will prefer its bundle to that of Seller 1. From Eq. 14, this will occur provided that

$$\begin{aligned}\sum_c v_c f(\beta_{2c}^{\text{dev}}) - \rho \sum_c \beta_{2c}^{\text{dev}} - p_2^{\text{dev}} \quad > \\ \sum_c v_c f(\beta_{1c}^*) - \rho \sum_c \beta_{1c}^* - p_1^*.\end{aligned} \tag{21}$$

Using Eq. 17 to substitute for p_1^* and solving for the value of p_2^{dev} that just makes the above relationship an equality, we obtain:

$$\begin{aligned}p_2^{\text{dev}} &= \sum_c v_c \left[f(\beta_{1c}^* + \beta_{2c}^*) - f(\beta_{1c}^*) - f(\beta_{2c}^*)\right] \\ &+ \sum_c v_c f(\beta_{2c}^{\text{dev}}) - \rho \sum_c \beta_{2c}^{\text{dev}}.\end{aligned} \tag{22}$$

Let $\Delta\pi$ denote Seller 2's change in profit if it deviates to a larger bundle. Then, using Eq. 15,

$$\begin{aligned}\Delta\pi &= \left[p_2^{\text{dev}} - r \sum_c \beta_{2c}^{\text{dev}}\right] - \left[p_2^* - r \sum_c \beta_{2c}^*\right] \\ &= p_2^{\text{dev}} - p_2^* - r\epsilon \\ &= v_{c'} \left[f(\beta_{2c'}^* + \epsilon) - f(\beta_{2c'}^*)\right] - (r + \rho)\epsilon.\end{aligned} \tag{23}$$

Letting $\epsilon \to 0$, we find that $\Delta\pi$ is positive (and hence the deviation is profitable) if:

$$\frac{\partial f(\beta_{2c'}^*)}{\partial \beta_{2c'}} > \frac{r + \rho}{v_{c'}} \tag{24}$$

Recall from Eq. 19 that, at the original outcome,

$$\frac{\partial f(\beta_{1c'}^* + \beta_{2c'}^*)}{\partial \beta_{2c'}} = \frac{r + \rho}{v_{c'}}. \tag{25}$$

Since $f(\cdot)$ is a concave function,

$$\frac{\partial f(\beta_{2c'}^*)}{\partial \beta_{2c'}} > \frac{\partial f(\beta_{1c'}^* + \beta_{2c'}^*)}{\partial \beta_{2c'}}. \tag{26}$$

Thus, for any outcome derived in the previous section, Seller 2 could increase its profit by deviating to a larger bundle and higher price. Since the game is symmetric, the same analysis could be applied to Seller 1 to show that it too has an incentive to try to increase its bundle.

The desire of both sellers to increase their bundle size would lead to the total supply of articles in each category exceeding the efficient level β_c^{tot}, as defined in Eq. 20. But, if this occurs, there cannot be a price equilibrium in which the consumers always buy from all s sellers unless the sellers earn zero profits and hence are indifferent as to whether or not their bundle is purchased. In general, even if all other sellers sell their respective bundles at cost, a seller could earn a positive profit by providing a better bundle and selling it for some low (but non-zero) mark-up. This is not possible only if a competitor (s') already offers the consumer's most preferred bundle: $\vec{\beta}_{s'} = \vec{\beta}_c^{\text{tot}}$. But s' maximizes its profit by selling the optimum bundle at cost only if at least one other seller is doing the same. Thus, when sellers choose bundles and prices simultaneously, the Nash Equilibrium involves at least two sellers offering the bundle $\vec{\beta}_c^{\text{tot}}$ for a price $r \sum_c \beta_c^{\text{tot}}$. These sellers do not necessarily offer identical articles, but they do provide the same total amount of content from each category. A consumer finds the bundles equally valuable, but does not want to purchase multiple bundles because a larger conglomeration does not add enough value to offset the increased clutter cost.

5.3 Summary

As expected, the socially efficient bundle size is smaller as the cost of creating content increases. A monopolist would produce this efficient bundle (since it is able to capture the entire consumers' surplus). For the oligopoly model, timing has a large impact on the resultant outcome. In the sequential version, sellers can predict the equilibrium prices that will result for any given product offering configuration. This allows for tacit collusion. Other papers have shown similarly that pre-commitment to investment levels such as in capacity or advertising [17, 19, 38, 42], location [24] or price (the Stackelberg leader-follower model) [45, 6] allow firms to influence the degree of competition in their industry. In the current model, when content decisions precede pricing decisions, there is a continuum of equilibria in which the sum of articles produced by an oligopoly equals the monopoly (and hence socially efficient) level. However, these outcomes are not equilibria if content and pricing choices are made simultaneously. In this case, sellers have an incentive to produce larger bundles and siphon business away from their competitors. However, if sellers overproduce, consumers will not buy all offered bundles. Thus, profits are driven to zero. In equilibrium, at least two sellers offer to sell the optimal bundle at cost.

6. SIMULATION

The game-theoretic analysis of the previous section assumed that the sellers chose either their prices or both their prices and their bundle composition simultaneously. It is also worthwhile (and probably more realistic) to investigate a scenario in which sellers asynchronously update their prices and bundling choices in response to the choices made by their competitors.

In this section, we assume that each seller s knows the buyers' parameters (their valuations and saturation functions), and that it also knows the current price and bundling parameters for its competitors. Using this information, seller s sets its p_s and β_{sc} to the values that maximize its expected profit, given the current state of the market. In other words, each seller employs a myopic best-response strategy that is optimal in the short-term, up until the moment when some other seller resets its parameters. The myopic best-response strategy, sometimes referred to as the "myoptimal" strategy, is attractively simple to describe and implement, and has been studied in several other models of software agent markets [22, 27, 35, 20].

In order to study the behavior of a market in which the sellers use an asynchronous, myoptimal strategy, we simulate its evolution from a given initial condition. The simulation proceeds as follows. At each discrete time step, a buyer or seller is randomly selected to act. If the selected agent is a buyer b, it experiments with all possible sets of bundles, evaluates Eq. 1 for each, and chooses to purchase q^*_{bs} bundles from each seller s, where q^*_{bs} represents the vector that maximizes b's surplus. In our simulation, we restrict q_{bs} to be either 0 or 1, *i.e.*, the buyer purchases at most one bundle from each seller, so the total number of bundle sets to explore is 2^S. If S is not larger than 10 or so, it is feasible for the buyer to search exhaustively over all of the options.

If, on the other hand, the selected agent is a seller s, then it re-evaluates its parameter settings using a myoptimal policy. It implements this by evaluating the expected profit (Eq. 2) for all possible p_s and β_{sc}, and then selecting the values at which the expected profit is maximized. Note that, in order to evaluate Eq. 2 for any given setting of p_s and β_{sc}, the seller must compute q_{bs} for each b. In other words, it must simulate each buyer's decision under each possible choice of p_s and β_{sc}. Depending on the various market parameters, this can be a very time-consuming operation. Suppose that β_{sc} is restricted to integer values in the range $0 \leq \beta_{sc} \leq \beta_{\max}$. Then, if the possible prices are taken to be of the form $0 < p = n\epsilon \leq p_{\max}$, where n is an integer and ϵ is the price quantum, then the seller must loop over $p_{\max}\epsilon^{-1}(\beta_{\max}+1)^C$ candidate values of p_s and β_{sc}. For each of these candidates, the seller must simulate B buyers, each of which is examining 2^S possible bundle sets and selecting the best to determine q_{bs}.

For example, one of the more computationally intensive simulations we have run used the parameters $S = 5$, $B = 500$, $C = 3$, $\beta_{\max} = 4$, $p_{\max} = 1.0$, and $\epsilon = 0.01$. Therefore, every time a seller re-evaluated its parameters, it had to loop over 6400 candidate settings of p_s and β_{sc}, each of which required 16000 evaluations of Eq. 1. Therefore a single time step of the simulation required 16000 evaluations of Eq. 1 if the selected agent was a buyer, and over 100 million evaluations if the selected agent was a seller.

The simulator permits us to follow the evolution of p_s and β_{sc} over time, along with other quantities of interest on which they depend. In addition to providing insight into market behavior that ensues when sellers are permitted to observe and respond asynchronously to one another's price and bundling choices, the simulator also allows us to extend our study to more than two sellers, heterogeneous preferences, and integer-valued β parameters.

First, consider a market with $S = 2$ myoptimal sellers, $C = 2$ categories, and 500 buyers with homogeneous intrinsic valuations $v_1 = 1$ and $v_2 = 0.5$ and a square-root saturation function. If the production cost r and the clutter cost ρ are both 0.2, then the optimal bundle for a monopolist is (21) (see Figure 1). Starting from a random initial condition, the simulation evolves as shown in Figure 4. The two sellers immediately enter into a battle over the bundle (21), driving one another's prices down until they settle at a price of 0.61, with is exactly one price quantum above the bundle production cost of $(2+1)r = 0.60$. Thus the myoptimal sellers reach (essentially) the game-theoretic equilibrium derived in section 5.2 for sellers that can set their price and bundle simultaneously: the sellers each choose the optimal monopolist bundle and drive one another down to (nearly) cost.

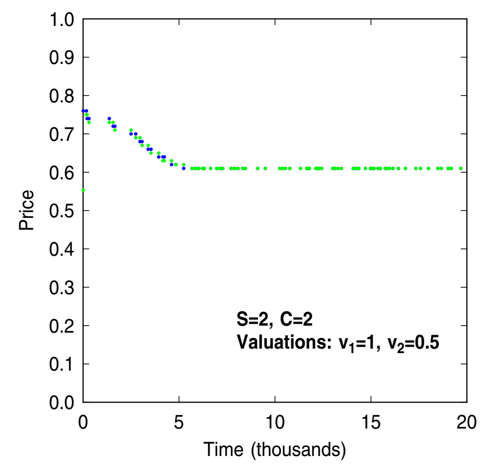

Figure 4: Simulated price dynamics for 2 sellers, 2 categories, and 500 buyers each with intrinsic valuations $v_1 = 1$, $v_2 = 1/2$. For all buyers, $\rho = 0.2$. For all sellers, $r = 0.2$.

With the simulator, we can explore markets with more sellers and more categories. Fig. 5 depicts the price dynamics for 5 sellers who can choose from among 3 categories. All other parameters are exactly as in Fig. 4, with the addition that $v_3 = 1/3$. During the first 1000 time steps, the sellers all compete for (200), but between times 1000 and 2000 they all switch over into a competition for the monopolist bundle—(210) in this case. Again, competition forces prices down to 0.61, which is one price quantum above the bundle production cost, $(2+1+0)r = 0.60$.

The simulator also permits us to explore what happens when consumer preferences are heterogeneous—a case for which game-theoretic analysis appears to be quite difficult. Figure 6 illustrates a market that is almost identical to that depicted in Figure 4, except that the consumers' intrinsic valuations v_{bc} are drawn uniformly from the unit interval. For $r = \rho = 0.2$, the optimal monopolist bundle is (11) (see Fig. 2). After a few rounds of price setting, the sellers enter into a price war over (11), and through successive undercutting the price heads down toward the bundle production cost, $2r = 0.4$. However, when the price gets down as low as 0.44, the best response is not to undercut to 0.43. Instead, the next seller to re-evaluate its price switches to the

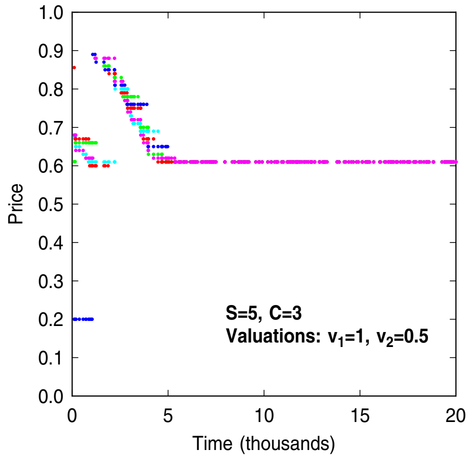

Figure 5: Simulated price dynamics for 5 sellers, 3 categories, and 500 buyers each with intrinsic valuations $v_1 = 1$, $v_2 = 1/2$, and $v_3 = 1/3$. For all buyers, $\rho = 0.2$. For all sellers, $r = 0.2$.

bundle (01) at price 0.35. The other seller responds to this by setting its bundle to (11) and its price to 0.70, whereupon the other seller responds by reverting to (11) and just undercutting to 0.69. This incites another price war cycle, and the process continues indefinitely. The cycles do not repeat one another perfectly because, while the individual sellers behave deterministically, they adjust their prices and bundle compositions in a random order.

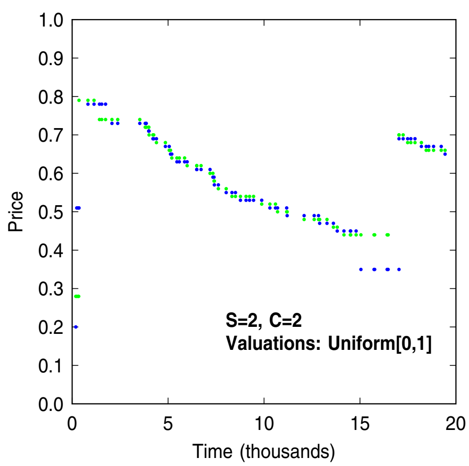

Figure 6: Simulated price dynamics for 2 sellers, 2 categories, and 500 buyers with uniformly distributed valuations. For all buyers, $\rho = 0.2$. For all sellers, $r = 0.2$.

Averaged over time, the profits for heterogeneous preferences are *not* nearly zero, as they were for homogeneous preferences. This is an interesting reversal of the situation for a single seller. A monopolist can extract all of the surplus from consumers that have homogeneous valuations, but can only extract a fraction of the surplus from consumers with heterogeneous valuations. For example, for homogeneous preferences of the form assumed in Fig. 4, the optimal price, bundle, and profit are 1.314, (21), and 0.714, while for heterogeneous preferences of the form assumed in Fig. 6 they are and 0.725, (11), and 0.119, respectively. In the corresponding duopoly, the measured prices, bundles, and profits[6] are 0.61, (21), and 0.005 for the homogeneous case.[7] For the heterogeneous case, the prices and bundle compositions cycle indefinitely, but averaged over these cycles the measured profit per buyer transaction is roughly 0.034—much *higher* than for the homogeneous case. Note that the profit in the homogeneous case is proportional to the price quantum, which may be arbitrarily small, while the profit in the heterogeneous case does not depend in any essential way on the size of the price quantum.

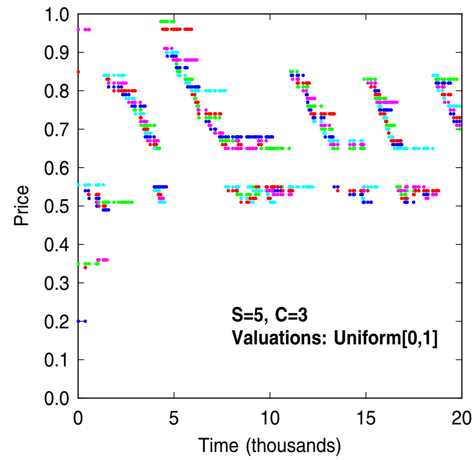

Figure 7: Simulated price dynamics for 5 sellers, 3 categories, and 500 buyers with uniformly distributed valuations. For all buyers, $\rho = 0.2$. For all sellers, $r = 0.2$.

Cyclical price and bundle composition wars can get more complex with more sellers and categories. Figure 7 shows the price dynamics when all parameters are kept as in Fig. 6, except that the number of sellers is increased to 5 and the number of categories is increased to 3. In this case, the optimal monopolist bundle is (111). By time step 3000, all 5 sellers have settled into a price war over the (111) bundle. However, near time step 4000, one of the sellers finds it more advantageous to switch to the (110) bundle and drop its price dramatically, from 0.68 to 0.55. Three other sellers quickly follow suit, successively undercutting one another. This effect shows up clearly as a vertical gap in the price dynamics between approximately 0.55 and 0.65. At this point, the fifth seller is the only one offering category 3. It finds that it can maximize its profit by adhering to (111) and jacking its price up to 0.98. As soon as it does so, the

[6]The profits reported here for multi-seller simulations are normalized by dividing the profit that accrues during a given interval by the number of buyer actions that are taken systemwide during that interval. This supports a fair comparison with the monopolist profits, which are computed from Eqs. 5 divided by the number of buyers.

[7]The profit of 0.005 can be understood as an even split between the two sellers of the price-quantum worth of profit per sale.

other sellers switch to (111) and attempt to undercut one another, and the cycle begins anew.

Even more complex cycles can be observed as the costs r and ρ are decreased because the optimal monopolist bundle grows larger. For example, if, in the market depicted in Fig. 7, r and ρ are reduced from 0.2 to 0.125, the sellers' bundle compositions cycle irregularly through (112), (121), (211), (102), (111), (101), and (011).

Cyclical price wars have been observed previously in a variety of models of agent economies in which the agents employ a myoptimal strategy [22, 35, 20]. The more complex cycles observed here, which involve bundle composition as well, are reminiscent of behavior seen previously in studies of a related information filtering model that included information categories [27, 26].[8] Cyclical price and bundle composition wars are symptomatic of an underlying multi-peaked *profit landscape*. We believe that such landscapes may occur in a broad array of markets. Even when sellers use strategies other than myopic best-response, a multi-peaked landscape can lead to non-equilibrium market dynamics [22, 27].

From another perspective, the fact that sellers' use of myoptimal strategies induces non-equilibrium market dynamics when buyers' preferences are heterogeneous indicates that, if sellers were to make their price and bundle choices simultaneously, there would be a mixed-strategy solution. Such a phenomenon occurs in a much different framework presented by Hopkins [23]. In this continuous-time best-response model, agents choose the price that performs best against the predicted play of their opponents, which is estimated from past observations. Hopkins finds that the strategy choices and average payoffs observed in the dynamic game are very similar to the game's unique mixed-strategy equilibrium.

Our model does satisfy the sufficient condition found by Dasgupta and Maskin [9] for a mixed strategy equilibrium to exist.[9] Unfortunately, analytic computation of the mixed-strategy game-theoretic solution is likely to be difficult if not impossible. Papers that make such calculations greatly simplify the structure of consumer demand and limit firm strategies to competing only on price [41, 44] or only on product characteristics [34, 40], but not both simultaneously. However, no-regret learning techniques have been shown to be capable of finding Nash equilibria for related market games in which there are several hundred possible strategies [21]. It would be of great interest to perform such a computation and compare the game-theoretic probabilistic equilibrium with the simulated results obtained here.

[8] One important difference between the information bundling model presented here and the information filtering model studied previously is that the analog of the bundle composition parameters β is a *probability* for an article to be included in an information stream, as opposed to an integer number of articles included in an information bundle.

[9] Dasgupta and Maskin prove the existence of a mixed strategy equilibrium for a robust class of games where agents have discontinuous payoff functions. Our model satisfies their sufficient condition that the sum of agents' payoffs be upper semi-continuous. A discontinuity in each firm's payoff function is present because a discrete jump in profit occurs if two firms choose the same content configuration and one firm reduces its price from just above to just below its competitor's price. However, since this marginal reduction in price does not lead to a discrete change in combined profits for the two firms, the upper semi-continuity property is maintained.

7. CONCLUSIONS

Recognizing that online content providers will have the opportunity to adjust both prices and content categories dynamically and automatically, we have introduced the element of categorization into the study of information bundling. Our ultimate aim is to gain a good understanding of how to create effective individual strategies for sellers in such markets. This in turn necessitates a fundamental understanding of the dynamics of such markets when several sellers are competing to offer related products.

Starting first with a monopolist, we characterized how the composition of the optimal bundle is affected by the sellers' production costs and the buyers' clutter (or consumption) costs. The optimal bundle size is small when costs are high, and increases as the costs are decreased. Next, when we introduced competition, we found that monopolist and oligopolist behavior are closely related in ways that depend upon the specific assumptions. When sellers choose their bundle and their price sequentially, and buyers' preferences are homogeneous, there is a continuum of Nash equilibria in which the total article production of the sellers is equivalent to that which would be offered by a monopolist. However, when sellers choose their bundle and their price simultaneously, and buyers' preferences are homogeneous, then every seller individually sets its bundle size to that of the monopolist, and prices get driven down to cost. For simultaneous choice of bundles and prices, and with our particular choice of heterogeneous preferences, a simulation of myoptimal sellers demonstrated that it is possible for sellers to sustain positive profits on average. Interestingly, in this case the market exhibits unending cycles in prices and bundle composition, with the monopolist bundle representing the largest bundle that is ever offered.

Monopolists prefer homogeneous consumer preferences to heterogeneous ones because they can exploit homogeneity to extract all of the surplus. However, the situation is reversed when there are two or more sellers. Homogeneity makes it possible for a seller to (temporarily) grab all of the market share by undercutting its rivals, but this short-sighted strategy leads to a price war that ultimately leads to negligible profits for all sellers. When consumers are heterogeneous in their preferences, no seller can completely satisfy the entire market, and therefore several different profitable niches may be available to sellers. In these simulations, it was possible for sellers to jump easily (without cost) to any niches that they desired, and they continually did so, creating a never-ending cycle of price and bundle composition wars.[10] However, despite the failure of the market to settle to an equilibrium, the sellers were able to make finite profits.

Even in electronic marketplaces, friction will exist at some level. It remains to investigate various forms of friction and their impact on market behavior. Even for online information goods, there are sunk costs (e.g. costs for creating the content in the first place) that may cause sellers to be less nimble in their choice of bundle composition. However, these costs will typically be less than for physical goods. Another source of friction occurs on the buyer side. Buyers may experience some cost for obtaining information about

[10] If the sellers were free to occupy several niches simultaneously, either by offering several different bundles at different prices, or by setting some more complex price schedule, it is conceivable that this phenomenon would no longer persist. This is definitely worth exploring in future work.

price and bundle composition, although again these costs are likely to be lowered by the Internet and technologies that exploit it, such as comparison shopping agents. If the bundles are sufficiently complex in nature, buyers may also experience some computational cost for optimizing their selection of bundles from different vendors. These types of costs may prevent buyers from fully optimizing their purchases. Price dispersion theory [12, 37, 46] suggests that this is another mechanism by which positive profits can be sustained.

Another important effect that remains as a topic for future research is the ability of sellers to offer several different configurations rather than just a single bundle. This opens up a vastly greater number of options for the sellers, most likely making the optimization required for best-response infeasible. Sellers would have to employ heuristic optimization approaches, and it will most likely be necessary to rely entirely on simulation approaches to understand the behavior of such markets.

Finally, it should be pointed out that the analysis and simulation presented in this paper assumed that both sellers and buyers have a great deal of knowledge of the state of the market. In reality, buyers and sellers (or software agents operating on their behalf) will have to learn or infer market parameters. The study of agent learning in such markets is just in its infancy [7, 25].

Acknowledgments

The authors gratefully acknowledge support from an IBM University Partnership grant and from the IBM Institute for Advanced Commerce. We are also grateful to Jeff MacKie-Mason for his guidance, and especially for his valuable input on our information bundling model and its relationship to prior models in the economics literature.

8. REFERENCES

[1] W. Adams and J. Yellen. Commodity bundling and the burden of monopoly. *Quarterly Journal of Economics*, 90:475–498, 1976.

[2] S. P. Anderson, A. de Palma, and J.-F. Thisse. *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge, Massachusetts, 1992.

[3] 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, 1998.

[4] Y. Bakos and E. Brynjolfsson. Bundling and competition on the internet: Aggregation strategies for information goods. Working paper series, MIT Sloan School of Management, April 1999.

[5] D. Bensaid and A. de Palma. Spatial multiple product oligopoly. Unpublished manuscript, Banque de France, 1993.

[6] M. Boyer and M. Moreaux. Being a leader or a follower: Reflections on the distribution of roles in duopoly. *International Journal of Industrial Organization*, 5(2):175–192, 1987.

[7] C. H. Brooks, S. Fay, R. Das, J. K. MacKie-Mason, J. O. Kephart, and E. H. Durfee. Automated search strategies in an electronic goods market: Learning and complex price scheduling. In *Proceedings of the First ACM Conference on Electronic Commerce*. ACM Press, November 1999.

[8] 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, 1998.

[9] P. Dasgupta and E. Maskin. The existence of equilibrium in discontinuous economic games, i: Theory. *Review of Economic Studies*, 53:1–26, 1986.

[10] C. d'Aspremont, J. J. Gabszewicz, and J.-F. Thisse. Product differences and prices. *Economics Letters*, 11:19–23, 1983.

[11] A. Denzau, A. Kats, and S. Slutsky. Multi-agent equilibrium with market shares and ranking objectives. *Social Choice and Welfare*, 2:95–117, 1985.

[12] P. Diamond. A model of price adjustment. *Economic Theory*, 3:156–168, 1971.

[13] B. C. Eaton and R. G. Lipsey. The principle of minimum differentiation reconsidered: Some new developments in the theory of spatial competition. *Review of Economic Studies*, 42:27–49, 1975.

[14] J. Farrell, H. K. Monroe, and G. Saloner. The vertical organization of industry: Systems competition versus component competition. *Journal of Economics and Management Strategy*, 7(2):143–182, 1998.

[15] S. A. Fay. Competition between firms that bundle information goods. Working Paper, November 1999.

[16] P. C. Fishburn, A. M. Odlyzko, and R. C. Siders. Fixed fee versus unit pricing for information goods: Competition, equilibria, and price wars. Working Paper, January 1997.

[17] D. Fudenberg and J. Tirole. The fat cat effect, the puppy-dog play, and the lean and hungry look. *American Economic Review*, 74(2):361–366, May 1984.

[18] J. J. Gabszewicz and J.-F. Thisse. Spatial competition and the location of firms. In J. J. Gabszewicz, J.-F. Thisse, M. Fujita, and U. Schweizer, editors, *Fundamentals of Pure and Applied Economics. Volume 5: Location Theory*. Harwood Academic Publishers, Chur, Switzerland, 1986.

[19] R. J. Gilbert and M. Lieberman. Investment and coordination in oligopolistic industries. *The Rand Journal of Economics*, 18:17–33, 1987.

[20] A. Greenwald and J. Kephart. Shopbots and pricebots. In *Proceedings of 16th International Joint Conference on Artificial Intelligence*, volume 1, pages 506–511, August 1999.

[21] A. R. Greenwald and J. O. Kephart. Probabilistic pricebots. Unpublished manuscript, available at www.research.ibm.com/infoecon, 2000.

[22] J. E. Hanson and J. O. Kephart. Spontaneous specialization in a free-market economy of agents. In *Proceedings of the Artificial Societies and Computational Markets Workshop (Agents '98)*, 1998.

[23] E. Hopkins. A note on best response dynamics. *Games and Economic Behavior*, 29:138–150, 1999.

[24] H. Hotelling. Stability in competition. *Economic Journal*, 39:41–57, 1929.

[25] J. O. Kephart, R. Das, and J. K. MacKie-Mason. Two-sided learning in an agent economy for

information bundles. In *Agent-mediated Electronic Commerce, Lecture Notes in Artificial Intelligence*, Berlin, 2000. Springer-Verlag.

[26] J. O. Kephart, J. E. Hanson, D. W. Levine, B. N. Grosof, J. Sairamesh, R. B. Segal, and S. R. White. Dynamics of an information filtering economy. In *Proceedings of the Second International Workshop on Cooperative Information Agents*, 1998.

[27] 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):1–23, 1998.

[28] A. Lerner and H. W. Singer. Some notes on duopoly and spatial competition. *Journal of Political Economy*, 45:145–186, 1973.

[29] X. Martinez-Giralt and D. J. Neven. Can price competition dominate market segmentation? *Journal of Industrial Economics*, 36:431–442, 1988.

[30] C. Matutes and P. Regibeau. Compatibility and bundling of complementary goods in a duopoly. *The Journal of Industrial Economics*, 40(1):37–53, 1992.

[31] R. P. McAfee, J. McMillan, and M. D. Whinston. Multiproduct monopoly, commodity bundling, and correlation of values. *The Quarterly Journal of Economics*, 114:371–383, 1989.

[32] D. Neven. On Hotelling's competition with non-uniform customer distributions. *Economics Letters*, 21:121–126, 1986.

[33] M. J. Osborne and C. Pitchik. The nature of equilibrium in a location model. *International Economic Review*, 27:223–237, 1986.

[34] M. J. Osborne and C. Pitchik. Equilibrium in Hotelling's model of spatial competition. *Econometrica*, 55:911–922, 1987.

[35] J. Sairamesh and J. O. Kephart. Price dynamics of vertically differentiated information markets. *Decision Support Systems*, 28:35–47, 2000.

[36] M. A. Salinger. A graphical analysis of bundling. *Journal of Business*, 68:85–98, 1995.

[37] S. Salop and J. Stiglitz. A theory of sales: A simple model of equilibrium price dispersion with identical agents. *American Economic Review*, 72(5):1121–1130, December 1982.

[38] S. C. Salop. Strategic entry deterrence. *American Economic Review*, 69(2):335–338, 1979.

[39] A. Shaked. Existence and computation of mixed strategy Nash equilibrium for 3-firm location problem. *Journal of Industrial Economics*, 31:93–96, 1982.

[40] A. Shaked. Existence and computation of mixed strategy nash equilibrium for 3-firms location problem. *Journal of Industrial Economics*, 31:93–96, 1982.

[41] Y. Shilony. Mixed pricing in oligopoly. *Journal of Economic Theory*, 14:373–388, 1977.

[42] A. M. Spence. Entry, capacity, investment and oligopolistic pricing. *Bell Journal of Economics*, 8:534–544, 1977.

[43] T. Tabuchi and J. F. Thisse. Asymmetric equilibria in spatial competition. *International Journal of Industrial Organization*, 13:213–227, 1995.

[44] H. R. Varian. A model of sales. *American Economic Review*, 70:651–659, 1980.

[45] H. von Stackelberg. *The Theory of the Market Economy.* Oxford University Press, New York, 1952. Translated from the German.

[46] L. L. Wilde and A. Schwartz. Comparison shopping as a simultaneous move game. *Economic Journal*, 102:562–569, 1992.