Bundling and Recommendation for Information Brokerage

D.J.A. Somefun[*] J.A. La Poutré[*,†]

[*]CWI, Centre for Mathematics and Computer Science

P.O. Box 94079, 1090 GB Amsterdam, The Netherlands.

[†]School of Technology Management, Eindhoven University of Technology

P.O. Box 513, 5600 MB Eindhoven, The Netherlands.

E-mail: {koye, hlp}@cwi.nl.

23rd September 2002

Abstract

In this paper, we discuss some of the consequences on-line dynamic bundling and/or pricing of (information) goods, and (automatic) recommender systems can have for information brokerage. We argue that dynamic bundling/pricing enhances especially the value extracting (or profit generating) capacity of an information broker. Recommender systems, on the other hand, enhance through, for example, customer lock-in especially the value generating capacity of an information broker. More traditional (automatic) recommender systems have a number of drawbacks. We outline how recommendation based on sales statistics can circumvent these difficulties. We discuss especially the advantages and challenges of integrating dynamic bundling/pricing into such recommender systems.

Keywords: value creation; value extraction; information brokerage; dynamic pricing; recommender systems

1 Introduction

An information broker gathers/buys information from various sources and sells it to his customers. Traditionally, due to savings in production and transaction costs the information is bundled and sold via subscriptions. Examples of this practice are newspapers, scientific journals, and subscriptions services such as Reuters' Stockmaster service. The coming of the Internet Economy changed the necessity of bundling. Due to the digitization of information goods— pushed by the existence of the Internet Economy— the costs of producing and selling an additional information good has become virtually zero. Consequently, it is currently economically feasible to deliver, for example, individual news items or individual stock quotes to consumers and have them pay per item. On the other hand, conditional on consumers' preferences, bundling various information goods might still be useful because it can facilitate the extraction of consumers' valuation (see Section 2.1).

In this changed environment an information broker has to reconsider his business strategy. More specifically, he has to reconsider the way he bundles and positions his information goods. As a consequence of the very low production and transaction costs selling, for example, only individual items, larger bundles, or a few smaller bundles aimed at particular market niches can all be profitable strategies. What the best strategy is depends on consumers' actual preferences.

Current developments in the machine learning community provide an extra incentive for information broker to reconsider his business strategy. Typically, a seller of information goods does not have complete information about consumers' preferences. By employing machine learning techniques, a seller can learn more about his customers. It is even possible to apply some of these techniques on-line, i.e., the obtained knowledge is directly and automatically put to use. Two important application areas for these on-line techniques are (on-line) dynamic bundling and/or pricing of (information) goods and (automatic) recommender systems.

Dynamic bundling/pricing entails automating (and therefore potentially speeding up) the search process of finding— via trial and error— more profitable bundling/pricing schemes.[1] Automatic recommender systems are machine learning systems specialized to recommend products in (electronic) commerce applications (Schafer et al., 2001). In this paper, we discuss some of the consequence dynamic bundling/pricing of (information) goods and (automatic) recommender systems can have for information brokerage. Moreover, based on this discussion we argue that recommender systems which use sale statistics are particularly promising for information brokerage; especially if the systems integrate dynamic bundling/pricing into the recommender system.

In this paper, we make the observation that dynamic bundling/pricing enhances especially the value extracting (or profit generating) capacity of an information broker, whereas recommender systems enhance the value creating potential of an information broker. Within a competitive setting the important drawback of dynamic bundling/pricing is that it can increase the competition with rivaling firms. Recommender systems, on the other hand, can reduce competition by generating value through customer lock-in. An important disadvantage of more traditional (automatic) recommender systems is, however, that they require significant interaction with customer which is rather time consuming for the customers and costly in terms of (human) resources. Moreover, what customers say about their preferences might not coincide with their actual behavior. In this paper, we therefore outline how recommendation based on sales statistics can circumvent these difficulties. An additional advantage of recommendation based on sales statistics is that it does not necessarily require information about individual customers. Given the growing concern about

[1] Due to the virtually zero costs of changing the price of information goods, it is economically feasible for an information broker to frequently update the bundling/pricing scheme, therefore he can actually apply dynamic bundling/pricing.

customers' privacy this can be of great practical importance.

Like recommendation based on sales statistics, on-line dynamic bundling/pricing makes use of sales statistics to learn either implicitly or explicitly more about customers' preferences. Therefore, it can be beneficial to make what is learned exchangeable by integrating dynamic bundling/pricing into automatic recommender systems. We discuss the advantage and difficulties of such an integrated recommender system. Moreover, we briefly sketch a possible architecture for such an integrated recommender system.

The remainder of this paper is organized as follows. In Section 2, we discuss when bundling is advantageous and how (with the use of machine learning techniques) the seller of information goods can apply bundling. In Section 3, we use the framework of Amit and Zott (2001) to analyze the value creating potential of information brokerage, especially the situation of information brokerage under imperfect competition. In Section 4, we discuss integrating dynamic bundling/pricing into recommender systems that recommend based on sales statistics. Conclusions follow in Section 5.

2 Bundling of Information Goods

2.1 When to Bundle

Bundling is the practice of combining two or more items together and selling them as one product. There are various bundling strategies possible. Pure bundling refers to the practice of either offering consumer the complete bundle or nothing at all. With pure unbundling, no bundle is offered, consumers can obtain the desired bundle of goods by buying the individual items. An additional option is mixed bundling, where consumer can either buy the whole bundle or the desired individual items. Besides these three bundling strategies there are obviously a whole host of other bundling strategies possible.[2]

Supply side considerations, demand side considerations and strategic considerations are roughly the three reasons for a seller to apply bundling.[3] On the supply side, bundling can result in savings in production and transaction costs. The custom of bundling hard copies of (scientific) articles into a journal is an example where bundling, instead of selling the individual items, means a significant reduction in the transaction and production costs. An additional supply side reason for bundling is complementarities among the bundled components (e.g., the bundling of scientific articles on a particular topic).

On the demand side, bundling can be an important instrument for extracting consumers' valuations. For example, a film distributor could sell to movie theaters a block of two films (e.g., a popular and less popular movie). This strategy might result in higher profits than selling these movies separately. The effective-

[2] We slightly abuse the definition of bundling by henceforth both interpreting pure bundling and not offering any goods at all as bundling options.

[3] See Mankila (1999) for an in debt discussion of the reasons for price bundling.

ness of bundling as a mean to extract consumers' valuations depends, among other things, on consumers' valuations for the offered goods. The bundling of two goods could, for example, increase seller's profit when consumers' valuations for the two goods is negatively correlated (Stigler, 1963; Schamlensee, 1984).

In imperfectly competitive markets (as opposed to monopolistic markets) there is also a strategic incentive for bundling. A monopolist in one market could, for instance, exclude competitors by bundling the monopoly product with a good that is sold in a competitive market. An infamous example of this practice is Microsoft's tied selling of their operating system and internet browser.

Traditionally, because of cost saving most information goods are sold in bundles. Examples of this practice are newspapers, scientific journals and subscriptions to services such as Reuters' Stockmaster service. The coming of the internet and the low digitization costs of information goods, however, make cost saving less of an issue. As a consequence, it is currently economically feasible to deliver, for example, individual news items or individual stock quotes to costumers and have them pay per item. Bundling of information goods might, nevertheless, still be a profitable strategy because of demand side issues and strategic reasons.

Besides making costs saving less of an issue the digitization of information goods also reduces the marginal costs of an information good virtually to zero. In a series of papers, Bakos and Brynjolfsson argue that this property of almost zero marginal costs could make it again attractive to apply pure bundling instead of pure unbundling for the sales of information goods. In Bakos and Brynjolfsson (1999a, 1999b) they consider the case of a monopolistic seller and in Bakos and Brynjolfsson (2000) they consider the case of a seller operating under imperfect competition.

The work of Bakos and Brynjolfsson (1999b) provides a good insight in why a seller can extract consumers' valuations via (pure) bundling. They develop a probabilistic framework that enables the analysis of bundling for very large numbers of (information) goods. The underlying assumption is that consumers' valuations for a bundle are independently and identically distributed. They show that with virtual zero marginal costs and for the most common demand curves it is advantageous to offer a very large bundle of goods instead of selling all the items separately. The intuition underlying their result is that as the number of information goods increases, the law of large numbers ensures that an increasing number of consumers will have a value for the offered bundle closer to the mean of the underlying distribution. Hence, an increasing number of consumers are willing to pay a price close to the mean value of the offered bundle. For "reasonable shapes" of the demand curve this implies a higher profit for the seller.

An additional option, besides pure bundling and pure unbundling, is mixed bundling. Mixed bundling can be more profitable in the case of more heterogeneity among consumers' valuations for the various combinations of the individual items (Chuang & Sirbu, 1999; Bakos & Brynjolfsson, 1999b). Note that besides these three options a host of other bundling combinations are possible. In fact, a seller of n (information) goods can in total offer 2^n different bundle combinations (including not offering any bundle at all). That is, in the most general case the seller considers a pricing schedule with 2^n variables.[4] Even with complete information about consumers valuation the problem of finding the optimal bundle prices becomes intractable for large values of n.[5] Thus for all practical purposes it make sense, for the seller, to only consider a very restricted number of bundling options.

2.2 Applying Bundling

The traditional reason for bundling information goods is cost saving. In the previous section, we argued that due to the digitization of information goods cost saving is currently less of an issue. Besides strategic considerations, bundling of information goods can be beneficial because it can facilitates the extraction of consumers' valuations. From a cost perspective it has become economically feasible for a seller of n information goods to offer 2^n different bundle combinations. For all practical purposes, however, it only make sense for the seller to consider a restricted number of bundling options. The seller might, for example, distinguish between bundles based on the number of items in a bundle and not on the identity of the items in the bundle. The most general pricing schedule for this approach has only n variables (instead of 2^n variables).

Typically, the seller has only (very) limited information about the distribution underlying consumers' valuations. Thus, to actual implement bundling requires that the seller refines an initial estimate of the relationship between prices and profits via trial and error. To search for the best bundle/price offers requires a trade-off between exploration and exploitation. To learn the optimal bundle/price offers for a complex pricing scheme (e.g., a pricing schedule with n variables) takes time; time during which the seller cannot capitalize on what he has already learned.

Brooks and Durfee (2000) conjecture that in many cases sellers of information goods are better off specializing by offering a limited number of different bundles that appeal to a particular niche of the population. Offering a large number of different bundles will require a lot of exploration to determine how all these bundles are priced. Brooks et al. (1999) use computer experiments to analyze the trade-off between exploration and exploitation more explicitly. They consider the case where the monopolistic seller distinguishes between bundles based on the number of items in a bundle (and not on the identity of the items in the bundle).[6]

They study the one-variable pricing schemes of pure bundling and linear pricing (i.e., consumers pay the same price for each item they choose to receive). Both bundling and linear pricing are one-variable pricing schemes because either the price of the bundle or the price of buying an additional item has to be determined. Additionally, they study the two variable pricing schemes of mixed

[4] A consumer pays the price $p(i)$ for the i^{th} bundle with $0 \leq i < 2^n$.

[5] Hanson and Martin (1990) manage to solve the bundling price problem up to $n = 21$, by formulating the problem as a mixed integer linear programming problem.

[6] Note that the most general pricing schedule for this problem has n variables.

bundling and two-part tariff pricing (i.e., consumers pay a subscription fee plus a price per item). Mixed bundling requires a two variable pricing scheme because both the bundle price and the price of buying an individual item have to be determined. Similarly, two-part tariff pricing requires a two variable pricing scheme because the subscription fee and the price of buying an additional item have to be determined. The last pricing scheme they consider is the n variable pricing schedule, which is the most general pricing strategy. The seller chooses a price for each possible quantity without restriction, i.e., the pricing scheme is nonlinear.

The conducted computer experiments show for two different learning mechanism, $n = 10$ or $n = 100$, and 1000 iterations that most of the improvement in the profit comes when the seller uses a two variable pricing scheme instead of an one variable pricing scheme. Kephart et al. (2001) use the same framework to show that if relatively frequent and unpredictable demand shocks occur then less complex pricing schemes such as two-part tariff and mixed bundling are more profitable than nonlinear pricing.

2.3 On-line Dynamic Pricing and Bundling

Since typically a seller of information goods has (very) limited information about the distribution underlying customers' valuation, applying bundling implies learning, in most cases. Sellers might automate aspects of this learning process by using machine learning techniques to develop on-line dynamic bundling and/or pricing algorithms. That is, the process of refining via trial and error the used pricing scheme is conducted by a "learning" algorithm. Moreover the most recently obtained information is automatically incorporated in the pricing schedule. The use of an on-line dynamic bundling/pricing algorithm makes frequent adjustments of the price (at little additional costs) possible. Consequently, past experience is incorporated into the price more frequently. Hence, more exploitation is possible. More exploitation might in particular be important in an environment with relatively frequent and unpredictable demand shocks. Moreover, in a setting of imperfect competitive, the need to apply dynamic pricing might simple occur because the competition is using dynamic pricing algorithms.

The use of dynamic bundling/pricing algorithms might increase the competition between sellers of information goods. Greenwald and Kephart (1999), for example, show how the use of various simple dynamic pricing algorithms can lead to price wars in a competitive setting.[7] For a similar set up, but different dynamic pricing algorithms, prices could converge to an equilibrium situation (Greenwald & Kephart, 2001). Thus the impact of selling bundles via dynamic pricing algorithms depends on the dynamic pricing algorithm the various sellers use.

Sellers of information goods could prevent price wars, or other negative influence caused by fierce competition, by finding a market niche. They might

[7]Dasgupta and Das (2000) obtain similar results for slightly different dynamic pricing algorithms.

opt for offering sub-bundles of information goods that do not coincide with the bundles offered by the competitors. Whether or not the search for market niches enables sellers of information goods to reduce competition depends, among other things, on the heterogeneity of consumers and how much is known about the profit landscape (cf. Hanson & Kephart, 1998; Kephart et al., 1998; Kephart & Fay, 2000; Gazzale & MacKie-Mason, 2001).

So far we have focused on how bundling and the use of dynamic pricing algorithms can help the seller to extract consumers' valuations. In the next section we discuss how a seller operating under imperfect competition can create consumers' valuation. A benefit of realizing value creation is that it provides an information broker with additional means— besides finding a market niche— to reduce competition.

3 Value Creation

3.1 Four Value Drivers

Extrapolating current developments on the internet might lead to predicting the coming of an economists' utopia of frictionless electronic trade, i.e., low search costs, strong price competition, low margins, low deadweight loss.[8] Clearly, a frictionless market is no utopia for firms. They can bring friction back to the market place by competing on more than just price characteristics. More specifically, firms should focus on all aspects of the value creation potential of electronic business.

Amit and Zott analyze the value drivers in electronic business (Amit & Zott, 2001). The four key value drivers they determine are (transaction) efficiency, complementarities, lock-in and novelty.

- Transaction efficiency increases when the costs per transaction decreases (where Amit and Zott define "cost" broadly). Electronic trade, for example, increases efficiency by reducing the information asymmetry between buyers and sellers. Low search costs on the internet reduce asymmetric information by making it relatively easy to compare various sellers (e.g., the use of shopbots on the internet, cf. Bakos & Brynjolfsson, 2000).

- Goods are complementary whenever bundling them together generates more value than the total value of having each of the goods separately. Amit and Zott give the example of e-bookers. This on-line travel organization gives its customers access to weather information, currency exchange rate information, and appointments with immunization clinics. Amit and Zott argue that these services enhance the value of the core product.

- Lock-in refers to the ability to persuade customers to engage in repeated transacions. The two main components of lock-in are switching costs

[8] Deadweight loss measures the amount by which consumers are made worse off by paying a price above the marginal costs.

and positive network externalities.[9] On-line vendors can, for instance, introduce switching costs by customizing their web site with the use of data mining methods. Network externalities occur in the context of electronic business when the value created for customers increases with the size of the customer base (i.e., the number of customers increases).

- Novelty involves the introduction of new products, processes, or services on the internet.

The first three value drivers provide a good framework for analyzing how an information broker can generate value. Clearly, novelty is also important. The actual creation of value through novelty is, however, too case specific for the purpose of this paper.

The bundling and dynamic pricing of information goods mainly aims at value extraction instead of value creation. Bundling and dynamic pricing could, however, also result in value creation, especially through (transaction) efficiency and complementarities. For instance, more frequently updating the bundling/pricing scheme with dynamic bundling/pricing algorithms could result in offers more tailored to the individual needs of the consumers at very little additional cost. Obtaining the same level of adaptation to current consumer demand without the use of dynamic bundling/pricing algorithms involves much higher costs of monitoring demand developments. For example, in the absence of dynamic bundling/pricing actual people have to do the data analysis and (partly) based on this analysis have to determine the next bundling/pricing scheme. To obtain the same monitoring level they have to do this with the same frequency the dynamic pricing algorithm recomputes the pricing schedule. Thus, updating the bundling/pricing scheme with dynamic pricing algorithms can result in transaction efficiency. Moreover, whenever the offered bundles contain complementary goods, bundling might also result in complementarities.

Within the context of information brokerage, another probably more important way of creating value (than bundling/pricing) is to recommend information to customers.[10] We will show that lock-in is one channel through which recommending of information generates value. Lock-in has the advantage of reducing the competition with rivals. Hence, by recommending, an information broker obtains additional means—besides finding a market niche— to reduce competition.

Recommending information to customers could eventually reduce their search costs for obtaining the desired information. These search costs could be significant if, for example, the "market value" of the information is low, the information is hard to categorize in advance, and/or at the same time a lot of seemingly similar information is being offered for sale. We can distinguish between passive and active recommendation. With passive recommendation customers are

[9]Network externalities occurs whenever the benefit, or surplus, that an agent derives from consuming a good depends on the number of other agents consuming the same kind of good.

[10]In Section 4 we discuss how bundling/pricing of information could influence the value creation capacity of an information broker indirectly by facilitating the recommendation of information.

given tools that facilitates the search for finding the desired information. Passive recommendation can contain aspects of all three value drivers (i.e., efficiency, complementarities and lock-in due to switching costs). Passive recommendation can, for example, entail selling information goods through an interface that is user friendly, customized, and that grants the use of a good search engine. Consequently, consumers can enjoy lower search costs. Lower search costs mean that transactions are done more efficiently. Moreover, offering customers the use of a search engine can complement the consumption of certain information goods. Finally, a customized user interface makes it less appealing for customers to switch between sellers, i.e., it increases switching costs.

With active recommendation the information broker actually suggests information to his customers. By utilizing a recommender system the information broker can realize active recommendation. Recommender systems use product/consumer knowledge to select products from the vendor's database that correspondent with the interest of a given consumer. The product/consumer knowledge is either hand coded knowledge supplied by experts or "mined" knowledge learned from the behavior of consumer. In case of mined knowledge we speak of automatic recommender systems.[11] For an information broker, in particular automatic recommender systems are interesting. Automatic recommender systems are machine learning systems specialized to recommend products in (electronic) commerce applications. They allow for a high level of personalized recommendation at little cost relative to the value of the offered information goods.

Like passive recommendation active recommendation can contain aspects of all three value drivers (i.e., efficiency, complementarities and lock-in due to switching costs). However, for active recommendation lock-in effects might be more significant because active recommendation can also generate network externalities. In the next subsection we will discuss— within the context of an information broker— automatic recommender systems and how they can result in value creation, especially network externalities. (For brevity we will henceforth drop the adjective automatic.)

3.2 Recommender Systems

The heart of an information brokerage recommender system should be an information filtering system. An information filtering system divides a large-volume data stream into substreams. The criteria according to which the division takes place is based on a profile.[12] The profile can be created and updated by directly interacting with the customers (e.g., questionnaires). Additionally, the profile can be obtained and updated by automatic adaptation via implicit feedback. If, for example, a profile coincides with a customer then the information filtering systems could select the information items that match the interest of the given customer best, based on the customer's past purchasing behavior.

[11] Cf. Schafer et al. (2001) for a general discussion of recommender systems in e-commerce

[12] See (Kutschinski & Poutré, 2001) for a discussion of profiling techniques.

A profile does not necessarily have to coincide with a customer/user. A group of customers could also make up a valid profile given that there is some way of (roughly) distinguishing the various customers' categories. For an information brokerage system profiles could, for instance, be derived from certain information categories (e.g., news items on Computer technology, E-commerce, sports, politics, etc.). Based on the profile type (one user versus multiple users) we can roughly distinguish between three types of information brokerage (recommender) systems.[13]

1. *Collective system*. The system keeps track of so called stereotypes. It composes and updates the stereotypes based on *anonymous* data, i.e., data from which it is not possible to derive individual user history (e.g., aggregate sales statistics). A stereotype roughly coincides with the behavior and preferences of certain groups of customers. The system uses the developed stereotypes as a basis for its recommendations. If, for example a customer buys a particular news item then the system could recommend other news items to the customer that belong to the same stereotype (as the bought news item).

2. *The personalized system*. The system keeps track of (individual) user profiles. It composes and updates a profile based on the past behavior (and provided feedback) of the user. By utilizing a user profile the system keeps track of the long-term behavior and preferences of individual customers. Based on the developed user profile it is possible for such a system to recommend. If, for example, news arrives that closely matches the interest of a customer (defined by the user profile) then the system could notify the customer.

3. *The hybrid system.* This system combines (1) and (2). The method of collaborative (or social) filtering underlies the hybrid system. It develops and keeps track of both individual user profiles and stereotypes. Based on their user profile users are assigned to one or more stereotypes. Due to the combination of user profiles and stereotypes the system can recommend a particular news item to a customers whenever other consumers that belong to the same stereotype already showed an interest in the news item.

With the collective system it is not possible to trace consumers directly back to a particular profile. It tries to recommend based on domain specific knowledge. A simple application is the use of aggregate sales data. For instance, if a particular piece of information is in high demand relative to other information in that same category then the system could recommend this piece of information to customers who express an interest in the same information category. (See Section 4 for a more detailed discussion of recommending based on sales statistics.) In this paper we focus on the use of sales data as the domain specific knowledge therefore we will call the collective system henceforth the *market oriented* system.

[13]For brevity we will henceforth drop the adjective recommender.

3.2.1 Value Creation Capacity

By reducing the search cost of customers, all three types of information brokerage systems create value via the value driver (transaction) efficiency. Additionally, all three generate value through (customer) lock-in. They, however, differ in the way they create customer lock-in. The hybrid system creates customer lock-in through network externalities and switching costs whereas the other two only create customer lock-in through either customer lock-in or switching costs.

- (*Network Externalities*) Both the hybrid and the market oriented system create customer lock-in through network externalities. The larger the customer base the more accurate customers' stereotypes are developed. Hence, the better the hybrid system can recommend. Similarly, the larger the customer base the more accurate the market oriented system can recommend.
- (*Switching Costs*) The personalized and hybrid system create customer lock-in through switching costs. They have the capacity to learn more and more about the individual users of the system. Consequently, it becomes more expensive for a customer to switch to a competitor because the competitor has to relearn the customer's preferences. Clearly, a market oriented system can also (indirectly) generate switching cost due to more accurate recommendation to *groups* of customers. The direct cause of these switching costs, however, has more to do with network externalities.

The hybrid system generates value through the most diverse number of value drivers. More specifically, the hybrid system seems particularly promising for reducing competition because it generates both switching costs and network externalities. Hence, making competitors less attractive for customers due to switching costs and less likely to achieve the same level of search cost reduction due to network externalities.

3.2.2 Privacy Issues

From the consumers perspective the hybrid system has the important drawback that it invades customers' privacy, moreover it uses this private information to advice other customers. In the extreme case, the system could implicitly use a particular customer's expertise of filtering out the right information to advice others. In the case of, for example, a financial broker, a particular customer, say an investor, might consider the type of information he consults propriety information. Most likely he will not be willing to share this with other customers of the information brokerage system. The personalized system also has customers' privacy as an important drawback. It does, however, not use this private information to advice others. The main advantage of the market oriented system is that it does not invade customers' privacy at all.

3.2.3 Concluding Remark

Thus from the seller's perspective the hybrid systems appears most promising for generating value. Especially, because it seems most suited for reducing competition. To anticipate privacy issues an information broker might, nevertheless, consider implementing a recommender system that allows customers' participation at three levels of anonymity (which may or may not be hierarchical). First level participation only involves recommendations based on a market oriented system. Second level participation (also) involves the use of a personalized system and third level participation (also) involves the use of a hybrid system.

4 Integrating Bundling and Recommending

In the previous two sections we have argued that the use of machine learning techniques can facilitate both the extraction and creation of consumers' valuations. More specifically, with the use of machine learning techniques an information broker can on-line learn which bundling/pricing scheme results in a higher profit given consumers' current preferences. Moreover, machine learning techniques also enable the information broker to apply active and automatic recommendation. Thus, so far we have treated recommending and the dynamic bundling/pricing of goods as two separate approaches. In this section we discuss recommendation based on sales statistics. Especially for this type of recommender systems integrating the dynamic bundling/pricing of goods into (active and automatic) recommendation can be beneficial.

4.1 Recommendation based on Sales Statistics

4.1.1 Sales Statistics and Data Mining

Simply put, the idea of recommendation based on sales statistics is that customers' buying behavior could reveal their preferences. In Section 3.2 we discussed three types of recommender systems: the market oriented (or collective) system; the personalized system; and the hybrid system. Within the context of the market oriented system we already briefly mentioned recommendation based on sales statistics. All three systems can make use of sales statistics. The important distinction between the market oriented system on the one hand and the personalized and hybrid system on the other hand, however, is that the former only makes use of anonymous sales data (i.e., data that cannot directly be linked to individual customers).

A great advantage of recommendation based on sales statistics is that it does not require customers to instruct the system about their preferences. Initialization of a recommender system that requires significant interaction with customers is rather time consuming for the customers and costly in terms of (human) resources. Moreover, what customers say about their preferences might not coincide with their actual behavior. From a privacy perspective, another additional advantage is that recommendation based on sales statistics does not necessarily require the initializing and updating of the profiles of individual customers.

Based on sales statistics it is, for example, possible by using data mining techniques to identify various product groups which roughly coincides with customers' categories. Whenever a customer buys an item belonging to a particular product group (or in some other observable way shows an interest in a particular product group) a recommender system can recommend other goods belonging to that same product group. Moreover, the recommender system could use the fact that customers are buying or not buying the recommended items as a feedback. This type of recommender system does not (necessarily) require the storage of customers' profiles. Clearly, on top of this type of recommender system it is also possible to do recommendations that require the use of customers' profiles. A recommender system could, for instance, by default recommend newly released news items to a customer whenever based on customer's profile these items belong to an "interesting" product group.

4.1.2 Dynamic Pricing and Bundling

For a seller of information goods it might be advantageous to integrate the dynamic bundling/pricing of goods into (active and automatic) recommendation. Especially the updating of customers' profiles can become more effective whenever the information broker explicitly offers various product bundles for sale and frequently experiments with the composition and price of these bundles. By experimenting the seller can speed up the process of learning the preferences of individual customers and/or groups of customers. Hence, the seller can improve the quality of the recommendations. Moreover, the fact that bundles of information goods are relatively low valued makes (a lot) experimentation possible without the seller running a great risk of reducing his total revenue significantly.

An additional advantage is that by buying a bundle which captures some of a customer's preferences the customer is saved the trouble of searching and paying for all the individual items that constitute the bundle. Thus, whenever the offered bundles closely match various customers' categories they could reduce the search and transaction costs for customers belonging to these categories. Finally, note that integrating the dynamic bundling/pricing of goods into (active and automatic) recommendation might not only increase the effectiveness of the recommenders system in creating valuation for the customers, it might also facilitate the search for appropriate product bundles.

4.1.3 Difficulties and Attention Points

There are a number of difficulties with integrating the dynamic bundling/pricing of goods into a recommender systems. Without any pretence of being complete we want to mention three of these difficulties.

1. Customers might be distrustful towards using such an integrated recommender system because it might use whatever it learns about them to

pursue higher profits. To obtain the products they are interested in customers might, consequently, end up paying a higher price. This complaint is particularly true for the hybrid recommender system and (to a lesser degree true for) the personalized recommender system because both explicitly store the profiles of individual customers. In order to convince customers of using such an integrated recommender system (especially whenever it is based on the hybrid or personalized system) requires that the use of such a system is also beneficial to them and that it remains beneficial to them.

2. Sales statistics of the offered bundles are not necessarily directly telling the recommender system something about customers' interests in a particular product group. Customers could, for example, have chosen a different product group if other options where available. Moreover, a change in the revenue of various product bundles might be caused by a change in the pricing scheme and might have little to do with customers being really interested in a particular bundle. The challenge is to develop an integrated recommender system that by striking the right balance between exploitation and exploration in the usage of the various pricing schemes can, nevertheless, learn to identify and link various relatively homogeneous consumer groups to the corresponding product bundles or groups.

3. An integrated recommender system could be biased towards focusing on the larger customers' categories because they generate higher revenue than customers' categories with fewer members. Consequently, the system might not be so good at recommending customers belonging to smaller consumer categories. Aggregation these smaller consumer categories might, nevertheless, result in an important group of customers. A possible way to avoid such a bias is by not having a strict relation between customers' stereotypes and product bundles. In addition to sales statistics the system could then use other relevant and available information to update these stereotypes.

Successfully integrating the dynamic bundling/pricing of goods into (active and automatic) recommendation based on sales statistics is not an obvious and easy task. We think that it is, nevertheless, promising enough to at least warrant further research. To further materialize the idea of an integrated recommender system we give a brief sketch of a possible architecture for such a system in the next subsection. With this more specific example at hand we will also discuss how an integrated system can avoid some of the difficulties mentioned above.

4.2 An Integrated Recommender System

A way to develop an integrated recommender system is to have both the feedback received from recommendations and dynamic bundling/pricing of goods contribute to the updating and maintaining of customers' stereotypes/product bundles. One option is to identify a collection of subbundles with customers' stereotypes. This collection of subbundles represent the currently conjectured building blocks for the actual product bundles. The actual bundles are obtained by combining one or more customers' stereotypes (or subbundles). Whenever customers buy a particular news item (or subbundle) the integrate systems could recommend other items in that same product bundle. The system can update customers' stereotypes and offered product bundles based on the feedback of customers buying or not buying the recommended items. Moreover, it could also use the sales statistics of the currently available bundles to update customers' stereotypes and the offered bundles. In figure 1 we draw the basic flowchart of the integrate recommender system. In the remainder of this subsection we will use this example to discuss the idea of integrating dynamic pricing and automatic recommending in more detail.

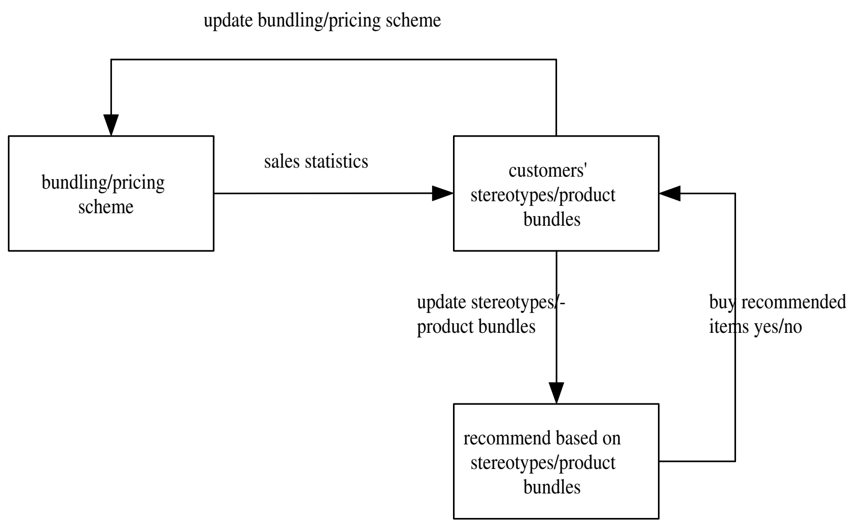

Figure 1: An example of a system that integrates dynamic bundling/pricing of goods into (active and automatic) recommendation based on sales statistics.

Initiating such an integrated recommender system can be particular difficult because there are so many bundling options (with n information goods 2^n different bundles are possible), moreover very little is know about the various consumers' stereotypes. A reasonable approach could be to start off with a limited number of stereotypes, which— based on the already available knowledge of consumers' preferences— are expected to result in a reasonably profitable partition of the potential customers. The offered bundles could at least initially coincide with these stereotypes. Due to the provided feedback, gradually, a richer and more diverse collection of customers' stereotypes can be developed. Consequently, also more different product bundles become possible. Furthermore, by also gradually increasing the number of different bundles customers, can get use to the richer variety of product bundles.

4.2.1 (Dis)advantages of Integrated System

In Section 3.2 we distinguished between three types of recommender systems (i.e., the market oriented, personalized and hybrid system). All three can be used as the underlying system for the integrated recommender system. The drawback of the market oriented and the personalized system is, however, that they will only register a positive feedback of customers buying recommended items if customers buy the items almost directly after they have has been recommended to them. Whenever customers buy the recommended items later (e.g., during a different session) the two systems cannot correlate the sales of the recommended items to the sales of the item (or items) that lead to the recommendation. By creating and updating customers' individual profiles both the hybrid system and the personalized system can keep track of all the news items customers buy. However, unlike the personalized system the hybrid system uses this information to update customers' stereotypes. Therefore, it can (unlike the other two systems) also correlate the sales of a recommended item whenever customers buy the recommended item at some later point in time.

From the customers' perspective using a hybrid system has the important drawback that customers' (individual) profiles are being used to determine the pricing/bundling scheme. In the worst-case, participating in such a recommender system could result in customers paying more for news items than they otherwise would have payed. To a lesser degree customers can have the same complaint about integrating bundling/pricing and recommending in general. To successfully integrate bundling/pricing and recommending requires that the information broker earns customers trust. In the case of a market oriented and personalized system it might suffices if— maybe via a trusted third party intermediating— customers can actually check if their profiles are being stored and if so (which is the case for the personalized and hybrid system) how this information is put to use. An additional way of earning customers trust is by trying to exclude changes in the bundling/pricing scheme which are solely based on strategic consideration (e.g., excluding tied selling). The information broker can do this, for instance, by a priori excluding certain pricing schemes (or making these pricing schemes less likely to occur).

The sales statistics of the currently available bundles and the customers buying or not buying a recommended item provide (possibly among other things) the feedback for the integrate recommender system. Based on this feedback the system updates customers profiles and the bundling/pricing scheme. Caution is, however, advisable when interpreting the feedback. A relatively low revenue of a particular bundle might, for example, indicate replacing or removing this bundle because customers are not interested in this group of products. On the other hand, the low revenue might simply be caused by a bundle price which is too high or too low.

The problem with interpreting feedback of an integrated system is that the system tries to learn both customers stereotypes and the best bundling/pricing scheme. Thus, in order to learn something about customers preferences based on sales statistics requires that the bundling/pricing scheme is not too complex and is not updated too frequently.

5 Concluding Remarks

Due to the digitization of information goods, it is currently feasible from a cost perspective to deliver and have customers pay per item. Bundling information goods might, however, be a good selling strategy because it can facilitate the extraction of consumers' valuations. The effectiveness of bundling as a mean to extract consumers' valuations depends, among other things, on consumers' valuations for the offered goods. With the use of on-line machine learning techniques an information broker can (directly and automatically) learn which bundling/pricing schemes results in the highest profit given consumers' current preferences. Due to the use of these techniques frequent adjustments of the bundling/pricing schemes are possible (at little additional costs). Consequently, they can enhance the value extracting capacity of an information broker.

The downside of exclusively using machine learning techniques for bundling/-pricing is that it can result in fierce price competition with rivaling firms. In order to reduce competition an information broker should also pay attention to (other aspects of) the value creating potential of information brokerage. We argue that for an information broker automatic recommender systems can particular be useful for generating value. One channel through which recommender systems generate value is by reducing the search costs for customers. Moreover, automatic recommender systems have the advantage of allowing for a high level of personalized recommendation at little cost relatively to the value of the offered information.

We distinguish between three types of systems that could underly (automatic) recommender systems for information brokerage. The third system, the hybrid system, seems most promising for generating value. Especially, because it has the most potential for customer lock-in through switching costs and network externalities. From the customers' perspective the hybrid system has the important drawback that it invades customers' privacy, and uses this private information to advice others. For some business application this could be a serious drawback. The second system, the personalized system, invades customers' privacy to a lesser degree. Moreover, the first system, the market oriented system, does not invade customers' privacy at all. Therefore, to anticipate privacy issues, an information broker might consider implementing all three systems, hence letting customers choose the preferred level of anonymity.

A promising implementation for any of the above three types of (automatic) recommender system is to recommend bundles based on sales statistics. A great advantage of recommendation based on sales statistics is that it does not require customers to instruct the system about their preferences. Initialization of a recommender system that requires significant interaction with customers is rather costly. Moreover, what customers say about their preferences might not coincide with their actual behavior. From a privacy perspective another additional advantage is that recommendation based on sales statistics does not necessarily require the initializing and updating of the profiles of individual customers (i.e., customers could remain anonymous) .

To enhance the accuracy of such a recommender system it might be a good idea for future research to integrate dynamic bundling/pricing of goods into recommendation based on sales statistics. The advantage of such an approach is that by experimenting with various product bundles an information broker can speed up the process of learning the preferences of individual customers and/or groups of customers. The aim of such an integrated recommender system should, however, not be to optimizing both the bundling/pricing scheme and the quality of the recommendations. Instead, dynamic pricing/bundling should be integrated into a recommender system so that customers perceive it as a service provided by the information broker. It allows customers to buy the desired items at once and (most likely) at a lower cost than buying all the individual items. The more the recommender system knows about the preferences of the customer the higher the quality of this service. Thus integrating pricing/bundling into a recommender system could potentially further enhance the value creating capacity of especially recommender systems based on sales statistics.

To conclude, in this paper we discuss some of the consequences the digitization of information goods and especially the use of machine learning techniques can have for information brokerage. We argue that these developments have both an impact on the value extraction and the value creation aspect of information brokerage.

References

Amit, R., & Zott, C. (2001). Value creation in ebusiness. *Strategic Management Journal*, *22*, 493–520.

Bakos, Y., & Brynjolfsson, E. (1999a). Aggregration and disaggregation of information goods: Implications for bundling, cite licensing and mircopayment systems. In B. Kahin & H. Varian (Eds.), *Internet publishing and beyond: The economics of digital information and intellectual property.* MIT Press.

Bakos, Y., & Brynjolfsson, E. (1999b). Bundling information goods: Pricing, profits and efficiency. *Management Science*, *45*(12).

Bakos, Y., & Brynjolfsson, E. (2000). Bundling and competition on the internet: Aggregation strategies for information goods. *Marketing Science*, *19*(1), 63–82.

Brooks, C. H., & Durfee, E. H. (2000). Toward automated pricing and bundeling of information. In *Knowledge based electronic markets '00.* (http://www.igec.umbc.edu/kbem/final)

Brooks, C. H., Fay, S., Das, R., MacKie-Mason, J. K., Kephart, J. O., & Durfee, E. H. (1999). Automated strategy searches in an electronic goods market:

Learning complex price schedules. In *Proceedings of the amc conference on electronic commerece* (pp. 31–41). ACM Press.

Chuang, J. C.-I., & Sirbu, M. A. (1999). Optimal bundling strategy for digital information goods: Network delivery of articles and subscridtions. *Information Economics and Policy*, *11*(1), 147–176.

Dasgupta, P., & Das, S. (2000). Dynamic pricing with limited competitor information in a multi-agent economy. In O. Eztion & P. Scheuermann (Eds.), *Coorperative information systems: 7th international conference* (Vol. 1906, pp. 291–310). Eilat, Israel: Springer.

Gazzale, R. S., & MacKie-Mason, J. (2001). Endogenous differentiation of information goods under uncertainty. In *Proceedings of the 3rd acm conference on electronic commerce* (pp. 95–104). ACM Press.

Greenwald, A., & Kephart, J. O. (2001). Probabilistic pricebots. In *Fifth international conference on autonomous agents.*

Greenwald, A. R., & Kephart, J. O. (1999). Shopbots and pricebots. In *Proceedings of sixteenth international joint conference on artificial intelligence* (Vol. 1). Stockholm.

Hanson, J. E., & Kephart, J. O. (1998). Spontaneous specialization in a free–market economy of agents. In *Agents-98 workshop on artificial societies and computational markets.* Minneapolis/St. Paul. (http://www.cs.toronto.edu/ grigoris/ascma.html)

Hanson, W., & Martin, R. K. (1990). Optimul bundle pricing. *Management Science*, *36*(8), 155–174.

Kephart, J. O., Brooks, C. H., & Das, R. (2001). Pricing information bundles in a dynamic environment. In *Proceedings of the 3rd acm conference on electronic commerce* (pp. 180–190). ACM Press.

Kephart, J. O., & Fay, S. (2000). Competitive bundling of categorized information goods. In *Proceedings of acm conference on electronic commerce (ec-00).* ACM Press.

Kephart, J. O., Hanson, J. E., & Sairamesh, J. (1998). Price and niche wars in a free-market economy of software agents. *Artificial Life*, *4*(1), 1–23.

Kutschinski, E., & Poutré, H. L. (2001). *Scientific techniques for interactive profiling* (Tech. Rep.). CWI.

Mankila, M. (1999). *Price bundling applied to retail banking* (FE Rapport No. 367). Göteborg University, School of Economics and Commercial Law. (online available at "http://swoba.hhs.se/gunwba/")

Schafer, J., Konstan, J. A., & Riedl, J. (2001). Electronic commerce recommender applications. *Data Mining and Knowledge Discovery*, *5*(1–2), 115-153.

Schamlensee, R. L. (1984). Gaussian demand and commodity bundling. *Journal of Business*, *57*(1), S211–S230.

Stigler, G. J. (1963). United states v. loew's inc.: A note on block-booking. *Supreme Court Review*, 152–157.