

This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: Innovation Policy and the Economy, Volume 7

Volume Author/Editor: Josh Lerner and Scott Stern, editors

Volume Publisher: MIT Press

Volume ISBN: 0-262-10121-1

Volume URL: <http://www.nber.org/books/jaff07-1>

Conference Date: April 19, 2006

Publication Date: February 2007

Chapter Title: Innovation Incentives for Information Goods

Chapter Author: Erik Brynjolfsson, Xiaoquan (Michael) Zhang

Chapter URL: <http://www.nber.org/chapters/c0035>

Chapter pages in book: (p. 99 - 123)

Innovation Incentives for Information Goods

Erik Brynjolfsson, *MIT Sloan School*

Xiaoquan (Michael) Zhang, *MIT Sloan School*

Executive Summary

Innovations can often be targeted to be more valuable for some consumers than others. This is especially true for digital information goods. We show that the traditional price system not only results in significant deadweight loss, but also provides incorrect incentives to the creators of these innovations. In contrast, we propose and analyze a profit-maximizing mechanism for bundles of digital goods which is more efficient and more accurately provides innovation incentives for information goods. Our "statistical couponing mechanism" does not rely on the universal excludability of information goods, which creates substantial deadweight loss, but instead estimates social value created from new goods and innovations by offering coupons to a relatively small sample of representative consumers. We find that the statistical couponing mechanism can operate with less than 0.1 percent of the deadweight loss of the traditional price-based system, while more accurately aligning incentives with social value.

I. Introduction

Background

Innovation is an important driver of firm productivity and social welfare for goods in general, and for digital information goods in particular. The ultimate outcome of these innovations, whether for music, software or other information goods, is improved consumer value. For the society, as long as the expected cost of introducing an innovation is smaller than the expected benefit (social welfare), the innovation should be encouraged.

Firms, however, may not always want to innovate in these circumstances. In general, they innovate only if the expected profit is higher than the expected cost. Furthermore, they may undertake some

innovations which are profitable but do not increase net social welfare. Firms' profit objectives are usually not completely aligned with the society's social welfare objectives, and when this happens, there is deadweight loss.

In the traditional price system, the objective of firm profit is aligned with the objective of social welfare only when the price of a good is set at the marginal cost of producing the good. This is usually unattainable in the real world, and we show below that even when social efficiency is ensured (firms setting price equal to the marginal cost), innovation incentives are still not correctly provided to the firms from a social planner's point of view. For digital information goods, where marginal cost of copies approaches zero, the price system is particularly problematic. Not surprisingly, business models for digital information are often chaotic, confusing, and unstable.

In this paper, we establish a simple framework to study how an innovation can change the consumer's valuations, and we identify two types of innovations: uniform enhancement and targeted innovation. We show that the traditional price system cannot generally provide correct innovation incentives for firms to innovate, and a better price system should be rewarding creators based on their *social contributions*. Our proposed mechanism addresses this issue for digital goods with the help of the unique property of digital goods, namely, the near zero marginal cost of these goods. Interestingly, it is this very property that creates serious pricing issues for the various digital goods industries. For example, digitized music has been blamed for the decline in music CD sales since 2001. The availability of digital music is said to threaten the incentives for innovation and creativity itself in this industry. It has engendered a ferocious backlash, with thousands of lawsuits, fierce lobbying in Congress, major public relations campaigns, sophisticated digital rights management systems (DRMs), and lively debate all around. Music is not the only industry affected. Software, news, stock quotes, magazine publishing, gaming, classified ads, phone directories, movies, telephony, postal services, radio broadcasting, and photography are just a few of the other industries also in the midst of transformation. Two predictions can be made with near certainty about the next decade: the costs of storing, processing and transmitting digital information will drop by at least another ten-fold and the vast majority of commercial information will be digitized.

The debate reflects two opposing economic ideas. On one hand, the near-zero marginal costs of reproducing digital goods suggests that static welfare, the sum of consumer and producer surplus, would be maximized by making these goods available at zero price. In that way, all consumers with a value greater than the marginal cost, i.e., zero, would have access to them, and deadweight loss would be minimized. On the other hand, a zero price would provide no revenues to the creators of the goods, and thus no incentives for their creation in the first place, leading to potentially even larger losses of social welfare over time.

Thus, the debate centers on who will be impaled on the two horns of the dilemma: should creators be deprived of the rewards from their creations or should users be deprived of goods which cost nothing to produce? Either approach is demonstrably suboptimal (e.g., Lessig 2004). It would seem impossible to have both efficiency and innovation when it comes to digital goods. Improving one goal appears to be inextricably intertwined with hurting the other goal.

Preview of the Paper

In this paper, we explore the possibility of a third way. In particular, we develop and analyze a method for providing improved incentives for innovation to the creators of digital goods. We show that it is possible to decouple the payments to the innovators from the charges to consumers while still maintaining budget balance. In this way, we can deliver strong incentives yet unhindered access to the goods for almost all interested consumers. In fact, we find that our system actually provides *better* incentives for innovation than the traditional price system, even if the traditions system is bolstered by powerful DRMs and new laws to enhance excludability and thus monopoly power.

We argue that in some cases it may be misguided to try to force the old paradigm of excludability onto digital goods without modification. Ironically, DRMs and new laws are often used to strip digital goods of one of their most appealing and economically-beneficial attributes—the ease of widespread use. At the same time, we take seriously the need to reward innovators financially if we wish to continue to encourage innovation and creativity.

The essence of our mechanism is to (1) aggregate a large number of relevant digital goods together and sell them as a bundle and then (2)

allocate the revenues from this aggregation to each of the contributors to the bundle in proportion to the value they contribute, using statistical sampling and targeted coupons. We do this in a way which is fully budget-balancing (meaning no outside subsidy needed for the system to work) efficiency losses as small as 0.1 percent of the traditional price system. Furthermore, our mechanism provides substantially better incentives for content creation than a "perfect" implementation of the traditional price based system where goods are sold individually and creators keep 100 percent of the revenues.

Large digital collections are increasingly common as much Internet content moves from free to fee-based systems and as new forms of digital content, such as satellite radio, emerge. Consider the bundles that constitute XM radio, Cable TV, AOL content, Rhapsody music, *Consumer Reports* reviews, JSTOR academic articles, and Microsoft Office software.

Bundling has been analyzed in some depth in the academic literature, including a cluster of articles specifically focusing on the bundling of digital information goods (e.g., Bakos and Brynjolfsson 1999, 2000 and the references therein). A key finding from the literature is that in equilibrium, very large bundles of information goods can provide content that is accessible to the vast majority of the consumers in the relevant market. It will not be profitable to exclude (via pricing) any consumers except the small fraction who have improbably low valuations for an improbably large number of the goods in the bundle. Thus, bundling can dramatically increase economic efficiency in the allocation of information goods to consumers.

Given the prior literature on bundling information goods, our paper focuses on the second part of the mechanism, which involves designing a system for allocating revenues from such a bundle. This is necessary because by its very nature, bundling destroys the critical knowledge about how much each of the goods in the bundle is valued by consumers. Did the consumer subscribe to XM radio for the classical music or for some other piece of content that was in the bundle? How much did the consumer value each of these components? Unlike for unbundled goods, the consumer's purchase behavior for the bundle does not automatically reveal the answers to these questions. This creates a problem when it comes time to reward the creators and providers of the component goods. Surveys, usage data and managerial "instinct" can all help allocate revenue to reward content creators, but none is likely to be as accurate as a true price-based system.

Our mechanism re-introduces prices, but only for a tiny fraction of consumers. For instance, in a large-scale implementation, only 1,000 consumers out of several million would face any prices for individual goods, typically via special coupons. Because of the law of large numbers, this allows us to get fairly accurate, unbiased assessments of value of the good as long as these consumers are chosen randomly, or better yet, representatively. However, because the vast majority of consumers do not face any non-zero price for individual goods, they incur virtually no inefficiency. Specifically, 99.9 percent of users have access to any given good as long as their value for that good is greater than zero and their values for all other goods in the bundle are not simultaneously unusually low.¹

In particular, our paper introduces a "statistical couponing mechanism" and argues that it is technically feasible and that it can dominate any of the approaches debated thus far. Barriers to diffusion and assimilation of this approach are likely to include overcoming knowledge barriers and some measure of organizational and institutional learning. Our analysis is meant to be a first step in addressing these obstacles. Notably, if this innovation succeeds, it should actually increase the pace of *future* innovations by improving incentives for the creation of useful digital goods. At a minimum, a broader discussion of this type of approach should change the terms of the existing debate about business models for digital goods.

In the remainder of this section, we review the related literature. Section II analyzes the incentives for innovation under the standard price system and shows how they are incorrect. Section III discusses some possible ways to address this issue using IT, as well as the weaknesses of each of these alternatives. Section IV introduces our statistical couponing mechanism in detail and provides simulations that demonstrate its high efficiency when the number of consumers is reasonably large. Section V provides some remarks on the feasibility of our mechanism and section VI concludes with a brief summary.

Related Literature

The academic literature related to our analysis is somewhat sparse. Some of the closest research is the work on a monopolist facing an unknown demand curve (e.g., Aghion et al. 1991) where it is shown that the seller can experiment by pricing to different buyers sequentially and updating the price accordingly. In addition, as discussed

later in our paper, Spence (1976) discusses some related problems with incentives for investments in improving quality.

We are not aware of any systems which fully implement *both* parts of our mechanism, although bits and pieces are used in various industries and applications. For instance, as noted above, there are many examples of bundling for digital goods. Revenue allocation similar to our approach is more difficult to find. The American Society of Composers, Authors and Publishers (ASCAP) does seek to monitor the consumption of its members' works and distribute its revenues to each creator in rough proportion to this consumption. However, they generally do not use direct price data and thus typically work under the implicit assumption that all songs have equal value to each listener.

William Fisher's (2004) book explores various solutions to the music piracy problem brought about by the new peer-to-peer technology. Specifically, he proposes to replace major portions of the copyright and encryption-based models with a "governmentally administered reward system." He correctly points out that to assess the correct level of these rewards, what we really need is not the number of downloads, but the "frequency with which each recording is listened to or watched" (i.e., the real value to consumers). Fisher's proposal is similar to the Nielsen TV sampling approach, and he proposes to implement special devices to estimate the frequency of listenings of each recording. He also suggests that the frequency should be multiplied by the duration of the works, and that consumer's intensity of enjoyment, obtained through a voting system, should be taken into consideration to make more precise estimates of the valuations.

This proposal, if carried out, could be superior to the current practice taken by ASCAP (BMI, SESAC, etc.) to compensate the creators of musical works, and it comes very near to the ideal of learning consumers' valuations and distribute money accordingly; but it also suffers from several inherent problems. First, unlike from Nielson TV sampling, people may use different devices to enjoy the same digital content. For example, a song can be played with an MP3 player in the car, a CD player in the home entertainment system, or a DVD drive on a computer. Second, and more critically, as shown in the public goods literature, a voting system such as that proposed by Fisher is not reliable because individual hidden incentives can induce voters to misrepresent their true values. For instance, consumers might falsely claim to have an extremely high or low value for a good in an attempt to influence the voting. In essence, the Fisher approach still does not provide a

reliable, incentive-compatible way to determine the true value of each good to consumers.²

II. Incorrect Innovation Incentives

Providing correct innovation incentives can be an issue for information goods. It is important to note that innovation incentives are often dramatically incorrect in the traditional pricing mechanism. This is exacerbated not only by the very low marginal costs of information goods, but also by another property, their enormous malleability and flexibility. Unlike most physical goods made of atoms, goods made of bits can easily be redesigned and reconfigured. Accordingly, unlike in traditional manufacturing and service industries, the core production workers, at companies that produce information goods like software work on changing the design of existing products and introducing new ones, not on manufacturing and distributing copies of existing designs, which is relatively trivial by comparison. This means that innovations in information goods can be highly targeted to specific consumer segments, if the seller so desires.

The traditional price system based on excludability does an impressive job in allocating resources and encouraging innovation. We argue that the traditional pricing mechanism, however, does not ordinarily provide correct innovation incentives to producers of information goods.

Suppose that the seller can invest in trying to create an innovation which improves consumers' valuations of her product. The investment can be in the form of improving product quality, functionality, or educating users to use the product more effectively. We now give a closer look at the innovation incentives of the seller.

In the next sections, all results are depicted with figures of arbitrary demand curves.

Uniform Enhancement

We begin with the simple case of an innovation that affects all consumers equally. In particular, suppose the innovation can increase each consumer's valuation by δ . This is equivalent to moving the demand curve upward by δ .

When the demand is shifted upward, the monopolistic seller will be charging a new price $p' = p^* + \varepsilon$ that maximizes her profit. With this

innovation, she can expect to gain some additional profit indicated by the shaded area in figure 4.1. In the figure, although the potential value the seller has created for society is the area between the two demand curves, in the traditional price system (with a single price for the good), she gets additional revenue and profit indicated by the shaded area. This shaded area is also, therefore, the amount of incentive for creating the innovation—the seller will pursue such innovations if the expected value is greater than the expected cost. The areas representing the value of the innovation to society and the value of the innovation to the seller are not necessarily equal. Part of the seller’s profit from the innovation comes from transferring surplus between the consumer and the seller, which has no net benefit to society. On the other hand, part of the profit also comes from reducing the deadweight loss to a certain extent, which does improve social welfare. Thus, depending on the exact shape of the demand curve, the incentives for innovation can be inaccurate.




Figure 4.1
A uniform upward shift of demand curve.

Targeted Innovation

Incentives are particularly misaligned for innovations which can be targeted, so they affect only a small subset of consumers' valuations. In particular, the innovation may be less significant so that only some consumers with valuation near some \tilde{v} are affected. For instance, consider three cases. a software developer could invest in adding features that could

1. make satisfied users of its product even more satisfied,
2. increase the value to consumers whose values were just below the market price, turning them into buyers, or
3. increase the value of non-buyers, but not enough to turn them into buyers.

Suppose that the developer has a limited budget and can only pursue one of these three types of innovations. Even if innovations of type (1) or (3) might create more value for society, the traditional price system will only provide incentives for innovation (2).

Figure 4.2 shows this graphically.³ When the seller takes some effort to innovate and creates some extra social value indicated by the area ABC , we can easily see that in the traditional price system, given




Figure 4.2
Social benefit/loss vis-à-vis seller innovation.

a fixed level of δ , the possible region of innovation cannot be far away from the optimal price p^* . This narrow range can be indicated by a pair of values: \bar{v}_L and \bar{v}_H . In other words, for all values lower than \bar{v}_L or higher than \bar{v}_H , the seller can not extract enough profit to justify the cost of innovation, so only innovation (2) will be carried out under the traditional price system. This is very intuitive. In the traditional market, if the seller sells goods to consumers with valuation higher than \bar{v}_H , it does not help to increase their valuations further because that will only contribute to consumer surplus, and the seller will not be able to extract the added value. Similarly, for the potential consumers with lower valuations (lower than \bar{v}_L , to be precise), the seller will not take the effort to innovate because these people will not be converted to consumers. For small δ , the range (\bar{v}_L, \bar{v}_H) is very small, and even in this range, innovation may not always be socially desirable.

Thus, in the traditional price mechanism, the seller has too little incentive to create innovations that mainly benefit consumers with very low or very high valuations. If your valuation is substantially above or below the equilibrium price for a good, don't expect the good's provider to put significant effort into innovating to specifically address your needs.

Interestingly, even while under-investing in certain types of innovations, the seller also has incentive to over-invest in other types of innovations. To see the socially wasteful incentive of innovation in the traditional price system, consider the case of the consumers' valuations near the optimal price. For example, if the seller takes an effort to innovate and increases the valuation for some consumers from p^* to $p^* + \delta$, but not for any other consumers, then her optimal effort is proportional to the triangle indicated by AJF , but her gains are as large as the entire rectangle indicated by DJF . When δ is small, the ratio of the seller's incentives to the optimal incentives can become arbitrarily large. In other words, meaning that the seller can have radically excessive incentives to innovate for people whose valuations are close to the optimal price p^* .

This is a striking result, that innovation for people whose valuation is just above the optimal price will reduce consumer surplus, yet this is exactly the range where it is most profitable for the sellers to innovate. An innovation in this neighborhood costs many times more than the value it creates would still be profitably pursued.

Discussion

The issue with targeted innovation is but one manifestation of misallocated (or narrowly-allocated) resources for creating values for consumers. If we look at one product as a collection of functional features, then the creator will be too focused on innovations around those features that cater to the marginal consumer (and thus ignoring possible innovations for consumers with much higher or much lower valuations). In a product bundle, through the traditional price system, the bundler will immediately get a positive feedback if she introduces a new product catering to the marginal consumer as this will turn some non-buyers to buyers. However, if she introduces an improvement to the bundle that caters to the higher valued consumers, she cannot see a corresponding profit. In the long run, this feedback process will discourage the bundler to introduce anything far away from catering only to the marginal buyers.

Spence (1976) studied the inability of prices to convey information about quality improvements in products. He argued that if firms are not perfect price discriminators (i.e., if firms are not paid according to the social value they create), then the profits are not equal to its net contribution to surplus, and since "profitability is the criterion by which products are selected or rejected in a market system, ... this may not always lead to desirable results." Compared with deadweight loss, this type of inefficiency has largely been neglected in the literature. In the next sections, we will be examining the special property of digital goods, and we propose a pricing mechanism that avoids exactly this kind of inefficiency as well as the traditional deadweight loss inefficiency.

III. Bundling and Mechanisms for Providing Incentives for Digital Innovation

Bundling Can Reduce Deadweight Loss

As noted in the introduction, if the marginal cost of producing a good is zero, charging any price greater than zero for that good can be socially inefficient: some consumers (e.g., those with valuations less than the price but greater than zero) will be excluded from consuming the good even though it would be socially beneficial for them to have access to it.

The zero marginal cost of reproducing digital goods has created huge pricing problems for various industries: it takes work, talent, and luck to create a successful CD, but once digitized as an MP3 file, any piece of music can be reproduced with virtually zero cost. The music industry has been profoundly influenced by this property of digital goods, and we are very likely to see more industries follow suit. Technology has enabled us to distribute digital goods more efficiently, but we must find the right mechanism to encourage their creation and to allocate them. Without a good mechanism for digital goods, we will not be able to provide sufficient innovation incentives for the creation of these digital goods.

It is shown in Bakos and Brynjolfsson (1999) that, in certain cases, bundling can be a partial solution for pricing of digital goods. By the law of large numbers, it is easier to find an optimal price for a bundle of digital goods than for each individual good. In equilibrium, the profit maximizing price for a large bundle will be set low enough so that virtually all consumers interested in any of the goods in the bundle will choose to buy the whole bundle (even if they use only a small fraction of its components). For instance, most PC users buy Microsoft Office, even if they don't use all its applications, or even not all of the features of the applications that they do use. While there may be anti-competitive implications (see Bakos and Brynjolfsson 2000), such bundling does give the socially desirable result of dramatically reducing the dead-weight loss because fewer consumers are excluded from using any of the bundled goods in equilibrium. In essence, once consumers pay a lump-sum to purchase the bundle, they can consume any of the goods in the bundle at zero marginal cost. Thus, when the cost of reproducing the goods is close to zero, bundling can provide close-to-optimal allocation of goods to consumers (Bakos and Brynjolfsson 1999).⁴

However these benefits come at a major cost. Bundling inherently destroys information about how each of the component goods is valued by consumers. Is the bundle selling because of the fresh sounds of a new artist or due to the lasting appeal of a traditional favorite? Without this information, it is impossible to allocate revenues to the providers of content in a way that accurately encourages value creation. Selling goods individually would automatically solve this problem, but as discussed above, individual sales create enormous inefficiencies because they exclude some users with positive value from access to the good.

Basically, bundling helps to address the innovation incentive problem by offering a viable business model to reward the creators of digital

goods with much less deadweight loss than à la carte pricing. However, bundling introduces another problem of innovation incentives—this mechanism does not give us a natural solution to distribute the revenue to provide correct incentives for each of the goods' creators. We will discuss the revenue distribution problem in the next section.

To illustrate the problem, we can make a simple comparison between items sold in Wal-Mart and songs sold in an online subscription service to digital goods. Every item in Wal-Mart will go through the POS scanner, so Wal-Mart knows if the blue jeans from Levi's sell better than the jeans from Eddie Bauer's, and this information can be used to quickly adjust purchasing and pricing. If Levi's produces better jeans, the price system will automatically reward the company with more revenues. This is a very desirable situation for all parties: the consumers can have access to products they like, Wal-Mart can respond to the market very quickly and ensure a competitive advantage, and most importantly, the creators of better products can get automatically rewarded.

When we observe a consumer subscribing to a bundle of digital goods, however, we do not automatically know from his purchase which song, feature, or other component he values more, thus the creator of the favorite component can not be properly rewarded. It is interesting to note that any form of bundling creates this problem, no matter the components are digital or not. For example, when we see people buying a subscription to cable TV, we do not automatically know which channels they value more than others; when people buy the Microsoft office bundle, we do not automatically know whether Word or Excel is more valuable.

The Revenue Distribution Problem

The ideal revenue distribution mechanism would be one which somehow determined each good's demand curve, and distributed the revenue among the content providers in proportion to the social value of each good to all consumers. This value can be calculated by integrating the area below each good's demand curve. Various mechanisms used to derive demand curves proposed in the literature all fail here because bundle pricing does not automatically provide ways to observe the market's response to a price change of individual goods.

If the benefits created by each good cannot be observed or calculated, then a host of inefficiencies may result. First, the content providers may not have enough incentives to produce creative products in the first

place, and consumers will eventually be harmed. Second, without a good signal of consumers' preference, content providers may produce content, but not the content that best fit the consumers' taste. Third, in any effort to overcome these problems, the collection of content producers may lead the potential bundler to adopt other strategies such as à la carte pricing. However, such strategies re-introduce the deadweight loss problem discussed at the beginning of section I.

In the following subsections, we discuss the costs and benefits of several ways to distribute revenue to address this challenge, culminating with our proposed statistical couponing mechanism.

Payment Determined by Number of Downloads

In the context of digital information goods, it is often natural to assume that the seller may be able to observe the number of times that each good is accessed. This gives us the following approach.

If one is willing to assume that the number of accesses signals popularity and popularity is a measure of value, we can infer the value by the number of accesses. Traditionally, this scheme is broadly used in the market of digital goods such as music, movie, TV shows, and software. For example, each episode of *Friends* got about 29 million viewers per week in its last year, which was far more than most other TV shows; as a consequence, each of the six stars was paid \$1.2 million per episode, far more than most other TV actors.

More formally, suppose we have n goods in the bundle, the price for the bundle is B . Also suppose there are m buyers of the bundle, each represented by j ($j = 1, \dots, m$), then the total bundle revenue is $R = B \cdot m$. We assume the system can record the number of downloads of buyer j for good i : d_{ij} , then the provider of content i should be paid:

$$\text{revenue}_i = R \cdot \sum_{j=1}^m \frac{d_{ij}}{\sum_{k=1}^n d_{kj}}. \quad (1)$$

This method is extremely easy to implement. In fact, the last equation implies that the bundler does not even have to keep record of all the downloads made by the m buyers, she can simply record d_i , the number good i has been downloaded.

This method is powerful when all the goods are approximately equal in value. However, if goods differ in value (bundling very cheap "Joke-A-Day" with more expensive "Forrester Research Report"), then

pricing based on number of downloads is misleading. (The Joke-A-Day may be downloaded more times than the Forrester Research Report, but aggregate value of the latter may be much higher). Another problem with this method is that it gives dishonest content providers a way to distort the values by manipulating the number of downloads of their own content. This has been a problem, for instance, with some advertising-supported content where prices are based on thousands of impressions recorded (Hu 2004).

Payment Determined by Downloads Combined with a Stand-Alone Price

Number of downloads itself is not a good measure of consumer valuation in many cases. Assuming there also exists a stand-alone price for every information good in the bundle, and assuming these prices are all fair prices, we can then derive an improved mechanism to distribute the revenue.

Consider the market introduced in the first subsection of part III. Suppose each item i ($i = 1, \dots, n$) in the bundle also has a stand-alone price p_i .

Building on the equation from the third subsection, an improved way to distribute the revenue is through the following formula:

$$\text{revenue}_i = R \cdot \frac{\sum_{j=1}^m p_i d_{ij}}{\sum_{k=1}^n p_k d_{kj}} \quad (2)$$

that says that the revenue to distribute to content provider i should be a proportion of the total revenue ($R = B \cdot m$), and the proportion is determined by the sum of each consumer's valuation of good j .

This method has the advantage of being more precise compared to the previous solution. Indeed, if "Joke-A-Day" is sold separately, its price will probably be much lower than that of "Forrester Research Report." The disadvantage of this method is that a fair and separate price may not always be readily available. If the distribution of revenue is set according to this method, and when bundling becomes a major source of revenue, there is room for content providers to misrepresent the stand-alone price. Furthermore, this approach implicitly assumes that the value from each good is proportional to the stand-alone price. This will only be true, however, if the price paid by the marginal consumer of each good is proportional to the average price that would be paid by all consumers of that good, for all goods.⁵

Discussion

As discussed above and in the introduction, merely adding an accounting framework to the traditional price system does not guarantee a socially efficient outcome of distributing the digital goods and providing correct innovation incentives to the sellers. In the next section, we propose a mechanism that goes beyond the excludability-based traditional price system.

IV. Statistical Couponing Mechanism

Description of the Mechanism

As discussed in the last section, the ideal way to provide correct incentives is to learn consumers' valuations for each good and make corresponding payments. Since bundling itself obscures consumers' valuations for individual goods, here we propose a mechanism to reveal the demand curve for each good by issuing targeted coupons to a small sample of consumers. For large populations, it is possible for the targeted sample to be large enough to be representative statistically while still being small enough to be fairly unimportant economically. Our mechanism is substantially different from the traditional use of coupons as a marketing method to price discriminate consumers. Instead, coupons in our mechanism are similar to the price experiments suggested in the optimal pricing literature.

Suppose the monopolistic bundler offers a bundle of information goods to a group of consumers. In order to derive the demand curve for one of the components, we choose $m \cdot n$ representative consumers and issue each of them a single coupon, where n is the number of price levels covering the range of the valuations, which we call "coupon levels," and m is the number of coupons to be offered for each of the price levels in total, which we call "sample points." While $m \cdot n$ should be large enough to make statistically valid inferences, it can nonetheless be a very small fraction (e.g., 1/1000 or less) of the total set of consumers buying the bundle.

If a consumer receives a coupon with face value \tilde{v} , then he can either choose to ignore the coupon and enjoy the complete bundle or choose to redeem the coupon and forfeit the right to use a particular component of the bundle, which is indicated on the coupon. So upon observing a consumer's action, the bundler can learn whether that consumer's

valuation for that component is higher or lower than the face value of that particular coupon. Aggregating the m consumers' valuations will give the bundler a good estimate of the valuations at that price; summarizing the results for the n coupon levels, the bundler can plot a fairly accurate demand curve.⁶ The area under the resulting demand curve, when scaled up to the whole population, is the total social valuation for that particular good, and also the maximum revenue which that good can contribute to the bundle revenue. Using the same method for all the components, the bundler can learn the approximate social valuation and revenue potential of each of the goods in the bundle. She can then distribute the revenue among the content providers according to their contribution share to the total valuation. Let R be the total revenue from selling bundles, and v_i be the social value of the component i in the bundle, content provider of i should be paid

$$\text{revenue}_i = R \frac{v_i}{\sum_{j=1}^N v_j}, \quad (3)$$

where N is the total number of content providers.

Comparison with Other Methods for Providing Innovation Incentives

This method compares favorably to the traditional price mechanism. The traditional price mechanism subjects 100 percent of consumers to the inefficiency of positive prices. However, only data from a small fraction of consumers are needed to get fairly accurate estimates of the value created and contributed by each good. The greater precision obtained by increasing the sample declines asymptotically to zero while the cost for subjecting each additional consumer to a positive price remains just as high for the last consumer sampled as the first one. When balancing the costs and benefits, the optimal sample size is almost surely less than 100 percent. Secondly, the proposed couponing mechanism actually provides a *more* accurate estimate of the overall demand curve than any single-price traditional system. Because multiple different prices for coupons are offered, a much more accurate overall picture of demand can be obtained than simply revealing the demand at a single price, as conventional prices do. As discussed in section I, this has large and important implications for dynamic efficiency and innovation incentives.

One can also compare our couponing mechanism with the well-known Vickrey-Clarke-Groves (VCG) mechanism. Unlike VCG, our mechanism does not give us exact valuations for each consumer. However, in general, approximate demand functions of the components will suffice, and by increasing the sample size, the accuracy can be made almost arbitrarily precise. Our mechanism is superior to the VCG mechanism in several ways. (1) Truth-telling is a robust and strong equilibrium in the couponing mechanism, in the sense that each consumer simply compares his valuation with the coupon's face value; he is not required to assign correct beliefs on all other people's votes. (2) In the VCG, if one respondent misreports his value (due to irrationality or due to error), the consequence may be very severe for the rest of the people. Furthermore, coalitions of consumers can "game" the VCG to their advantage. However, in the couponing mechanism, the effects on others from a consumer's misreport are minimal. (3) The couponing mechanism is fully budget balancing, unlike the VCG. (4) The couponing mechanism is more intuitive than the VCG for real world problems.

Simulation Results for the Mechanism

It can be shown that for any one of the components in the bundle, given a large number of randomly chosen respondents and level of coupons, the above mechanism gives an empirical demand function that arbitrarily approximates the true demand function (see Brynjolfsson and Zhang 2004). We can also run simulations to see the effectiveness of this mechanism (see figure 4.3).

The use of our coupon mechanism gives us empirical estimates of the inverse demand curves for each of the distributions, and we define the error rates to be the percentage differences between the area under the empirical demand curve and the area under the true demand curve. Figure 4.3 shows the result of the coupon mechanism applied to the uniform distribution; other distributions yield similar figures. We see that error rate is declining with more coupon levels and with more sample points for each coupon value. With just 20 coupon levels, the error rate is as low as five percent. Adding more sample points for each coupon value also helps to improve the precision. For example, with 40 coupon levels, sampling 20 consumers for each coupon level (for a total of 800 respondents) gives us an error rate of ten percent, and sampling 80 consumers improves the error rate to be near five percent. From the error

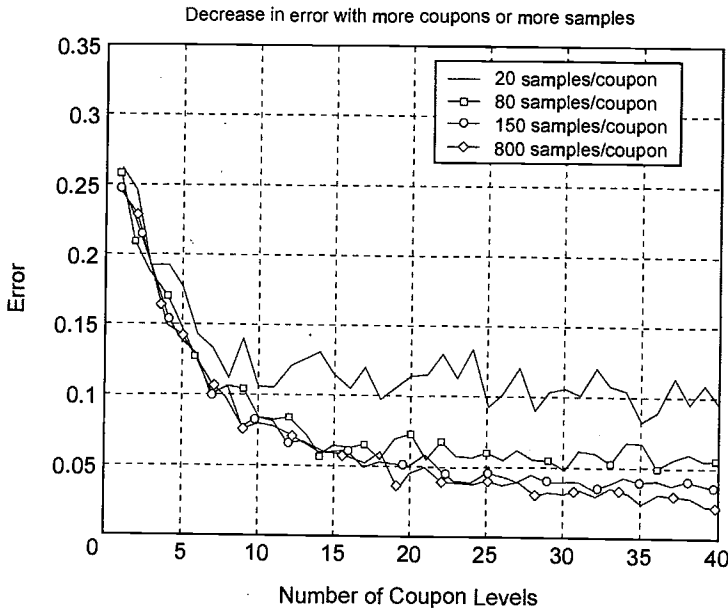


Figure 4.3
Simulation results for the couponing mechanism.

rate curves, we can also see that when sampling 20 consumers, adding coupon levels beyond ten does not improve the precision significantly; also, when sampling 80 consumers, adding coupon levels beyond 15 does not improve the precision significantly. This observation tells us that we have to add coupon levels and sampling points simultaneously in order to achieve the best result estimating the social values of goods. Error rate converges toward zero more quickly for fatter demand curves (the ones with a higher expected value). In our simulations, for some demand curves, with just five coupon levels and 20 sample points (for a mere 100 respondents), the mechanism can give us an error rate below 0.1 percent. Thus, sampling just 100 consumers can provide almost as accurate an estimate of demand as sampling the entire population of consumers of the good, which could be in the millions.

The deadweight loss is proportionately smaller, too. Consumers who cash-in the coupons forgo access to the corresponding good, which creates a deadweight loss (unless the consumer's value was exactly zero). For such a consumer, this decision is analogous to facing a market price, with similar costs, benefits and overall incentives. In contrast to the traditional price approach, however, our mechanism only subjects a

fraction of consumers to this decision, so only a fraction choose not to buy, and the total deadweight loss is only a fraction of what it used to be.

This mechanism can be used to solve the revenue distribution problem discussed in section III, and we can show that this mechanism can also help to avoid the innovation incentive issues arising in traditional price systems:

1. If an innovation can increase only some consumers' valuations, the traditional price system does *not* provide correct incentives for the producer to innovate for people with relatively high or relatively low valuations. In contrast, the proposed mechanism always gives the producer socially desirable level of incentives to innovate, and
2. As shown in the analysis in section II, the traditional price system gives the producer too high an incentive to innovate where it is most harmful to the social welfare, and no incentive elsewhere; the proposed mechanism induces the producer to make socially desirable innovation efforts.

V. Remarks on Feasibility

This paper contributes to establishing a more efficient approach to create, distribute and consume digital goods. The theoretical foundation proposed here is just the first step toward this goal; in order to build viable business models, we need to address some practical issues to be discussed below.

In this paper, couponing has been analyzed solely as a mechanism for revealing existing demand, not for influencing it. Of course, in practice, couponing may also be viewed as a form of advertising that increases demand. If it increases demand for some goods, and not for others, then the estimated values may be biased in a non-uniform fashion.

There is a related, more conspicuous problem: due to the heterogeneity in people's tastes, some goods are surely downloaded less than some others (consider an analyst's research report; maybe only tens out of millions of consumers would want to download it). If we do not offer enough sampling points, there will be a bigger error in estimating demand for these less popular goods.

Both issues can be addressed by a practice we call "passive couponing." Under "passive couponing" regime, only those who downloaded a good will be offered a coupon for that good. After downloading, the consumer learns all the product characteristics, so the informative role

of couponing as advertising is no longer valid. For goods downloaded by the majority of people, we can choose a small fraction out of them to offer coupons, and for goods downloaded only by a few, we may offer coupons to most or all of them. In either case, subsequent access to that good, or similar goods, can be restricted for consumers who prefer to redeem the coupon instead. By discriminating coupons offered to different types of goods, we can get a better overall estimate of the specific demands.⁷

In previous sections, we did not consider the issue of contract duration. It is likely to be unnecessary to permanently block access to a good for consumers who redeem the corresponding coupon. Temporary blockage will generally suffice. We can put this question into the context of subscription-based business models. Suppose the bundle is to be paid by month (e.g., \$20/month), then for time-critical information goods (e.g., news, stock quotes, etc.), we can offer the coupons by month, too (e.g., "Take this \$1 coupon and give up access to CNN news for the next month"). For those less time-critical information goods (e.g., music, software updates, etc.), we can offer the coupons by longer periods (e.g., "Take this \$10 coupon and give up access to music by Madonna for the next year").

What if the valuations are not independent as assumed in the paper? If two goods are substitutes, offering a coupon for one of them will only help us to estimate the incremental value that it brings to the bundle, and this is also true for the other good, so we will be paying less for the two creators than the value they bring into the bundle. Similarly, for complements, we overestimate total value of the goods. If we can identify clusters of goods that are substitutes or complements to each other, we can offer coupons for individual clusters and use the proposed mechanism to estimate the share of contribution by each cluster. This will ensure that a cluster of content providers will be paid a fair overall payment. Within a cluster, each individual content provider can be paid according to the estimated share of incremental value they bring to the cluster.

VI. Conclusion

Major innovations in technologies often engender innovations in business organizations. The digitization of information is no exception. We seek to advance the debate on how best to allocate digital goods and reward their creators by introducing and exploring a novel mechanism and analyzing its implications.

Our approach eliminates the marginal cost of consuming digital information goods for the vast majority of consumers via massive bundling. For very large aggregations, this preserves most of the static efficiency which could be achieved with a zero price policy. However, in the long run, the more important issue is how to create incentives for ongoing innovation. Indeed, our living standards, and those of future generations, depend far more on continuing innovation than on simply dividing up the existing set of digital goods. In this area, the proposed statistical couponing mechanism shows particular promise. We find that our approach can provide substantially better incentives for innovation than even the traditional monopoly price system bolstered by artificial excludability (e.g., via DRMs, laws, etc.). In particular, the traditional price system, in which each good is sold for a specific price with the proceeds going to the monopolist creator, focuses virtually on incentives on a very narrow band of consumers—those just on the margin of buying. In fact, the price system provides *too* strong incentives for innovations that help this narrow group of consumers. Rents transferred to the creator from such innovations exceed the social benefits. In contrast, our approach, using statistical sampling and couponing, can provide incentives which are nearly optimal for every type of innovation.

In summary the mechanism we introduce,

- potentially has orders of magnitude less inefficient than the traditional price system,
- is budget balancing, requiring no external inflows of money,
- works with existing technology and existing legal framework,
- requires no coercion and can be completely voluntary for all parties, since it is fully incentive compatible,
- doesn't assume that innovators will continue to innovate even without financial rewards,
- can be implemented and run in real-time, and
- is scalable to very large numbers of goods and consumers (in fact, it works better for larger numbers).

Our approach also has weaknesses and challenges. Compared to giving away all digital goods for free, our approach will exclude a small number of consumers and create some inefficiency as a result. More importantly, our approach does require the creation of new business institutions or models, which is never easy. Specifically, an entity is

needed to manage the statistical sampling and couponing, analyze the resulting data, and allocate payments to the content owners accordingly. Near misses for this type of entity already exist. For instance, ASCAP does much the same thing already for broadcast music, but without accurate price information. Nielsen and similar organizations provide usage information, but again without accurate price information. There are organizations that regularly collect and distribute large sums of money to member companies based on various algorithms. The Federal Deposit Insurance Corporation, which does this for banks, is one example. Some cooperatives are also run this way. Last but perhaps not least, the government regularly makes these types of transactions. It should be stressed, however, that our mechanism does not require any government role since all of the participants (consumers, content creators, bundlers) have incentives to participate completely voluntarily and it adheres to the existing legal framework. This stands in contrast to the proposal by Fisher (2004) or the varied proposals to change copyright or other laws.

By offering this new framework and analysis, with a new set of opportunities and challenges, we hope to lay the foundation for future research on the critical question of providing incentives for innovation in the creation of digital content and implementing mechanisms to deliver that content to consumers efficiently. Furthermore, the problems that we analyze with existing institutions for providing innovation incentives apply to a greater or lesser degree to many other products and services, not just digital goods, and variants on the mechanism we describe can also be useful in those other contexts.

We expect that the next ten years will likely witness a scale of organizational innovation for creating and distributing digital goods surpassing even the remarkable pace of the last ten years. New coordination mechanisms, such as the innovation incentive approach described and analyzed in this paper will flourish. With a proactive attitude toward technology-enabled organizational innovation, we believe that academia can speed this process by framing the issues, and by providing tools and analyses.

Acknowledgements

We thank participants in seminars at MIT, the Workshop on Information Systems and Economics, the International Conference on Information Systems, and the National Bureau of Economic Research Conference on Innovation Policy and the Economy for helpful comments. The MIT

Center for eBusiness, the National Bureau of Economic Research and that National Science Foundation (IIS-0085725) provided generous research funding.

Endnotes

1. The efficiency properties of large bundles of information goods is analyzed in in Bakos and Brynjolfsson (1999).
2. The public goods mechanism design literature seeks to provide a remedy to the voter misrepresentation problem. Specifically, the VCG mechanism can be shown to induce truth-telling by all participants. However, it has two fatal flaws. First, it is not budget-balancing—significant inflows (or net penalties) are generally needed. Second, it is quite fragile. Each participant must believe that all other participants are truth-telling or he will not tell the truth himself. Accordingly, while VCG design is intriguing in theory, it is rarely, if ever, seen in practice.
3. A more formal analysis can be found in our companion paper, Brynjolfsson and Zhang (2005).
4. If a single seller cannot provide a large enough number of information goods to achieve the benefits of massive bundling, it can be worthwhile to have a content aggregator to negotiate with multiple sellers to offer a bundle of information goods from multiple sources.
5. Barro and Romer (1987) explore how similar proportionalities can explain a number of pricing anomalies.
6. An alternative approach to revealing consumer demand would be to reverse the default consumption rights of the consumers which were targeted. The targeted consumer would be required to pay the specified offer price in order to obtain access to the selected good. Consumers who did not pay the relevant price would not have access to that good. As with the coupons, only consumers with a value greater than the relevant offer prices would choose to consume the good, thereby revealing the demand curve.
7. What if a good is only downloaded by one consumer? First of all, in this case, this good is not important in the bundle, the bundler can exclude it in the future. Second, the bundler can offer this consumer a different coupon in each period with the face value determined by a random draw. Within some periods of sampling, the bundler can still extract the true value, the math works exactly the same as in the proof of proposition 1. It can also be easily shown that there is no incentive for the consumer to misreport his value in each period.

References

- Aghion, P, P. Bolton, C. Harris, and B. Jullien. 1991. "Optimal Learning by Experimentation." *Review of Economic Studies* 58(4): 621–654.
- Bakos, J. Y., and E. Brynjolfsson. 1999. "Bundling Information Goods: Pricing, Profits and Efficiency." *Management Science* 45(12): 1613–1630.

- Bakos, J. Y., and E. Brynjolfsson. 2000. "Bundling and Competition on the Internet: Aggregation Strategies for Information Goods." *Marketing Science* 19(1): 63–82.
- Barro, R., and P. Romer. 1987. "Ski-Lift Pricing, with Applications to Labor and Other Markets." *American Economic Review* 77: 875–890.
- Brynjolfsson, Erik, and Xiaoquan (Michael) Zhang. 2004. "Slicing the Gordian Knot: A Novel Mechanism for Providing Innovation Incentives for Digital Goods." Proceedings of the 25th International Conference on Information Systems (ICIS), 517–532, December, Washington, D.C.
- Fisher, W. 2004. *Promises to Keep: Technology, Law, and the Future of Entertainment*. Stanford: Stanford University Press.
- Hu, Y. 2004. "Performance-based Pricing Models in Online Advertising." Working Paper. Sloan School of Management, Massachusetts Institute of Technology.
- Lessig, L. 2004. *Free Culture: How Big Media Uses Technology and the Law to Lock Down Creativity*. New York: Penguin Press.
- Spence, M. 1976. "Product Differentiation and Welfare." *American Economic Review* 66(2): 407–414.

