Producer Innovation Incentives and Revenue Distribution of Bundled Products

Erik Brynjolfsson and Xiaoquan (Michael) Zhang*

Abstract

Digital goods can be reproduced costlessly. Thus a price of zero would be economically-efficient for consumers. However, zero revenues would eliminate the economic incentives for creating such goods in the first place. We develop a novel mechanism which tries to solve this dilemma by decoupling the price of digital goods from the payments to innovators while maintaining budget balance and incentive compatibility. Specifically, by selling digital goods via large bundles the *marginal* price for consuming an additional good can be made zero for most consumers. Thus efficiency is enhanced. Meanwhile, we show how statistical sampling can be combined with tiered coupons to reveal the individual demands for each of the component goods in such a bundle. This makes it possible to provide accurate payments to creators which spurs further innovation. In our analysis of the proposed mechanism, we find that it can operate with an efficiency loss of less than 0.1% of the efficiency loss of the traditional price-based system. It is not surprising to find that innovation incentives in our mechanism are improved relative to the zero-price approach often favored by content consumers. However, it is surprising to find that the incentives are also substantially better than those provided by the traditional system based on excludability and monopoly pricing which is often favored by content owners. The technology and legal framework for our proposed mechanism already exist, the key issues of implementing it are organizational.

*Brynjolfsson: MIT Sloan School of Management and MIT Center for Digital Business (erik@mit.edu). Zhang: HKUST Business School and MIT Center for Digital Business (zxq@mit.edu).

Introduction

Digital goods are different. Unlike other goods, perfect copies, indistinguishable from the original, can be created at almost zero cost. With the advent of the Internet, mobile telephony, satellite communications, broadband and related technologies, these goods can be distributed to almost anyone in the world at nearly zero cost as well. Furthermore, when a person consumes a digital good, he or she doesn't reduce the stock for anyone else.

This should be a virtual nirvana. Yet, ironically, low cost digital goods are seen as a mortal threat to the livelihoods of many individuals, companies and even whole industries. For example, digitized music has been blamed for a whopping 52% decline since 2000.[1] 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 PR 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 that are also in the midst of transformation. It's said to be difficult to predict the future, but a few 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 100-fold and virtually all commercial information will be digitized. While our colleagues in computer science, both in academia and industry, deserve much praise for this, it is incumbent upon information systems researchers to understand the business, social and economic implications of these changes. Unfortunately, these implications have proven far less predictable.[2] What's more, we should go beyond prediction and seek to develop methods for maximizing the benefits from technological innovations while minimizing the costs.

Two schools of thought have dominated the debate on the economics of digital goods. One school stresses the benefits of the traditional market system. Clear property rights allow creators to exclude users from access to their creations. Users who wish to benefit from a

[1] Source: New York Times. `http://www.nytimes.com/2010/01/07/arts/music/07sales.html`, accessed June 2010.

[2] Although, as noted by Sorenson and Snis (2001), and Lyman and Varian (2004) among others, we can predict with some confidence that there will be an increasing need for, and existence of, computer supported codified knowledge and information, and the concomitant institutions for managing this information.

creation must therefore pay the creator. This payment in turn assures that a) the goods go to those individuals with the highest value for the good and b) that the creator has incentives to continue to create valuable goods. This system has been pretty successful in modern market-based economies. To many people, it seems natural to apply the same principles to digital goods, typically via a some combination of law (e.g. the Digital Millennium Copyright Act), technology (e.g. DRMs) and social education.(e.g. the software industries ongoing anti-piracy public relations efforts).

Another school of thought thinks this approach is all wrong. "Information wants to be free" some of them argue. More formally, the point can be made that since digital goods can be produced at zero marginal cost, the textbook economic principle of efficiency: "price equals marginal cost" demands that price should never be greater than zero. After all, society as a whole is only made worse off if a user is excluded from access to a digital good which could have been provided without reducing the consumption of anyone else. While appealing, this approach begs the question of how to provide incentives for the goods creators. While some creators might continue to create for the shear joy of it, for indirect economic benefits such as enhancing their reputation or competency, or out of altruism, economic systems and business models which rely solely on these motivations have historically not fared as well as those which provide more tangible rewards to innovators and creators.

Thus, the debate can be thought of as over which group should be impaled on the two horns of the dilemma: should users be deprived of goods which cost nothing to produce or should creators be deprived of the rewards from their creations? Either approach is demonstrably suboptimal (Lessig 2004, for example). 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.

In this paper, we argue there is a third way. In particular, we develop and analyze a method for providing optimal 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 slice the Gordian knot and deliver strong innovation incentives to sellers 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 when bolstered by powerful DRMs and new laws to enhance excludability and thus monopoly power.

We argue that it is 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 — low marginal costs. 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 and which provides accurate incentives for innovation with efficiency losses as small as 0.1% of the traditional price system.

Bundling has been analyzed in some depth in the academic literature (McAfee, McMillan and Whinston 1989), including a cluster of articles specifically focusing on the bundling of digital information goods (Bakos and Brynjolfsson 1999, Bakos and Brynjolfsson 2000, Geng, Stinchcombe and Whinston 2005, Fang and Norman 2006, and the references therein). A key finding from the literature is that in equilibrium, very large bundles will 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 those with very low valuations for all the goods in the bundle. Thus, bundling can dramatically increase economic efficiency in the allocation of information goods to consumers.

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 are valued. Did I subscribe to XM radio for the classical music or for jazz in the bundle? How much did I value each of these components? Unlike for unbundled goods, my purchase behavior for the bundle does not reveal the answers to these questions, creating 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, but none is likely to be anywhere as accurate as a true price-based system. Our mechanism re-introduces prices, but only for a tiny fraction of consumers. For instance, only thousands of consumers out of several million would face any prices for individual goods, typically via special coupons. This allows us to get accurate, unbiased assessments of value but because the vast majority of consumers do not face any non-zero price for individual goods, they incur virtually no inefficiency. Specifically, 99.9% 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 extremely low.

The academic literature related to this part of our analysis is quite sparse. Some of the closest research is the work on a monopolist facing an unknown demand curve (Aghion, Bolton, Harris and Jullien 1991) where it is shown that the seller can experiment by pricing to different buyers sequentially and updating the price accordingly. Some of the works on optimal market research is also relevant (Jain, Mahajan and Muller 1995).

We are not aware of any systems which fully implement *both* part 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. However, 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 have no direct price data, and thus must work under the implicit assumption that all songs have equal value to each listener.

Thus, our paper both introduces a novel mechanism and rigorously analyzes it, finding 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.

The paper is organized as follows. Section 1 describes the basic assumptions and derives the asymptotic properties of massive bundling of information goods. Section 2 introduces the problem of revenue distribution in bundling and characterizes the different types of solutions to this problem. Section 3 shows that the traditional way of distributing revenue does not provide a socially desirable innovation incentive for goods. Section 4 proposes a mechanism to solve the revenue distribution problem and gives the convergence properties. It is shown that our proposed mechanism can induce correct innovation incentives. Section 5 discusses practical issues of using the mechanism in the real world, and Section 6 concludes with a brief summary and some implications.

1 Model Setup

Our goal here is to provide a theoretical framework to which we refer in later sections.

We consider a market with many providers of digital goods and many potential buyers. Digital goods are assumed to have a reproduction cost of zero.[3] If these goods are sold separately, then any price greater than zero will 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. This is commonly called deadweight loss. In this section, we briefly show how bundling can radically eliminate this inefficiency, albeit at the cost of introducing a different problem involving incentives for innovation.

Suppose a monopolistic bundler connects the producers and the buyers by designing an optimal pricing and revenue distribution policy to maximize the bundler's profit. Each buyer has (at most) unit demand for any of the information goods. Suppose a buyer's valuations of each of the goods in the bundle are i.i.d. draws from a random variable V in the range normalized to $[0, 1]$, and that the random variable has a cumulative distribution function $F(v)$, whose corresponding probability density function is $f(v)$. In other words, a buyer's value for one good (e.g. a song) is independent of his value for an unrelated good (e.g. a piece of news). At a price of p, the demand for one good will be $D(p) = Prob(v > p) = 1 - F(p)$, yielding revenue of $\pi(p) = p[1 - F(p)]$. This implies that the inverse demand curve is $P(z) = F^{-1}(1 - z)$, and the seller's problem is to solve:

$$\pi^* = \max_p \{p \cdot (1 - F(p))\} \tag{1}$$

Taking first order condition, we have $\frac{\partial \pi}{\partial p} = (1 - F(p) - p \cdot \frac{\partial F(p)}{\partial p}) = 0$, which can be rear-

[3]In our models, the condition of zero marginal cost is important. Digital goods typically satisfy this assumption easily. However, Bakos and Brynjolfsson(1999) showed that the main results for bundling continue to hold even for small marginal costs.

ranged to:

$$\frac{p^* \cdot f(p^*)}{1 - F(p^*)} = 1 \tag{2}$$

For the monopolistic bundler, it turns out that her profit maximizing decision is not difficult. The bundler's job is to find the optimal price for the sum of many random variables ($S_n = \sum_{i=1}^{n} v_i$). By the law of large numbers, it is easier to find an optimal price for the sum S_n than for individual goods v_i, because the distribution variance of S_n is decreasing as n becomes large (Bakos and Brynjolfsson 1999). Without relying on asymptotic conditions, Fang and Norman (2006) derive similar results for the finite-good case.

In particular, it can be shown, with Lemma 1 and Lemma 2 in the appendix, that for non-negative random variables, the expected value of the random variable V can be written as

$$E[X] = \int_0^\infty [1 - F(x)]\, dx. \tag{3}$$

Interestingly, this expression can be linked directly to the area under the demand curve. When price is v, demand is given by $D(v) = 1 - F(v)$, so the area under the demand curve is just $\int_0^\infty D(v)dv = \int_0^\infty [1 - F(v)]dv = E[V]$.

As shown by Bakos and Brynjolfsson (1999), in equilibrium, the profit maximizing price will be set low enough so that virtually all consumers interested in any of the goods in the bundle will 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 all of the features of the applications that they do use. While there may be anti-competitive implications to this fact (Bakos and Brynjolfsson 2000, Nalebuff 2004), such bundling does give the socially desirable result of dramatically reducing the deadweight loss because very few consumers are excluded from using any of the bundled goods in equilibrium. In essence, once consumers 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 provides close-to-optimal allocation of goods to consumers (Bakos and Brynjolfsson, 1999).

However these benefits comes at a major cost. Bundling inherently destroys information about how each of the component goods are 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.

Accordingly, the remainder of the paper studies the question of how to provide the correct rewards to content providers, and thereby give them financial incentives to create content.

2 The Revenue Distribution Problem

Bundling strategies help sellers to extract more consumer surplus. If one single seller cannot provide a large enough bundle of information goods, it is worthwhile to have one content aggregator to negotiate with multiple sellers to offer a bundle of information goods from multiple sources.

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 curve proposed in the literature all fail here because bundle pricing does not automatically provide a way 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, and consumers will eventually be harmed. Second, without a good signal of consumers' preference, content providers may not produce the content that best fit the consumers' taste. Third, in any effort to overcome these problems, the content producers may force the potential bundler to adopt other strategies such as pay-per-view (the case of iTunes). However, such strategies re-introduce the deadweight loss problem discussed at the beginning of section 1.

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

2.1 Payment determined by number of downloads

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

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* gets about 29 million viewers per week, which is far more than most other TV shows; as a consequence, each of the six stars gets paid \$1.2 million per episode, which is 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 ($j = 1, ...m$), 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:

$$revenue_i = \sum_{j=1}^{m} B \cdot \frac{d_{ij}}{\sum_{k=1}^{n} d_{kj}} = R \cdot \frac{d_i}{\sum_{k=1}^{n} d_k}.$$

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.[4]

This method is powerful in the context when all the goods are approximately equal in value. 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 consumer valuation 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 (Wilbur and Zhu 2009).

[4]Since no j appears in the final term.

2.2 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 subsection 2.1, suppose each item i ($i = 1, .., n$) in the bundle also has a stand-alone price p_i.

Building on the equation from subsection 2.1, an improved way to distribute the revenue is through the following formula:

$$revenue_i = \sum_{j=1}^{m} B \cdot \frac{p_i d_{ij}}{\sum_{k=1}^{n} p_k d_{kj}} = R \cdot \frac{p_i d_i}{\sum_{k=1}^{n} p_k d_k}, \tag{4}$$

which suggests that the revenue to distribute to content provider i should be a proportion of the total revenue that 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 are rooms 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. However, this will only be true if the price paid by the marginal consumer of each goods is proportional to the average price that would be paid by all consumers of that good, for all goods.[5]

2.3 Other Mechanisms

In William Fisher's book (Fisher 2004), he explores various solutions to the music piracy problem brought about by the new peer-to-peer technology. Specifically, he proposes to

[5]Barro and Romer (1987) explore how similar proportionalities can explain a number of pricing anomalies.

replace major proportions of the copyright and encryption-based models with a "governmentally administered reward system," and he correctly points out that 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 each recording is listened to. 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, should be superior to the current practice taken by ASCAP (and BMI, SESAC, etc.) to compensate music producers, and it comes very close to our ideal of learning consumers' valuations and distribute money accordingly; but it also suffers from several problems. First, different 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, as shown in the literature, votes are not reliable because individual hidden incentives may induce voters to misrepresent their true values. In essence, the Fisher approach still does not provide a reliable, incentive-compatible way to determine the true value of each good to consumers.[6]

3 Innovation Incentives

Before moving on to propose our mechanism to solve the revenue allocation problem of bundling, we will look at another related issue in this section. We show how innovation incentives are severely limited by the traditional pricing mechanism.

In particular, we show that, contrary to common belief, the traditional price system based on excludability does not provide correct innovation incentives to producers. Our subsequently

[6]The public goods mechanism design literature seeks to provide a remedy to the voter misrepresentation problem. Specifically, the Vickrey-Clarke-Groves (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 the truth himself. Accordingly, while VCG design is intriguing in theory, it is rarely, if ever, seen in practice.

proposed couponing mechanism not only solves the revenue distribution system, but also can be a socially desirable way to promote innovation for digital goods.

Suppose a seller can invest in an innovation that improves consumers' valuations of her digital good. The investment can be in the form of improving product quality, functionality or educating users to use the product more effectively. We now discuss two types of innovations (market innovation and targeted innovation) and demonstrate why the traditional pricing mechanism offers insufficient innovation incentives.

3.1 Market innovation

When an innovation is related to improved design, faster speed, higher safety, or increased efficiency, all users benefit. We call this type of innovation as a market innovation because everyone in the market derives higher valuation with the innovation. Suppose an innovation can increase each consumer's valuation by δ, then the density function of consumer valuations will be moved to the right by δ, this is equivalent to moving the demand curve upward by δ.

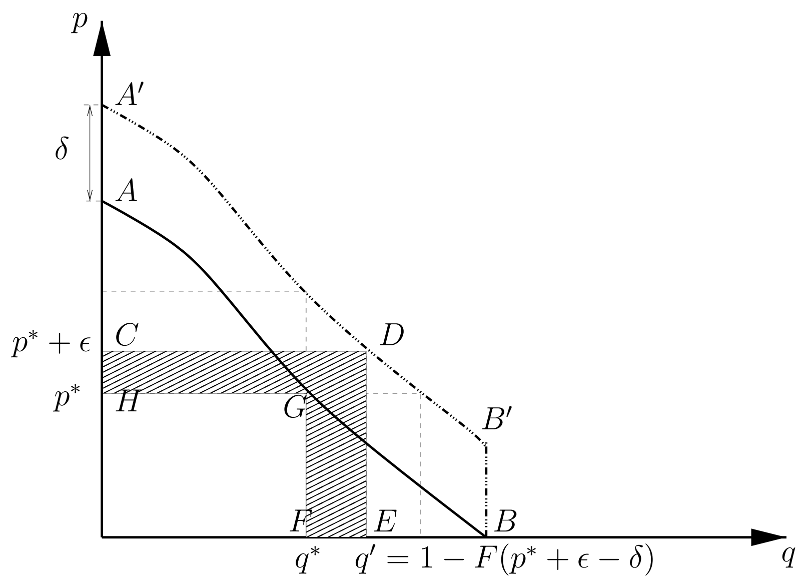

Figure 1: Market Innovation — Upward shift of demand curve

When the demand is shifted upward, the monopolistic seller will be charging a new price $p' = p^* + \epsilon$ that maximizes her profit. In Figure 1, the social value of this innovation is the area between the two demand curves (area $ABB'A'$). The reward of this innovation to the seller, however, is indicated by the shaded area (area $CDEFGH$)with the social value.

To the society, this innovation on the one hand reduces consumer surplus, but on the other hand also reduces the deadweight loss to a certain extent, so the overall social welfare effect is mixed. Depending on the shape of the demand curve, the seller's profit can be greater or less than the social value.[7]

When the demand is shifted upward, if the monopolistic seller charges a higher price of $p^*+\delta$, she will keep selling the optimal quantity q^*, or alternatively, she could keep charging the optimal price p^* and sell to more people (the demand will be $q' = 1 - F(p^* - \delta)$ now). We next show, in Lemma 3, that both strategies lead to the same expected profit for the seller.

Lemma 3: Marginally, the innovative monopolist seller can charge a higher price or enjoy a increased demand, and the two strategies are equivalent in terms of expected profit.

Lemma 3 naturally follows the optimality of p^*, we will be using this result in the next sections. From the figure, the seller should be charging a new price at $p^*+\epsilon$, with $0 < \epsilon < \delta$.

3.2 Targeted innovation

We have assumed above that the innovation can uniformly increase consumers' valuations of all types. In many situations, an innovation can only affect the valuation of a subset of consumers. For example, die-hard iPhone fans may find multi-tasking extremely desirable, but those who only use iphone as a cellphone may not find any value in multi-tasking. As a result, Apple's efforts in supporting multi-tasking can only be appreciated by a fraction of consumers. We call this type of innovation targeted innovation. With this type of innovation, only some consumers with valuation near some $\tilde{v}$ are affected.

To make concrete arguments, we model targeted innovation with two approaches. Each of these approaches are based on some technical assumptions. The insights obtained are mutually complementary.

3.2.1 Unit Mass Model

When an innovation is targeted to users with valuation of $\tilde{v}$, we assume it can change the density function of consumers' valuation by moving $\tilde{v}$ to the left to $\tilde{v} + \delta$. All users with

[7]Spence (1976, p.409) discusses a bias in pricing due to the shape of the demand curve. He illustrates a case in which a product with a low price elasticity can have a higher social value but lower profits than a product with a high price elasticity. In our model, depending on the shape of the demand curve, an innovation with certain social value may offer a too high or too low incentive for firms to pursue.

valuation in the interval $[\tilde{v},\tilde{v}+\delta]$ will also have the new valuation at $\tilde{v}+\delta$. This practically forms a unit mass at $\tilde{v}+\delta$ on the valuation density curve.

Figure 2 shows the corresponding change on the demand curve.

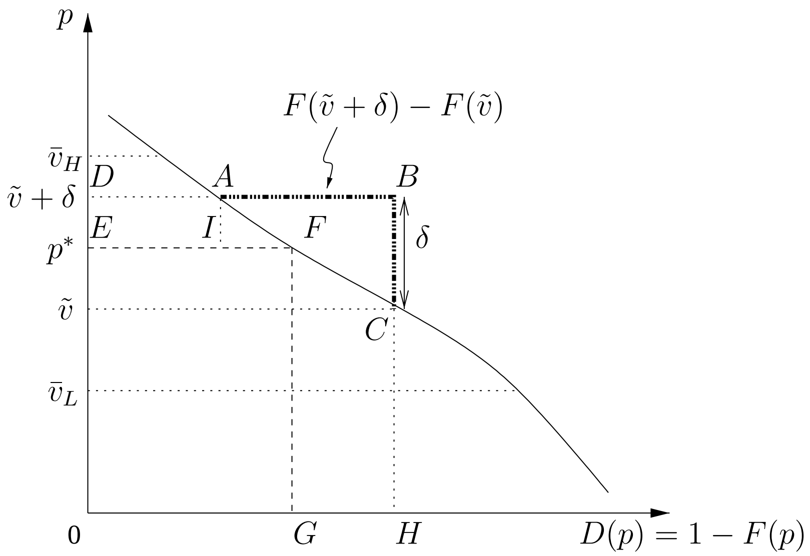

Figure 2: Targeted Innovation — Unit Mass Model.

The total social value of the good equals the area under the demand curve. If the seller makes a targeted innovation for some consumers with valuation $\tilde{v}$, the social gain of the innovation is thus denoted by the area ABC in Figure 2. When δ is small, $\Delta ABC \approx \frac{1}{2}\delta[F(\tilde{v}+\delta)-F(\tilde{v})] \approx \frac{1}{2}\delta^2 f(\tilde{v})$.

We shall need the following technical assumption to get a well-behaved demand curve.

Assumption I: $F(v)$ is twice continuously differentiable with $F(0)=1$, $F(1)=1$, $f(v)>0$ $\forall v>0$, and $\frac{1}{1-F(v)}$ is strictly convex for $v \in (0,1)$.

This assumption is implied by log-concavity of $1-F(v)$, which itself is implied by log-concavity of the density function $f(v)$. This assumption implies that the profit function $p \cdot [1-F(p)]$ is concave, and has a unique global maximum.

Log-concavity property is frequently assumed in the economics literature.[8] It is also well known that log-concavity of the density function implies the notions of IFR (increasing failure rate), and NBU (new better than used) in survival and reliability analysis literature

[8]See, for example, Laffont and Tirole (1988) in the context of games with incomplete information, Baron and Myerson (1982) in the context of theory of regulation, Myerson and Satterthwaite (1983) in the context of auction, and Johnson and Myatt (2006) in the context of informative advertising.

(Barlow and Proschan 1975). Bagnoli and Bergstrom (1989) give a good review. In our context, log-concavity is sufficient to guarantee that solutions are unique and well-behaved.

Given an innovation that increases some consumers' valuation by δ, there exists a pair of $\bar{v}$, such that the seller is indifferent between carrying out the innovation and not carrying out the innovation. For the indifferent seller,

$$(\bar{v} + \delta)\left[1 - F(\bar{v})\right] = p^* \left[1 - F(p^*)\right] \tag{5}$$

Solving for $\bar{v}$, we have two values, $\bar{v}_L$ and $\bar{v}_H$,[9] such that the optimal price $p^* \in (\bar{v}_L, \bar{v}_H)$ satisfies (5). Also, for all $v \notin [\bar{v}_L, \bar{v}_H]$, it must be that $(v + \delta)[1 - F(v)] < p^*[1 - F(p^*)]$, so the seller has no incentive at all to innovate for consumers with valuations outside the range $(\bar{v}_L, \bar{v}_H)$. For small δ, the range $(\bar{v}_L, \bar{v}_H)$ is very small, and even in this range, innovation may not be socially desirable.

For consumers with valuation in the range $(\bar{v}_L, \bar{v}_H)$, one can look at three distinct cases:

(1) The socially desirable region: $\tilde{v} \in [\bar{v}_L, p^* - \delta)$

In this case, the seller would want to charge a price $p = \tilde{v} + \delta$ and earn profit $\pi = (\tilde{v} + \delta)[1 - F(\tilde{v})]$. By Lemma 3, $\pi > (\bar{v}_L + \delta)[1 - F(\bar{v}_L)] > p^*[1 - F(p^*)]$. So the seller prefers to lower the price from p^* to $p = \tilde{v} + \delta$, and earn a higher profit. The reduction in price has two socially desirable effects. First, the consumer surplus is increased. People with valuation in the range $v \in (\tilde{v} + \delta, p^*)$ are no longer excluded from accessing the good; people with valuation in the range $(p^*, +\infty)$ can each enjoy an increased consumer surplus of $\Delta CS = p^* - (v + \delta)$. Second, deadweight loss is reduced, the change in deadweight loss is $\Delta DWL = [F(p^*) - F(\tilde{v} + \delta)](\tilde{v} + \delta)$. The reduction in deadweight loss is composed of two parts: First, for people with valuation in the range $(\tilde{v} + \delta, p^*)$, apart from the increase in consumer surplus, there is also reduction in deadweight loss due to the fact that their demand is satisfied. Second, for people with valuation in the range $(\tilde{v}, \tilde{v} + \delta)$, the innovation increases their valuation, and they are no longer excluded from purchasing the good.

(2) The socially undesirable region: $\tilde{v} \in [p^*, \bar{v}_H]$

In this case, the seller innovates for people with valuation just higher than the optimal price. By lemma 1, we know it is worthwhile for her to increase the price to $\tilde{v} + \delta$, there are two

[9] This result directly follows from the assumption of log-concave density function. Here we omit a formal proof of the existence and uniqueness of $\bar{v}_L$ and $\bar{v}_H$, which can be easily derived with the fixed point theorem.

socially undesirable effects associated with this. First, for consumers originally having a valuation above $\tilde{v} + \delta$, they each lose consumer surplus by $\Delta CS = v + \delta - p^*$. Furthermore, for people with valuation in the range $(\tilde{v}, \tilde{v} + \delta)$, although their valuation is increased due to the innovation, they no longer enjoy a surplus now. Second, for people with valuation in the range $(p^*, \tilde{v})$, they can no longer afford to buy the good now, so there is an increase in deadweight loss.

(3) The mixed region:$\tilde{v} \in [p^* - \delta, p^*)$

This case has mixed effects. People with valuation higher than $\tilde{v} + \delta$ suffer a reduction in consumer surplus by $\Delta CS = \tilde{v} + \delta - p^*$. In Figure 2, the loss is indicated by the area $ADEI$. For people with valuation in the range $(p^*, \tilde{v} + \delta)$, due to the innovation, they have a higher valuation now, but due to the increased price, they no longer enjoy a surplus (the area AIF). For people with valuation in the range $(\tilde{v}, p^*)$, their valuation is increased to $\tilde{v} + \delta$, but again, the seller gleans all the surplus due to innovation. A socially desirable side effect is that the deadweight loss is reduced because these group of people are able to use the product now. The area $FGHC$ indicates the social gain from reduced deadweight loss. In total, consumer surplus is hurt by the area $ADEF$, deadweight loss is reduced by the area $FGHC$, and the seller enjoys the extra value created by innovation indicated by area ABC.

In sum, with traditional price mechanism, the seller has too little incentive to create innovations that mainly benefit consumers with very low or very high valuations. A seller who is able to invest in targeted innovation is always putting resources to benefit the marginal consumers whose valuation is close to the monopolistic price.

To see the socially wasteful innovation incentives offered by the traditional price system, consider the following example. If a seller takes an effort to innovate and increases the valuation of some consumers from p^* to $p^* + \delta$, then her gain is $\delta[1 - F(p^*)]$. The ratio of her gain over her contribution is $incentive\_ratio_{Traditional} = \delta[1 - F(p^*)]/[\frac{1}{2}\delta^2 f(p^*)] = 2\frac{1 - F(p^*)}{\delta f(p^*)} = \frac{2p^*}{\delta}$, and $lim_{\delta \to 0} incentive\_ratio_{Traditional} = \infty$. This is a very shocking result, as shown above in case (2), the innovation for people whose valuation is just above the optimal price (the socially undesirable range) will bring about effects such as reduced consumer surplus and increased deadweight loss, yet this is exactly the range where it is most attractive for the sellers to innovate.

3.2.2 Demand Frontier Model

The unit mass model above is based on an assumption of how the valuation density function is changed by the innovation. One disadvantage of that formulation is the undesirable technical requirement of creating a unit mass on the density function. We now turn to examine the demand curve directly. We show that the traditional pricing mechanism introduces two types of social losses. First, when a firm is able to target some consumers, the target may not be chosen optimally. Second, given a targeted consumer segment, monopoly pricing is unable to realize all social gains because some consumers are priced out of the market.

For the original demand function $D(p) = 1 - F(p)$, denote the inverse demand function by $P(z)$. An innovation opportunity targeted to value $\tilde{v}$ can enhance the valuation of the consumers by a factor $\delta f(z;\tilde{q})$, where $\tilde{q}$ is the demand of consumers with valuation $\tilde{v}$.[10] The new inverse demand curve can be written as $P_{\tilde{q}}(z) = [1 + \delta f(z;\tilde{q})] \cdot P(z)$. Innovation is targeted in the sense that $f(z;\tilde{q})$ is bell shaped around $\tilde{q}$. It reaches maximal value of 1 at $z = \tilde{q}$. The decay function $f(z;\tilde{q})$ is assumed to be concave in $\tilde{q}$ and twice differentiable with respect to z and $\tilde{q}$. Furthermore, we assume that both $P(z)$ and $P_{\tilde{q}}(z)$ are well behaving inverse demand functions.

Assumption II: Both $P(z)$ and $P_{\tilde{q}}(z) = [1+\delta f(z;\tilde{q})] \cdot P(z)$ are log-concave, such that there is a unique profit maximizing production level.

Based on this setup, the social value created by the innovation can be denoted by

$$SV(\tilde{q}) = \int_0^1 \delta f(z;\tilde{q}) P(z) dz.$$

Given the magnitude of the innovation δ, the optimal target with respect to the social value creation is determined by the first order condition

$$\frac{\partial SV(\tilde{q})}{\partial \tilde{q}} = \int_0^1 \delta f_{\tilde{q}}(z;\tilde{q}) P(z) dz = 0. \tag{6}$$

Before the innovation, the monopolist charges the optimal price p^* and serve a market size of z^*. The valuation of the marginal consumer equals to the price. z^* satisfies the following first order condition

$$MR(z^*) = P(z^*) + z^* P(z^*) = 0.$$

[10] Formally, $\tilde{q} = D(\tilde{v}) = 1 - F(\tilde{v})$.

After the innovation, we have the profit function as $\Pi_{\tilde{q}}(z) = P_{\tilde{q}}(z) \cdot z = zP(z) \cdot [1 + \delta f(z; \tilde{q})]$. The new marginal revenue function is

$$MR_{\tilde{q}}(z) = MR(z) \cdot [1 + \delta f(z; \tilde{q})] + \Pi(z) \cdot [\delta f_z(z; \tilde{q})].$$

Given $\tilde{q}$, it is easy to see that the optimal level of production lies between z^* and $\tilde{q}$.[11] It is then straight-forward to have the following proposition.

Proposition 1 *If the target of the innovation could be chosen endogenously, the monopolist will choose to target at the demand level that is optimal before the innovation (i.e., $\tilde{q} = z^*$). In other words, the original marginal consumer will be targeted. The market size will not change as a result of the innovation.*

Proof. All proofs are in the appendix. ■

This result suggests that the monopolist would like to focus on the current market and seek to increase the value of the current marginal customer.

After the innovation, the new marginal consumers' valuation will be $P_{\tilde{q}}(z^*) = (1 + \delta)P(z^*)$.

To assess whether the innovation incentive based on profits is aligned with the socially optimal level, we present the following result.

Proposition 2 *If the decay function can be represented as a quadratic function of the distance between z^* and $\tilde{q}$, or if $f(z; \tilde{q}) = f[(z - \tilde{q})^2]$, then we have that a profit-driven seller always targets a segment that is not socially desirable.*

What happens if $\tilde{q}$ is exogenously given and cannot be chosen by the seller? We showed that the new equilibrium market size $z^*_{\tilde{q}}$ should locate between $\tilde{q}$ and z^*. If $\tilde{q} < z^*$, a larger market will be served as a result of the innovation. On the other hand, when $\tilde{q} > z^*$, the monopolist, to maximize profit, will increase the price so much that the original marginal consumer with value $P(z^*)$ will be priced out of the market.

[11] To see this, we just need to check that MR cannot be zero

4 The Couponing Mechanism

As discussed in Section 2, 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 derive the demand curve for each good by issuing targeted coupons to a small but statistically representative sample of consumers. Our mechanism is substantially different from the traditional use of coupons as a marketing method to price discriminate consumers. Coupons in our mechanism is 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, she could choose $m \cdot n$ representative consumers and issue each of them a coupon, where n is the number of price levels covering the range of the valuations, which we call "coupon levels" (one simple way to get these levels is to offer coupon values from $\frac{1}{n}\bar{V}$ to $\frac{n-1}{n}\bar{V}$ where $\bar{V}$ is the upper bound of consumer valuations for this good), and m is the number of coupons to be offered for each of the price levels, which we call "sample points" (there will be m consumers who can receive a coupon with face value $\frac{i}{n}\bar{V}$, $i = 1, ...n-1$). While $m \cdot n$ is large enough to make statistically valid inferences, it is nonetheless 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 an indicated component. So upon observing the consumer's action, the bundler can learn whether his valuation of the component is higher or lower than the face value of the coupon. Aggregating the m consumers' valuations will give the bundler a good estimate of demand at that price, summarizing the results for the n coupon levels, the bundler can plot a fairly accurate demand curve, and the area under the demand curve is the social valuation for the particular good. Using the same method for all the components, the bundler can learn the social valuation of each of the goods in the bundle. She can then distribute the revenue among the content providers according to their share of 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

$$revenue_i = R \frac{v_i}{\sum_{j=1}^{N} v_j} \tag{7}$$

where N is the total number of content providers.

This method compares favorably to the traditional price mechanism. The traditional price mechanism subjects 100% of consumers to the inefficiency of positive prices. However, only data from a small fraction of consumers are needed to get extremely 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%. 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 3, 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 couponing 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 couponing mechanism is superior to the VCG mechanism in several ways. First, 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. Second, in the VCG mechanism, 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 (as their payments are determined by all other bids). Furthermore, coalitions of consumers can game the VCG to their advantage. In contrast, in the couponing mechanism, the effects on others from a consumer's misreport are minimal. Third, the couponing mechanism is fully budget balancing, unlike the VCG. Finally, the couponing mechanism is more intuitive than the VCG for real world problems.

The following proposition asserts that the Couponing Mechanism indeed gives us correct demand curve estimations in expectation.

Proposition 3 *For any one of the components in the bundle, given a large number of randomly chosen respondents and levels of coupons, the above mechanism gives an empirical demand function* $\hat{D}(p) = 1 - \hat{F}_V(p)$ *that arbitrarily approximates the true demand function* $D(p) = 1 - F_V(p)$.

Proposition 3 gives an asymptotic result, we run simulations to see the effectiveness of this mechanism.

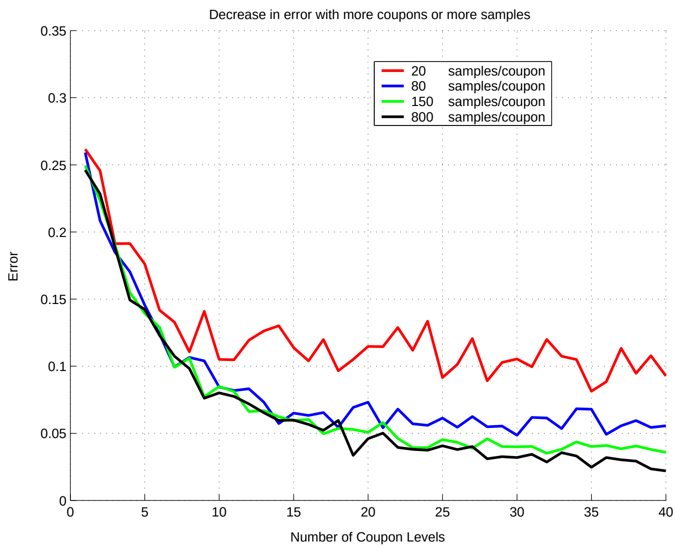

Figure 3: Simulation Results for the Couponing Mechanism

The use of the couponing mechanism gives us empirical estimates of the inverse demand curves for various valuation 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 3 shows the result of the couponing mechanism applied to the uniform distribution (other distributions yield qualitatively similar results). We see that error rate is declining with more coupon levels and with more sample points for each coupon value. It is remarkable that with just 20 coupon levels, the error rate can be as low as 5%. 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 10%, and sampling 80 consumers improves the error rate to be close to 5%. From the error rate curves, we can also see that when sampling 20 consumers, adding coupon levels more than 10 does not improve the precision significantly; similarly, when sampling 80 consumers, adding coupon levels more than 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 0 more quickly/slowly for fatter/thinner demand curves (the ones with a higher/lower expected value). In our simulations, for some demand curves, with just 5 coupon levels and 20 sample points (for mere 100 respondents), the coupon mechanism can give us an error rate below 0.1%. Thus, sampling just 100 consumers can provide almost as accurate an estimate of demand as sampling all the consumers of the good, which could be in the millions.

The deadweight loss is proportionately smaller, too. Consumers who cash-in the coupon 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. However, in contrast to the traditional pricing approach, the couponing mechanism only subjects a fraction of consumers to this choice, so only a fraction choose not to buy, and the total deadweight loss is a fraction at large.

This mechanism can be used to solve the revenue distribution problem discussed in section 2, and we will show next, with a few propositions, that this mechanism can also help to avoid the innovation incentive issues arising in traditional price systems.

Consider market innovation introduced in section 3.1. When the demand is shifted upward, the seller can get paid virtually the full amount of the extra valuation it created for the consumers. Let the original profit be $\pi = E[V] = \int_0^\infty [1 - F(v)]dv$, she can now earn $\pi' = \delta + \pi$. We will show in the next Proposition that the seller' innovation incentive in the bundling+couponing scheme is higher than that in the traditional market.

Proposition 4 *If an innovation can increase consumers' valuations uniformly higher, the proposed couponing mechanism gives the producer strictly greater incentives of innovation than does the traditional pricing mechanism.*

For targeted innovation introduced in section 3.2, it is obvious that, with couponing, the seller does not care how high or how low the targeted consumers' valuation is because she is paid according to the area under the demand curve. Combining bundling with couponing can provide balanced incentives for innovations targeted to any value, leading the developer to pursue any innovations whose expected benefits exceed expected costs.

In Figure 2, no matter where $\tilde{v}$ is , the reward to the seller is the area ABC, so she will not discriminate against consumers with low or high valuations. This brings us to Proposition 5.

Proposition 5 *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.*

Similarly, for the case of the proposed mechanism, the ratio of the expected return over the social contribution is $incentive\_ratio_{Bundling} = \frac{1}{2}\delta^2 f(\tilde{v}) / \left(\frac{1}{2}\delta^2 f(\tilde{v})\right) = 1$, which is fair. So we have the following proposition:

Proposition 6 *The traditional market 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.*

5 Discussion

Throughout this paper, we assumed the more general case that the demand curves of different goods look different. If the demand curves are all the same, or at least all parallel to each other, there can be easier mechanisms to distribute the revenue while ensuring to keep the innovation incentives of producers. When the goods all have similar social values (the areas under the demand curves are the same), Equation (7) becomes $\text{revenue}_i = R\frac{v_i}{\sum_{j=1}^{N} v_j} = R\frac{q_i v}{\sum_{j=1}^{N} q_j v} = R\frac{q_i}{\sum_{j=1}^{N} q_j}$, where q_i denotes the number of times that good i is consumed, and the payment to content provider i is solely determined by the number of downloads. Interestingly, even if the social values are not similar, as long as the demand curves of different goods have similar shapes, we can still use the number of downloads as a sufficient statistic to derive the correct revenue distribution rule. For example, imagine the simplest case that we have linear demands with slope k; different goods have different intercepts, but they are all parallel to each other. Once we observe the total quantity consumed for each good, we know the social value created by this good is just a quadratic function of this number. So, in the spirit of Equation (7), we have:

$$\text{revenue}_i = R\frac{v_i}{\sum_{j=1}^{N} v_j} = R\frac{q_i^2}{\sum_{j=1}^{N} q_j^2}.$$

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 more 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 a Forrester report, maybe only a dozen 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. It turns out that both issues can be easily addressed by a practice we call "passive couponing". Under the "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 ruled out. 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.[12]

In previous sections, we avoided the issue of duration of contracts. It is likely to be unneces-

[12]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 3. It can also be easily shown that there is no incentive for the consumer to mis-report his value in each period.

sary to permanently block from 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 sacrifice 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 downloading Madonna for the next whole year").[13]

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. For complements, we overestimate total value of the goods. First of all, non-independence will only affect the estimated share of contributions of each content provider, so the payment to each of the creators will be changed, but the benefit of innovation incentives will not be affected. Second, 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.

6 Conclusion

Revolutionary technologies often engender innovations in business organization. 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 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.

[13]The mechanism proposed here may not work as well with information goods whose value is time-invariant (e.g. Symphony No. 9 by Beethoven). Depending on the nature of the DRM system used, once someone downloads a copy of the work, there may be no point in offering coupons because the consumer might not need to download any more copies in the future. In this case, goods with time-invariant value may not be suitable for sale in a bundle.

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 mechanism shows particular promise. We find that our approach can provide substantially better incentives for innovation than even the a 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,

- has orders of magnitude less inefficiency 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 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, works better for larger numbers),

Our approach also has weaknesses and challenges. First of all, since massive bundling is a component of the mechanism, our approach only works for digital goods. As long as the marginal cost is not close to zero, the benefit of using bundling to ensure social efficiency will be non-existent. Second, this mechanism may not work for all information goods. In order to have the mechanism useful, the contents must be updated regularly and consumers should have a good estimate of the expected value of the future contents. If consumers can subscribe once and download all contents and sign off the service, the mechanism will be useless[14]. 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 (Fichman and Kemerer 1999). 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 which regularly collect and distributed 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 (Kremer 1998). However, it should be stressed, that our mechanism does not require any government role since all of the participants (consumers, content creators, bundlers) have incentives to participate completely voluntarily. This stands in contrasts 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.

We expect that the next 10 years will witness a scale of organizational innovation for creating and distributing digital goods surpassing even the remarkable pace of the last 10 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, "cookbooks", repositories and analyses.

[14]Modern digital rights management technology may help to alleviate the problem. Some services allow the subscribers to download the music but require them to log on at least once per month to varify the status of subscription, if a subscription gets expired, the content will no longer be accessed by the consumer.

References

Aghion, P., Bolton, P., Harris, C. and Jullien, B.: 1991, Optimal learning by experimentation, *Review of Economic Studies* **58**, 621–654.

Bagnoli, M. and Bergstrom, T.: 1989, Log-concave probability and its applications. memeo.

Bakos, J. and Brynjolfsson, E.: 1999, Bundling information goods: Pricing, profits and efficiency, *Management Science* **45**(12), 1613–1630.

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

Barlow, R. and Proschan, F.: 1975, *Statistical Theory of Reliability and Life Testing, Probability Models*, Holt, Rinehart, and Winston, New York, NY.

Baron, D. and Myerson, R.: 1982, Regulating a monopoly with unknown costs, *Econometrica* **50**, 911–930.

Barro, R. and Romer, P.: 1987, Ski-lift pricing, with applications to labor and other markets, *American Economic Review* **77**, 875–890.

Fang, H. and Norman, P.: 2006, To bundle or not to bundle, *Rand Journal of Economics* **37**, 946–963.

Fichman, R. and Kemerer, C.: 1999, The illusory diffusion of innovation: An examination of assimilation gaps, *Information Systems Research* **10**(3), 255–275.

Fisher, W.: 2004, *Promises to Keep: Technology, Law, and the Future of Entertainment*, Stanford University Press.

Geng, X., Stinchcombe, M. B. and Whinston, A. B.: 2005, Bundling information goods of decreasing value, *Management Science* **51**, 662–667.

Jain, D. C., Mahajan, V. and Muller, E.: 1995, An approach for determining optimal product sampling for the diffusion of a new product, *The Journal of Product Innovation Management* **12**, 124–135.

Johnson, J. P. and Myatt, D. P.: 2006, On the simple economics of advertising, marketing, and product design, *American Economic Review* **96**, 756–784.

Kremer, M.: 1998, Patent buyouts: A mechanism for encouraging innovation, *Quarterly Journal of Economics* **113**(4), 1137–1167.

Laffont, J.-J. and Tirole, J.: 1988, Dynamics of incentive contracts, *Econometrica* **55**(5), 1153–1175.

Lessig, L.: 2004, *Free Culture: How Big Media Uses Technology and the Law to Lock Down Creativity*, Penguin Press.

Lyman, P. and Varian, H. R.: 2004, How much information?, School of Information Management and Systems, University of California Berkeley. Retrieved from http://www.sims.berkeley.edu/research/projects/how-much-info-2003/ on April 27, 2004.

McAfee, R., McMillan, J. and Whinston, M.: 1989, Multiproduct monopoly, commodity bundling, and correlation of values., *Quarterly Journal of Economics* **104**, 371–383.

Myerson, R. and Satterthwaite, M.: 1983, Efficient mechanisms for bilateral trading, *Journal of Economic Theory* **28**, 265–281.

Nalebuff, B.: 2004, Bundling as an entry barrier, *Quarterly Journal of Economics* **119**, 159–187.

Sorensen, C. and Snis, U.: 2001, Innovation through knowledge codification, *Journal of Information Technology* **16**, 83–97.

Spence, M.: 1976, Product differentiation and welfare, *American Economic Review* **66**(2), Papers and Proceedings of the Eighty–Eighth Annual Meeting of the American Economic Association 407–414.

Wilbur, K. C. and Zhu, Y.: 2009, Click fraud, *Marketing Science* **28**(2), 293–308.

Appendix: Proofs

Proof of Proposition 1: Post-innovation maximum profit is a function of the target $\tilde{q}$, $\Pi^*(\tilde{q})$. The optimal innovation target $\tilde{q}^*$ satisfies

$$\frac{d\Pi^*(\tilde{q})}{d\tilde{q}} = \frac{\partial \Pi(z;\tilde{q})}{\partial \tilde{q}}|_{z=z^*(\tilde{q})} = \delta\Pi\left(z^*(\tilde{q});\tilde{q}\right) f_{\tilde{q}}\left(z^*(\tilde{q});\tilde{q}\right) = 0$$

by the envelop theorem. Thus $z^*(\tilde{q}^*) = \tilde{q}^* \Rightarrow \tilde{q}^* = z^*$.

Q.E.D

Proof of Proposition 2: Due to log-concavity of the inverse demand function, we only need to consider the value of the first order derivative at z^*, we have

$$\begin{aligned}
\frac{dSV(z^*)}{d\tilde{q}} &= \int_0^1 \delta\frac{\partial f}{\partial \tilde{q}}(z;z^*)P(z)dz \\
&= \int_0^1 -2\delta P(z)f'((z-z^*)^2)(z-z^*)dz \\
&= \int_0^1 -\delta P(z)df((z-z^*)^2) \\
&= -\delta P(z)f((z-z^*)^2)|_0^1 + \int_0^1 \delta f((z-z^*)^2)dP(z) \\
&= -P(0)f(z^{*2}) + \int_0^1 \delta f((z-z^*)^2)dP(z) < 0
\end{aligned}$$

This means that the social value created is decreasing at z^*. Thus, z^* can never be the socially optimal segment to target.

Q.E.D

Proof of Proposition 3:

We prove proposition 3 in two steps. First, we show that for each price level, the mechanism offers a consistent estimate of the true demand at that level. Second, we show given enough price levels, the demand curve can be arbitrarily closely approximated.

For one particular component in the bundle, the seller first chooses the number n of coupon levels, then, for each coupon level, sends coupons to m randomly chosen consumers. For a coupon with face value $\tilde{v}$ for the component, the respondent will take it only if he has a valuation lower than $\tilde{v}$. The probability of the coupon getting accepted is $Prob(V \leq \tilde{v}) = F_V(\tilde{v})$.

We now define indicator variables $Y_1, ..., Y_m$ where Y_i is 1 if the coupon with face value $\tilde{v}$ is accepted by the i^{th} consumer, and Y_i is 0 if otherwise. We have $Y_k = \begin{cases} 1 & if\ X_k \leq \tilde{v} \\ 0 & if\ X_k > \tilde{v} \end{cases}$, where $k = 1, ..., m$. Note that $Prob(Y_k = 1) = Prob(X \leq \tilde{v}) = F_V(\tilde{v})$, and $Prob(Y_k = 0) = Prob(X > \tilde{v}) = 1 - F_V(\tilde{v})$. For all the m people to whom we sent coupon $\tilde{v}$, we know the number of acceptance is $a_m = \sum_{j=1}^{m} Y_j$. Define $\hat{F}_V(\tilde{v}) = \dfrac{a_m}{m}$ as the empirical cdf at $\tilde{v}$, which gives the result of the experiments telling us what percentage of people accepts the coupon $\tilde{v}$. We can show the expected value of the empirical cdf is the true unknown cdf. $E[\hat{F}_v(\tilde{v})] = E[\dfrac{a_m}{m}] = \dfrac{E[a_m]}{m} = \dfrac{m \cdot E[Y]}{m} = E[Y] = 0 \cdot Prob(Y = 0) + 1 \cdot Prob(Y = 1) = F_V(\tilde{v})$. That completes the step 1.

Next consider the interval between any neighboring coupon's value levels. For explanatory purpose, we now assume that the seller sets equi-distance intervals on the value range [0,1], that is, the coupon values are $0, \dfrac{1}{n}, ...\dfrac{n-1}{n}$. Our result does not rely on this assumption, it holds as long as the distances are all weakly shrinking when adding more coupon levels.

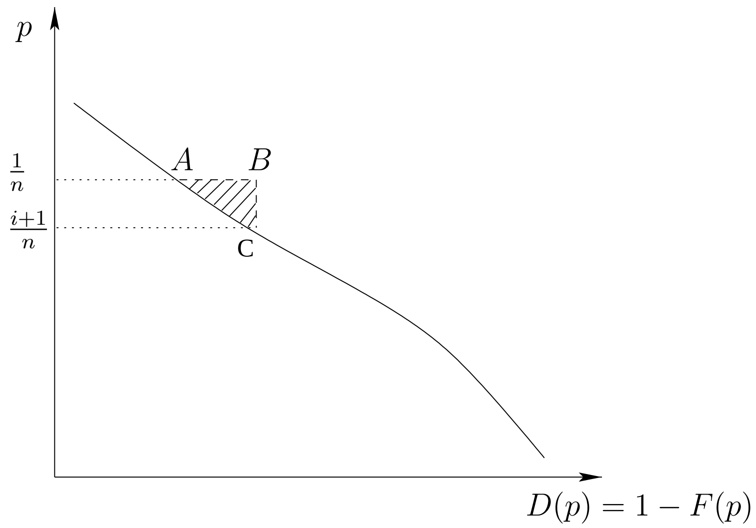

Figure A1. The upper bound of error in estimating demand

For neighboring coupon levels $\frac{i}{n}$ and $\frac{i+1}{n}$, the seller may estimate points A and C from step 1. She can simply connect the estimated points to approximate the demand curve between the two points. Since the demand curve is monotonically decreasing from 1 to 0, when estimating the area below the demand curve, the triangle ABC is the upper bound for the error. The area of ABC is $\Delta ABC = \dfrac{1}{2}(\dfrac{i+1}{n} - \dfrac{i}{n})[\hat{F}(\dfrac{i}{n}) - \hat{F}(\dfrac{i+1}{n})]$. We know $\hat{F}(\dfrac{i}{n}) - \hat{F}(\dfrac{i+1}{n}) \leq 1$, and given the assumption that $F_V(x)$ is continuously differentiable. We have $lim_{n\to\infty}(\hat{F}(\frac{i}{n})-\hat{F}(\frac{i+1}{n}))=0$, so we have $lim_{n\to\infty}\Delta ABC=\frac{1}{2}\left(lim_{n\to\infty}\frac{1}{n}\right)\cdot\left(lim_{n\to\infty}(\hat{F}(\frac{i}{n})-\hat{F}(\frac{i+1}{n}))\right)=0$, which suggests that when n is large enough, the error in estimation will converge to 0.

Q.E.D

Proof of Proposition 4:

In Figure 1, the demand curve is moved upward by δ, we need to show that the area between the two demand curves is larger than the area $CDEFGH$. We first show that the new optimal price can not be out of the range $(p^*, p^*+\delta)$. Suppose, for contradiction, that charging a price $\tilde{p} > p^*+\delta$ gives a higher profit than charging p^*(or equivalently $p^*+\delta$, due to Lemma 3), then mapping this back to the original demand curve tells us that charging a little bit higher than p^* can give us a higher profit, which can not be true since p^* is the optimal price in the original demand curve. Using the same argument, we can show that the new optimal price can not be lower than p^*. So we have, for the new optimal price, $p' = p^*+\epsilon \in (p^*, p^*+\delta)$, or equivalently, $0<\epsilon<\delta$.

Next, we only need to show that the increased profit from the traditional price mechanism is lower than what the couponing mechanism can provide.

The area between the demand curve is

$$A'B'BA = \int_0^\infty [1-F(p-\delta)]dp - \int_0^\infty [1-F(p)]dp = \delta + o(\delta)$$

where $o(\delta)$ is defined as $\text{Lim}_{\delta\to 0}\frac{o(\delta)}{\delta}=0$. The area $CDEFGH$ can be calculated as

$$\begin{aligned}
CDEFGH &= (p^*+\epsilon)[1-F(p^*+\epsilon-\delta)] - p^*[1-F(p^*)] \\
&= p'[1-F(p'-\delta)] - p^*[1-F(p^*)] \\
&= (p'-p^*)[1-F(p^*)] + p^*[F(p^*)-F(p'-\delta)] + (p'-p^*)[F(p^*)-F(p'-\delta)] \\
&= \epsilon[1-F(p^*)] + (p^*+\epsilon)[F(p^*)-F(p^*+\epsilon-\delta)] \\
&= \epsilon - \epsilon F(p^*) + p^*F(p^*) + \epsilon F(p^*) - p^*F(p^*+\epsilon-\delta) - \epsilon F(p^*+\epsilon-\delta) \\
&= p^*\left[F(p^*)-F(p^*+\epsilon-\delta)\right] + \epsilon\left[1-F(p^*+\epsilon-\delta)\right] \\
&= p^*I + \epsilon J
\end{aligned}$$

where $I \equiv F(p^*) - F(p^* + \epsilon - \delta)$,and $J \equiv 1 - F(p^* + \epsilon - \delta)$.

Since $0 < \epsilon < \delta$, we have $\text{Lim}_{\delta\to 0}\epsilon = 0$, $\text{Lim}_{\delta\to 0}\frac{\epsilon^2}{\delta} = 0$ and $\text{Lim}_{\delta\to 0}(\delta - \epsilon) = 0$, so

$$\text{Lim}_{\delta\to 0}\frac{I}{\delta - \epsilon} = \text{Lim}_{\delta\to 0}\frac{F(p^*) - F(p^* + \epsilon - \delta)}{\delta - \epsilon} = f(p^*).$$

Since $p' = p^* + \epsilon$ is the new optimal price, it must satisfy the optimal condition given in equation (2) , so we must have:

$$\frac{p' f(p')}{1 - F(p' - \delta)} = 1$$

where the term $F(p' - \delta)$ corresponds to the shifted demand curve.

Substituting for J, we have $J = 1-F(p^*+\epsilon-\delta) = 1-F(p'-\delta) = p'f(p') = (p^*+\epsilon)f(p^*+\epsilon-\delta)$. By continuity, we also know that $\text{Lim}_{\delta\to 0}f(p^* + \epsilon - \delta) = f(p^*)$,so we can write

$$\begin{aligned}
\lim_{\delta\to 0}\frac{CDEFGH}{\delta} &= \lim_{\delta\to 0}\frac{p^* I + \epsilon J}{\delta} \\
&= \lim_{\delta\to 0}\frac{p^*[F(p^*) - F(p^* + \epsilon - \delta)] + \epsilon[1 - F(p^* + \epsilon - \delta)]}{\delta} \\
&= \lim_{\delta\to 0}\frac{p^*(\delta - \epsilon)f(p^*) + \epsilon(p^* + \epsilon)f(p^*)}{\delta} \\
&= \lim_{\delta\to 0}\frac{p^*\delta f(p^*) + \epsilon^2 f(p^*)}{\delta} \\
&= p^* f(p^*)
\end{aligned}$$

It is now obvious that

$\text{Lim}_{\delta\to 0}\frac{CDEFGH}{\delta} = p^* f(p^*) = 1 - F(p^*) < 1 = \frac{A'B'BA}{\delta} = \text{Lim}_{\delta\to 0}\frac{A'B'BA}{\delta}$. This completes our proof that the area $CDEFGH$ is smaller than the area between the two demand curves.

Lemma 1. Given any nonnegative random variable Y with finite mean(i.e. a random variable for which $F_Y(y) = 0$ for $y < 0$, and $E[Y] < \infty$), $lim_{y\to\infty} yP(Y \geq y) = 0$.

Proof of Lemma 1: First note that, $yP(Y \geq y) = y(F_Y(\infty) - F_y(y)) = y\int_y^\infty dF_Y(z) \leq \int_y^\infty z dF_Y(z)$, the last inequality is due to the fact that y is a lower bound for all z when z goes from y to infinity.

Next, from the definition of mean, we know $E[Y] = \int_0^\infty zdF_Y(z) = \int_0^y zdF_Y(z) + \int_y^\infty zdF_Y(z)$. Taking the limit, we have

$$lim_{y\to\infty} \int_y^\infty zdF_Y(z) = lim_{y\to\infty}(E[Y] - \int_0^y zdF_Y(z)) = E[Y] - E[Y] = 0 \qquad (8)$$

So we have $lim_{y\to\infty} yP(Y \geq y) \leq lim_{y\to\infty} \int_y^\infty zdF_Y(z) = 0$.

Q.E.D

The above lemma enables us to write the expected value of a random variable in a very enlightening way.

Lemma 2. For a nonnegative random variable X, $E[X] < \infty$, the expectation can be written in the following form:

$$E[X] = \int_0^\infty [1 - F_X(x)]\, dx \qquad (9)$$

Proof of Lemma 2: By definition,

$$E[X] = \int_0^\infty x f_x(x)dx = \int_0^\infty xdF_x(x)$$

we now define $G_X(x) = 1 - F_X(x)$, then we have $E[X] = \int_0^\infty xd(1 - G_X(x)) = -\int_0^\infty xdG_x(x)$. Using integration by parts:

$$\begin{aligned} -\int_0^\infty xdG_X(x) &= -xG_X(x)|_0^\infty + \int_0^\infty G_X(x)dx \\ &= -(lim_{x\to\infty} xG_X(x) - 0 \cdot G_X(0)) + \int_0^\infty G_X(x)dx \\ &= \int_0^\infty G_x(x)dx = \int_0^\infty [1 - F_X(x)]dx \end{aligned}$$

Q.E.D

Proof of Lemma 3: We need to compare $\pi_p = (p^* + \delta)[1 - F(p^*)]$ and $\pi_q = p^*[1 - F(p^* - \delta)]$, and show that as $\delta \to 0$, they are equal.

Equivalently we need to show: $\lim_{\delta\to 0}(p^* + \delta)[1 - F(p^*)] = \lim_{\delta\to 0} p^*[1 - F(p^* - \delta)]$, which is

$\lim_{\delta \to 0} \frac{F(p^*)-F(p^*-\delta)}{\delta} = \frac{1-F(p^*)}{p^*} \Leftrightarrow f(p^*) = \frac{1-F(p^*)}{p^*}$, which is true due to equation (2), the optimality condition.

Q.E.D