Shared Information Goods

Yannis Bakos
New York University & MIT
bakos@mit.edu

Erik Brynjolfsson
Stanford University & MIT
erikb@mit.edu

Douglas Lichtman
University of Chicago
dgl@uchicago.edu

Abstract*

Once purchased, information goods are often shared among groups of consumers. Computer software, for example, can be duplicated and passed from one user to the next. Journal articles can be copied. Music can be dubbed. In this paper, we ask whether these various forms of sharing undermine seller profit. We compare profitability under the assumption that information goods are used only by their direct purchasers, with profitability under the more realistic assumption that information goods are sometimes shared within small social communities. We reach several surprising conclusions. We find, for example, that under certain circumstances sharing will markedly *increase* profit even if sharing is inefficient in the sense that it is more expensive for consumers to distribute the good via sharing that it would be for the producer to simply produce additional units. Conversely, we find that sharing can markedly *decrease* profit even where sharing reduces net distribution costs. These results contrast with much of the prior literature on small-scale sharing, but are consistent with results obtained in related work on the topic of commodity bundling.

I. Introduction

In an influential article published a decade ago,[1] Stan Besen and Sheila Kirby investigated the economic effects of "small-scale, decentralized reproduction of intellectual property" -- the types of information sharing that take place every time a consumer pirates a computer program, dubs a music CD, duplicates a journal article, or otherwise shares access to a purchased information good. Content producers had long claimed significant economic harm from these types of small-scale sharing[2] and had even "occasionally succeeded" in having legislation introduced that would compensate for its purported effects.[3] Yet the profit implications of small-scale sharing were poorly understood. So Besen and Kirby set out to model this phenomenon as an important first step toward evaluating content producers' claims, their business strategies, and the various legal responses.

Like most of the work that preceded it[4] -- and many of the papers since[5] -- Besen and Kirby's model focused on the relationship between consumers' marginal cost of sharing and sellers' marginal cost of producing original units. They argued that a critical determinant of producer welfare is the relationship between these two types of marginal cost. Specifically, they suggested that where consumers can distribute an information good via sharing more cheaply than its producer can distribute it via the production of additional original units, sharing will tend to increase seller profit; but where sharing is more expensive, seller profit will typically diminish.[6]

An example helps to illustrate their insight. All other things held equal, if consumers can pass a magazine from one reader to the next more cheaply than its publisher can print, package, and mail the appropriate number of individual copies, overall efficiency will increase. The savings, Besen and Kirby reasonably suggested, will make both producers and consumers better off.[7] If, however, consumers incur greater costs in sharing the magazine than its publisher would have incurred had he simply produced additional copies, efficiency and profit will both decrease.

In the ten years that have passed since Besen and Kirby's important contribution, small-scale sharing has continued to be an issue of commercial, political, and scholarly import. Content producers have continued to claim that sharing devastates profit.[8] Legislators have continued to propose and enact protective legislation.[9] And a long line of academic scholarship has continued to develop Besen and Kirby's core insight as to the relationship between production costs and profitability.[10]

Technology, however, has changed in these intervening years -- and that change has important and as-yet-unrecognized implications for this line of scholarship. Specifically, in many settings, technology is reducing both the marginal cost of producing original units and the

inconvenience costs of small-scale sharing to near-zero levels. This is true largely because information goods are increasingly available in digital form. They are composed of bits, not atoms -- and bits can be quickly, accurately, and inexpensively duplicated by consumers and producers alike.

The change, of course, does not invalidate Besen and Kirby's insight. Consumers and producers still likely face *different* marginal costs, and the difference still surely influences seller profit. Our point here is only that modern technology considerably reduces these costs -- a shift that has led us to wonder whether other factors might today be correspondingly more important to the question Besen and Kirby first raised.

In this paper, we therefore ask once again how various forms of small-scale sharing affect seller profit. We compare profitability under the assumption that information goods are used only by their direct purchasers, with profitability under the more realistic assumption that information goods are sometimes shared within small social communities. In contrast to prior work, however, we assume that the seller's marginal costs of original production and consumers' marginal costs of sharing are negligibly low. We find that two previously unexplored factors significantly determine sharing's effect on seller profit: an "aggregation effect" that tends to increase profit, and a "team diversity effect" that tends to diminish it. Our contribution is to identify these factors, integrate them into the prior literature, and explore some of their implications.

Like all of the prior work, our analysis begins with two competing intuitions. The first and more familiar is the idea that sharing harms producers by decreasing the number of original information goods sold. Some consumers who receive the information good through sharing, after all, would have purchased the good were sharing not an option. The second and more interesting idea, however, is that sharing confers a corresponding benefit: consumers are likely willing to pay more for information goods that they can then share and trade with others.[11]

To quantify these competing intuitions, we have to make some general assumption as to how consumers team together to share information goods. Two simple examples help to establish the range of possibilities:

> ***Example 1***: A video store owner buys a videotape that he expects to rent to numerous customers. Video store owners compete on price; video renters shop for the best deal. The competition ultimately serves to establish some market-clearing price for video rentals, and in the process matches videos to sets of renters. When buying a new video, the store owner will take into account its subsequent value to others in the rental market.

Example 2: Families purchase and share access to cable television. The head of the household first tries to estimate each family member's valuation, and then enters the marketplace willing to pay up to approximately that sum. This is true even where he or she expects to bear the full cost individually, never asking other family members to "ante up" and contribute their fair share.

In example 1, market forces determine with whom and under what conditions renters share videos. Price, for example, is set at the intersection of supply and demand. The market works because renters are free to rent from any video store owner and hence can haggle for the best deal. They do not care with whom they share a given video so long as they enjoy access. Example 2, by contrast, illustrates a different type of sharing. Here, sharing groups are defined by pre-existing social relationships. There is no market price; family members do not switch from one family to the next; and the family's willingness-to-pay simply reflects the sum of individual family member's reservation prices.

Market-mediated and family-based sharing are obviously two extreme cases; sharing often occurs in intermediate settings. Friends, for example, feel not only market pressure to maximize personal welfare, but also the pleasant constraints of social intimacy. Consider two friends engaged in a long-term pattern of buying and sharing music CDs. Acting out of self-interest, each friend will purchase only CDs that he himself values. At the same time, however, each will likely also account for the other's preferences when making any purchasing decision; one purpose of the purchase, after all, is to thank his friend for prior purchases by purchasing a new CD that the friend, too, will value.

In their paper, Besen and Kirby explicitly adopted a market framework similar to that set forth in example 1. [12] Most of the related papers have followed suit.[13] However, many of the examples that first motivated Besen and Kirby's work, and continue to motivate ours, seem better described by a predominantly social model. In this paper we therefore adopt the latter approach.

The resulting shift in emphasis has three important implications. First, unlike prior scholarship, in our work we allow each information good to be shared by a different number of consumers. Markets tend to establish an equilibrium "team size" determined by the relative costs and benefits of adding new team members.[14] In prior work, as a result, all sharing groups were assumed to be of the same size.[15] In our model – and, we believe, in many common forms of small-scale sharing -- team size is determined exogenously by social relationships. Friends are not left out just because their inclusion would increase costs, nor are strangers invited to share solely for the

purpose of further amortizing expenses. The model thus explicitly considers how heterogeneity in team size affects seller profit.

Second, the market-based models assume that no high-valuing consumer is ever denied access to a good enjoyed by a low-valuing consumer. This is true because, as Hal Varian succinctly explains, "[i]f this were not the case, one of the members of a [team] that didn't purchase the [good] would be willing to switch places with a member of a [team] that did purchase [the good], and pay the appropriate compensation."[16] In our model, team members do not trade places. Thus not all consumers with high values will have access to the good nor will all low-valuing consumers be excluded from access. For example, a casual music fan might happen to know a music aficionado and thus enjoy access to a broad music collection even if some other individual would have been willing to pay more for such access.

Third and most importantly, implicit in the market models is the assumption that every consumer who enjoys access to a shared good pays the market price for that opportunity. Consumers in these models never contribute unevenly toward group consumption; there is a single market-clearing price and everyone must pay it, either in cash or in kind. In reality, however, friends and family members often share information goods unevenly. The music aficionado, for example, might offer his friend access to ten CDs for every one he receives in return. Similarly, he might offer to pay for more than half of some shared music purchase. Unlike the market-based models, our approach allows for this type of diversity.

In the remainder of this article we show that, when consumers engage in social sharing of the type described above, and when the marginal costs of original production and sharing are both assumed to be negligible, sharing will at times substantially increase, but can also markedly diminish, producer profit. More formally, we analyze a basic setting in which shared goods are no more or less expensive to produce than are unshared originals; shared goods are no more or less valuable to consumers than are unshared originals; goods are shared in predetermined "social" teams of small size; and those teams are willing to pay up to the sum of what each team member would have been willing to pay individually. In this setting, we find that:

(1) all other things being equal, when goods are shared in teams of a constant size, sharing will almost always *increase* seller profit as compared to the profit earned in the absence of consumer sharing;

(2) when goods are shared in small teams of varied size, sharing will tend to decrease profit when the diversity in team size is greater than the diversity in individual consumer valuations;

(3) conversely, when goods are shared in small teams of varied size, sharing will tend to increase seller profit when the diversity in team size is less than the diversity in individual consumer valuations; and

(4) seller profit under sharing can be enhanced by a negative correlation between team member valuations (as where high-valuing consumers tend to share with low-valuing consumers), and also by a negative correlation between team size and team member valuations (as where low-valuing consumers share in large teams but high-valuing consumers purchase individually or in small teams).

Furthermore, when we extend the model to consider the possibility that shared goods might be inferior to unshared originals (say, because the good degrades, or because sharing imposes non-zero coordination and transaction costs), we find that:

(5) seller profit will sometimes *increase* as the value of the shared good diminishes.

These results can be explained by the interplay of two factors that have heretofore gone unexplored in the sharing literature. We term these factors the "aggregation effect" and the "team diversity effect." The intuition for the aggregation effect is that, in many situations, a team's valuation for a good has a probability distribution with lower variance than the distribution associated with individual members' valuations for that same good. For instance, individual family members might value Corel's WordPerfect at disparate and unpredictable levels. In some households, the mother might value it highly. In others, the high-valuer might be the father or a teenager. If Corel wanted to maximize profit while selling individual copies, it would have to distinguish between these high- and low-valuing consumers and then set prices accordingly. Typically, this is difficult or impossible, and so low-valuing consumers tend to be priced out of the market while high-valuing consumers generally retain a surplus even after paying the market price.

Were these individuals to team together and purchase WordPerfect as family units, however, the seller's problem would often be significantly reduced. Even without knowing which family members were of which type, Corel would likely be able to correctly guess that most families (say) are comprised of one or two high-valuing consumers and several other consumers who value the good at a low level. Corel would be better able to set an appropriate price -- and hence better able to

extract surplus from the market -- all without ever having to specifically identify the high- and low-valuers.

Phrased another way, under reasonable assumptions about the distribution of valuations in the original demand curve, team formation makes consumer valuations more predictable. Aberrant valuations from the original curve are dampened through combination with more middling values, a process that concentrates demand and makes it easier for the seller to price and sell his good. This is especially true when consumers team together with other consumers who value the good at higher or lower levels, but it remains true even when consumer tastes are similar.[17]

This aggregation effect is closely related to what Schmalensee termed the "reduction in buyer diversity"[18] that can result when a producer engages in commodity bundling. Commodity bundling is a practice whereby a seller chooses to sell several goods together in a single package instead of selling each good individually. A long line of scholarship[19] suggests that this can enhance profit since, similar to the above, consumer valuations for multiple products tend to have a probability distribution with a lower variance per good as compared to consumer valuations for each product individually.[20] Our point here is that, just as bundling can increase a seller's revenue by combining a single consumer's demand for several goods, sharing can increase a seller's revenue by combining several consumers' demand for a single good.[21] Under the right conditions, either type of aggregation can be a boon to the seller.[22]

We can illustrate the beneficial effects of demand aggregation by using a simple demand distribution where six consumers value an information good at $5, $7, $9, $11, $13, and $15, respectively. In the absence of price discrimination, a seller facing this demand allocation could extract at most $36 in revenues (by pricing the good at $9). Group these consumers into *any* pairings, however, and the resulting demand curve is easier to exploit. Indeed, the extreme case -- pairing the $5 with the $15, the $7 with the $13, and the $9 with the $11 -- allows the seller to extract the entire surplus by charging a price of $20. Even the pairing that is least favorable leaves the seller with $40 in revenues, an improvement over the case where only single consumer sales were allowed. It is important to note that this aggregation effect depends critically upon our assumption that a team's valuation is approximately equal to the sum of individual team member valuations.[23] This seems natural for the social sharing we study but less appropriate for market-mediated sharing.[24]

Of course, there is one significant limitation to this analogy between sharing and bundling: whereas a seller engaged in commodity bundling can choose both which and how many products to

bundle, a seller who permits sharing has little control over the related concepts of how many and which specific consumers team together to share. A seller could certainly influence team size and composition by (for example) making unauthorized duplication time-consuming or using simple copy protection techniques to deter or stigmatize information piracy, but the precise team patterns are largely out of the producer's control.

This imprecision -- what we call the "team diversity effect" -- is problematic for the seller. After all, demand aggregation is beneficial only because, all other things being equal, teams tend to have more predictable valuations for a good than do individual consumers. Wide variations in team sizes can undermine this predictability, once again making valuations more dispersed and hence more difficult for the seller to surmise. This effect is not evident in earlier scholarship because in those papers market-based models lead to an equilibrium where each information good is shared among exactly the same number of consumers.[25] The contrary view, in contrast, turns out to have significant implications for profitability.

Our formal analysis proceeds as follows. In Part II, we show that, under nearly all traditional demand assumptions, if consumers were to share information goods within equally-sized teams, sharing would always increase producer profit. We argue that sharing in this case is almost perfectly analogous to commodity bundling, and we suggest that, like bundling, sharing can have important effects on profit even when it has no effect on the technology of production. Indeed -- in contrast to the Besen-Kirby result -- we show that sharing can in this case increase profit even if sharing is "inefficient" in the sense that it is more expensive to distribute the good via sharing than it would be for the producer to simply produce additional units itself.

In Part III, we examine sharing in a more realistic form by first allowing team size to vary, and then studying various types of correlations that might exist between team member valuations. We begin by modeling the simplest situation that nevertheless captures the relevant complexity: the case where consumer valuations are uncorrelated, and consumers either purchase information goods in teams of size one (i.e. as individuals) or in teams of size two (i.e. they share them). We then provide a limiting theorem for when sharing will increase profit and use a series of simulations to show that the theorem provides a good approximation under a broad set of conditions. Finally, we consider two classes of correlated valuations: cases where members' valuations are correlated with one another (for example, when high-valuers tend to team with other high-valuers); and cases where valuations are correlated with team size -- as where, for example, consumers with high valuations tend to share only in small teams or not at all.

In Part IV, we analyze what happens when information goods received via sharing are substantially inferior to originals, either because the relevant copying technology is imperfect or because sharing imposes non-trivial transaction or coordination costs. While this will often make sharing less attractive, we show how a producer might actually use such degradation to *increase* profit, by taking advantage of the potential to reshape demand.

Part V concludes with a discussion of the limitations of this work.

II. Sharing in Teams of Constant Size

In this section, we introduce our basic model and show that, for a general set of conditions, so long as consumers share information goods within teams of constant size, sharing will always increase profit.

Consider a setting with a single seller providing an information good to a set of consumers Ω. Assume each consumer demands either 0 or 1 units of the good and that, while resale is not permitted (or is unprofitable for consumers), consumers can share the information good in teams. Specifically, consumers share the good in teams $\varphi_i \in \Phi$, where $\Phi = \{\varphi_1, \varphi_2, \cdots, \varphi_{|\Phi|}\}$ is a partition of Ω. In other words, for all φ_i and φ_j $(i \neq j)$, $\varphi_i \cap \varphi_j = \varnothing$ and $\bigcup_{\varphi_k \in \Phi} \varphi_k = \Omega$. We denote by $|\varphi|$ the size of team φ. For each consumer $\omega \in \Omega$, let $v(\omega, \varphi)$ represent that consumer's valuation when the good is shared within team φ. We will use $v(\omega)$ to denote ω's valuation when the good is not shared with other consumers, i.e., $v(\omega) = v(\omega, \{\omega\})$.

For the analysis presented in part II, we make certain simplifying assumptions which are relaxed, modified, or further considered in the subsequent parts of the paper.

A1: The marginal cost for additional copies of the information good is zero to the seller.

A2: Consumer valuations $v(\omega)$ are independent and uniformly distributed in [0,1].

A3: For all teams $\varphi \subseteq \Omega$, $v(\omega, \varphi) = v(\omega)$; that is, shared copies are perfect substitutes for unshared originals, and, further, they can be made costlessly within teams.

A4: Teams are all of the same size such that, for all $\varphi \in \Phi$, $|\varphi| = \hat{n}$ where $\hat{n} \geq 1$.

The first assumption, A1, is meant primarily to focus attention on the way sharing reshapes the demand curve. We do this by assuming away the seller's marginal costs of production. The assumption resonates since information goods -- like intellectual property more generally -- are

often expensive to create but inexpensive to reproduce. The general trend described here would hold, however, even if this assumption were weakened or removed.[26]

More significantly, assumption A2 asserts that there is no predictable mathematical relationship between team members' valuations for the shared good. That is, we imagine that some teams are composed of members with radically different valuations; others are made up of members with largely similar preferences; and, overall, the pattern is best represented by independent random variables.[27] It is interesting to note that the "correct" assumption here might be context-specific. Teenagers who share computer software, for example, likely have highly correlated tastes. The opposite might apply to the family's subscription to *Time* magazine. These possibilities are explored in part III.

Assumption A2 also requires that consumer valuations be uniformly distributed; that is, we assume linear demand. This specification is made for expositional purposes only. In the appendix, we demonstrate that our results apply to a much broader class of demand functions, including most of those commonly used in economic models.

The third assumption, A3, sets the framework for a later discussion of how an individual's valuation might vary depending on the number of consumers sharing the good. For instance, the value of a software program might be diminished by sharing, as when sharing means that access to the original distribution media and hardcopy documentation are less convenient. Further, sharing might itself impose non-trivial coordination or duplication costs, costs that would ultimately diminish team members' willingness-to-pay.[28] Assumption A3 temporarily forestalls discussion of these effects by stipulating that individual valuations are unaffected by sharing and that shared copies can be created at zero cost. As with assumption A1, this is more likely to be reasonable for information goods than for other types of goods.

Our fourth assumption, A4, provides a starting point for our analysis by focusing on the case where all teams are of uniform size. This is admittedly a strong and unrealistic assumption.[29] Some information goods are likely used only by their direct purchaser while others are surely shared among two, three, or several users. We employ this assumption here, however, because it helps to make clear an important piece of our initial argument: contrary to widespread perception, under certain circumstances sharing unequivocally increases direct profitability. We relax this constant-team-size assumption in part III.

If the four assumptions A1-A4 hold, we show (in an appendix) that sharing always increases profitability. In fact, more sharing -- here, larger teams -- yields even greater profit. More formally,

Proposition 1:
Given assumptions A1, A2, A3, A4, sharing (in teams of size $\hat{n} > 1$) always increases profit compared to the case of no sharing ($\hat{n} = 1$).

Proof: All proofs are in appendix 1.

Corollary 1a:
Under the conditions of Proposition 1, the larger the value of $\hat{n}$, the greater the profit.

As explained in the introduction, the intuition behind Proposition 1 is that, for many common distributions of consumer valuations, as the number of consumers sharing a good increases, the likelihood that the team will value the information good at some "moderate" value (as compared to the overall distribution of team valuations) also increases. That is, as team size grows, it becomes less and less likely that a team will have a significantly high or significantly low valuation relative to the other teams. Since demand curves are derived from the cumulative distribution functions for consumer valuations, this trend helps to make demand more elastic near the mean and less elastic away from the mean. We illustrate the process in Figure 1.


Distribution of valuations for teams of consumers
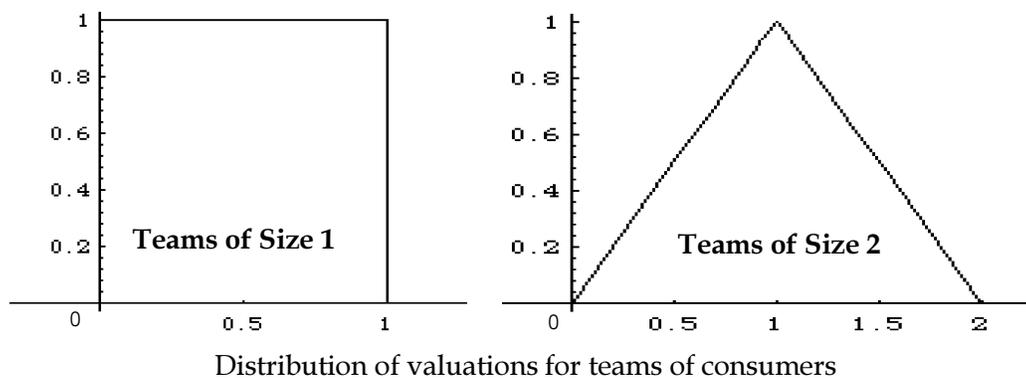
Figure 1: Distribution of valuations…

Note that this result is not restricted to linear demand, but instead applies to a broad range of demand functions. For example, Proposition 1 and Corollary 1a hold for any distribution of consumer valuations where (1) the distribution satisfies the "single-crossing" property developed in Bakos & Brynjolfsson,[30] and (2) the optimal price for the seller given sharing is less than or equal to the mean team valuation. These conditions are satisfied by most common demand functions, such as semi-log and log-log, as well as any demand function based on a Gaussian distribution of valuations. Thus the applicability of Proposition 1 is quite general.

Specifically, let $\mu = E[v(\omega)]$ denote the mean consumer valuation and let $v_n = \frac{1}{n}\sum_{i=1}^{n} v(\omega_i)$

represent the mean valuation per consumer for a team $\varphi = \{\omega_1, \omega_2, \cdots, \omega_n\}$ of size n. Assume that:

A2': Consumer valuations $v(\omega)$ are independent and identically distributed.

A5: Single-Crossing of Cumulative Distributions: The distribution of consumer valuations

is such that $\text{Prob}[|v_n - \mu| < \varepsilon] \leq \text{Prob}[|v_{n+1} - \mu| < \varepsilon]$ for all n and ε.

A6: The seller's optimal monopoly price is $p^* \leq \mu$.

Proposition 2:
Given assumptions A1, A2', A3, A4, A5, and A6, sharing (in teams of size $\hat{n} > 1$) always increases profit as compared to the case of no sharing ($\hat{n} = 1$).

Corollary 2a:
Under the conditions of Proposition 2, the larger the value of $\hat{n}$, the greater the profit.

Assumptions A1 and A3 combine to imply that sharing makes possible no new production efficiencies; that is, we assume that distribution via sharing is no cheaper than distribution via original production. As we discussed in the Introduction, in prior scholarship, sharing is profitable for the seller only where these efficiencies exist.[31] In our model thus far, by contrast, sharing will not only strictly increase profit when distribution and production costs are unchanged, but also can increase profit even when this factor pushes in the opposite direction. This is somewhat surprising. It means that even if sharing is technologically inefficient -- even if sharing is a more expensive means of distributing an information good than is original production -- sharing might nevertheless benefit information sellers.[32]

While Proposition 1 and Proposition 2 are mathematically true no matter how large the team size, the propositions are most relevant to the current context when team size remains relatively small. The intuition here is that demand aggregation itself is a realistic assumption only in small-scale social settings. Friends, we have argued, can be expected to both know and account for one another's preferences. Similarly, a head of household will likely be aware of and internalize family members' valuations. The same cannot be said, however, of larger groups. Large teams will find it difficult to coordinate joint purchases, with any attempt at cooperation likely to fall victim to free-riders, information asymmetry, and other problems of strategic behavior. Thus, whereas our work in

this section applies to small-scale social sharing, the market models of prior literature seem to better capture the interactions where sharing occurs primarily in large groups.

In summary, for small teams, our results suggest that, where team size is constant, sharing can significantly benefit information sellers even in the absence of production efficiencies. This result is fully consistent with related work in the commodity bundling context[33] -- as well it should be. After all, by stipulating constant team size, we have removed from consideration the one critical difference between sharing and commodity bundling: that, due to team size heterogeneity, sharing can be a far less predictable mechanism for combining demand. It is to this additional wrinkle that we now turn.

III. Teams of Varying Size or With Correlated Valuations

In the previous section, we showed that sharing tends to increase producer profit as long as consumers share information goods in teams of constant size, and valuations of team members are independently distributed. Here, we explore the more interesting question of what happens when consumers team together to share information goods in small teams of varying size,[34] have correlated valuations, or both. We begin by analyzing a simple model with only two team sizes -- either one or two members -- and with uniformly distributed, uncorrelated valuations. We then derive a limiting principle for determining the relative profitability of sharing for any arbitrary distribution of team sizes and any arbitrary distribution of consumer valuations. Next, we illustrate this principle with a variety of simulation results. Finally, we consider the implications of two types of correlations: correlations between individual team member valuations, and correlations between team size and team member valuations.

a. Teams of Two Different Sizes

By allowing team size to vary, we can distinguish and quantify the two countervailing effects sharing has on the shape of the demand curve. On one hand, as discussed above, aggregation through sharing can reduce diversity in consumer valuations and thus increase seller profit; we have termed this the aggregation effect. On the other hand, heterogeneity in team size (the "team diversity" effect) will tend to increase diversity and thus reduce profit. In general, either of these two effects may dominate, but, interestingly, when the diversity in team sizes is not too different from the diversity in individual valuations, the two effects roughly cancel out.

To see this, assume that for a team $\varphi \in \Phi$, the number of consumers in this team, n_φ, is distributed with a probability mass function $g(n_\varphi)$. To keep things simple, let us start by exploring the case where n_φ is either 1 (individuals) or 2 (pairs), and consumer demand is originally distributed uniformly (i.e. linear demand). Specifically, suppose:

A5': Consumer valuations are uniformly distributed in [0,1]. Furthermore, a fraction α of consumers ($0 \le \alpha \le 1$) are in teams of size 1 (i.e., do not share), and a fraction $1 - \alpha$ are in teams of size 2.

In other words, for all $\varphi \in \Phi$, $g(n_\varphi) = \begin{cases} \dfrac{2\alpha}{1+\alpha} & \text{for } n_\varphi = 1 \\[2mm] \dfrac{1-\alpha}{1+\alpha} & \text{for } n_\varphi = 2 \\[2mm] 0 & \text{otherwise} \end{cases}$.

Let x_1 and x_2 denote the valuations of teams of size 1 and 2 respectively. Then x_1 is uniformly distributed in [0,1] and x_2 has a probability distribution $h(x_2)$ that is derived from the convolution of two uniform distributions in [0,1], or $h(x_2) = \begin{cases} x_2 & \text{for } 0 \le x_2 \le 1 \\ 2 - x_2 & \text{for } 1 < x_2 \le 2 \\ 0 & \text{otherwise} \end{cases}$.

Thus if we denote by $q_1(p)$ and $q_2(p)$ the sales to single consumers and to teams of size 2, respectively, as a fraction of the total number of consumers $|\Omega|$, when the seller sets price p, we get

$$q_1(p) = \begin{cases} \alpha(1-p) & \text{for } 0 \le p \le 1 \\ 0 & \text{for } p > 1 \end{cases} \quad \text{and} \quad q_2(p) = \begin{cases} \dfrac{1-\alpha}{2}\left(1 - \dfrac{1}{2}p^2\right) & \text{for } 0 \le p \le 1 \\[2mm] \dfrac{1-\alpha}{4}(2-p)^2 & \text{for } 1 < p \le 2 \\[2mm] 0 & \text{for } p > 2 \end{cases},$$

$$\text{resulting in sales } q(p) = \begin{cases} \dfrac{1-\alpha}{2}\left(1-\dfrac{1}{2}p^2\right)+\alpha(1-p) & \text{for } 0 \le p \le 1 \\[3mm] \dfrac{1-\alpha}{4}(2-p)^2 & \text{for } 1 < p \le 2 \\[3mm] 0 & \text{for } p > 2 \end{cases}$$

$$\text{and profit } \pi(p) = p \cdot q(p) = \begin{cases} \dfrac{1-\alpha}{4}\left(2p-p^3\right)+\alpha\left(p-p^2\right) & \text{for } 0 \le p \le 1 \\[3mm] \dfrac{1-\alpha}{4}\left(4p-4p^2+p^3\right) & \text{for } 1 < p \le 2 \end{cases}.$$

To maximize profit, the seller will set $\dfrac{d\pi}{dp}=0$. The resulting first order condition is

$$\frac{d\pi(p)}{dp}=0=\begin{cases} \dfrac{1-\alpha}{4}\left(2-3p^2\right)+\alpha(1-2p) & \text{for } 0 \le p \le 1 \\[3mm] \dfrac{1-\alpha}{4}\left(4-8p+3p^2\right) & \text{for } 1 < p \le 2 \end{cases}.$$ When $1 < p \le 2$, the first order condition

requires $\dfrac{1-\alpha}{4}\left(4-8p+3p^2\right)=0$, which is satisfied when $p=2$. It is easily verified, however, that

this corresponds to a minimum. When $0 \le p \le 1$, the first order condition requires

$\dfrac{1-\alpha}{4}\left(2-3p^2\right)+\alpha(1-2p)=0$, or $3(1-\alpha)p^2+8\alpha p-2(1+\alpha)=0$. This gives

$p^* = \dfrac{-4\alpha+\sqrt{6+10\alpha^2}}{3(1-\alpha)}$ if $0 \le \alpha < 1$ and $p^* = \dfrac{1}{2}$ if $\alpha=1$.

Relative Profit
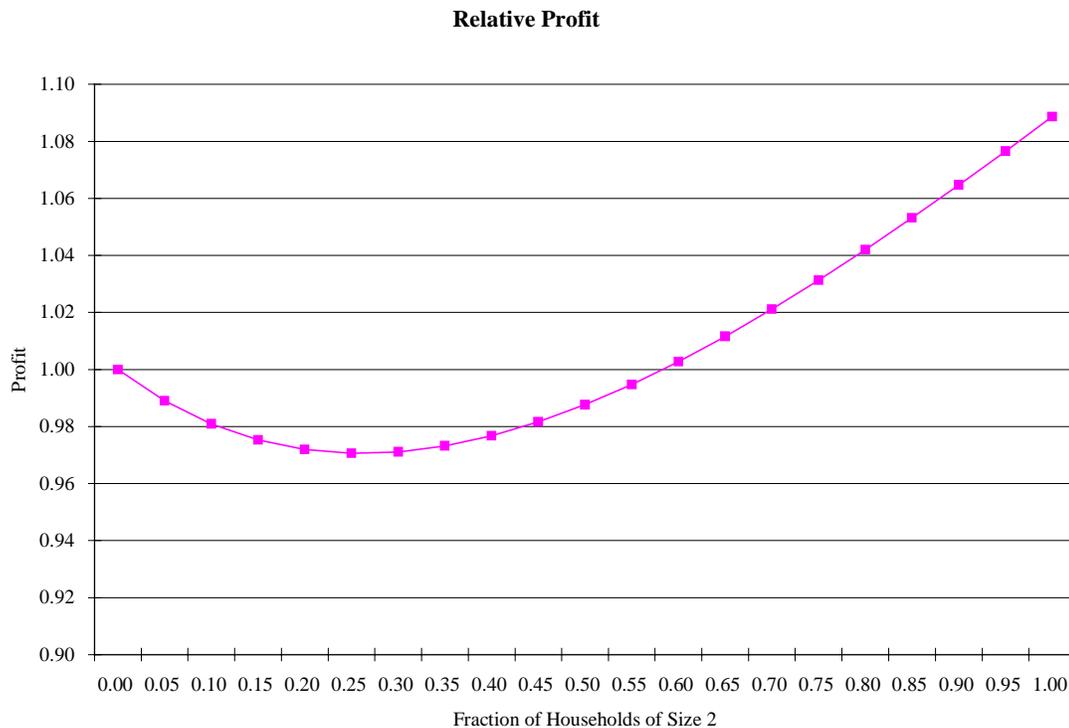
Figure 2: Profit when sharing occurs in teams of size one or size two. The fraction of consumers purchasing in teams of size two increases from left to right, with the left side capturing one extreme (all consumers purchase as individuals), and the right side reflecting the other (all consumers purchase as part of a two-person team). Profit under a no-sharing regime is normalized to 1.

Figure 2, above, shows profit as a function of the fraction of teams that share in households of size 2. We have normalized profit to 1.0 for the case where there is no sharing (i.e. all consumers purchase as individuals, so α equals 1). When everyone shares in teams of size 2, by contrast, profit is approximately 1.089. As long as the fraction of households that share is above a critical value ($1 - \alpha < 2 - \sqrt{2} \approx 0.59$), sharing increases profit. As we expect from the analysis in Part II, profit in this example is highest when all consumers are in teams of size two and hence the team diversity effect is minimized. However, when there is more than one household size, sharing can reduce profit. The most unfavorable distribution of team sizes -- when about one quarter of the teams are of size 2 -- causes a profit drop of about 3%. In this case, the increase in buyer diversity introduced by having multiple team sizes slightly outweighs any favorable effects from demand aggregation.

b. Arbitrary Distributions of Team Sizes and Valuations

As is surely already clear, the interplay between the aggregation effect and the team diversity effect can either enhance or diminish profit. We showed in Part II, for example, that when team size is constant, the aggregation effect dominates and profit always increases. When team size varies widely, by contrast, the team diversity effect may instead prevail. In this part, we explore more general cases in the middle of the range -- cases where both effects have relevance, and hence profitability can move in either direction.

Each such case will be defined by two parameters. The first parameter will be some specific assumption as to the pattern of team sizes. Consumers might (as above) always share the good in teams of size one or two; or they might tend to share it in teams such that the distribution is (say) normally distributed around a mean size of three. The second parameter will be another distribution, this one representing consumer valuations in the original demand curve. The original demand distribution is relevant since, if consumer preferences are relatively predictable even without sharing, the marginal benefits of aggregation can at most be small. Conversely, if individual consumer preferences initially vary widely, aggregation can significantly improve profit even if team size diversity is also substantial.

With these factors in mind, we begin by developing a limiting theorem that establishes a rough relationship between profit, the diversity inherent in the original demand curve, and team size heterogeneity. Specifically, we show that if member valuations are independent and team sizes are large but varied, then profit is almost entirely determined by the distribution of team sizes. We then show in the simulations that follow that, even though the theorem is derived under a large team assumption, it nevertheless offers a useful approximation for the cases of interest here, cases where team size is varied but relatively small.

This limiting theorem is important for two reasons. First, it establishes a general rule of thumb as to when sharing will be profit-enhancing, profit-diminishing, or limited in its effect. Specifically, the theorem suggests that, when teams are of large size, profit under any team pattern will closely approximate the profit a seller would have earned were individual consumer valuations distributed according to that same pattern. Thus, if team sizes are uniformly distributed and consumer demand was originally distributed according to (say) a certain Gaussian distribution, an information seller with the option of producing a sharable good is choosing between two demand alternatives: allow sharing and thereby exchange the Gaussian demand distribution for the uniform one; or prevent sharing and sell into the Gaussian distribution. The relative profitability of various

demand distributions is easy to calculate; our theorem therefore provides a workable test to determine whether sharing will increase or decrease profit in a given market.

Second, in certain circumstances, the distribution of team sizes might be better known than the distribution of consumer valuations. In such cases, a seller may be able to take advantage of the limiting theorem to in effect transform demand from a relatively unknown pattern (original demand) toward a more familiar distribution.

Consider a probability distribution function f that allows only nonnegative valuations, has a finite mean μ_f, and finite variance σ_f^2. Let π_f^* be the optimal profit (as a fraction of the total area under the demand curve) for a seller facing individual consumer valuations distributed according to f. Define g to be a probability mass function representing a distribution of team sizes with mean μ_g

such that: $g(n) = \begin{cases} \int_{x=(n-1)\mu_f/\mu_g}^{n\mu_f/\mu_g} f(x)dx & \text{if } n > 0 \\ 0 & \text{otherwise} \end{cases}$ [35].

In this setting, under assumptions A1 and A3, if the valuations of individual consumers are independently distributed in any arbitrary distribution with a finite mean and variance, and if consumers share the information good in teams, and the distribution of team sizes has probability mass function g, as μ_g increases, the optimal profit π_g^* received by the seller (again expressed as a fraction of the total area under the demand curve) converges on expectation to π_f^*:

Proposition 3:
$$\lim_{\mu_g \to \infty} \mathrm{E}[\pi_g^*] = \pi_f^*.$$

Thus, when consumers share the good in large teams, the primary determinant of seller profit is the distribution of team sizes rather than the distribution of consumer valuations. The following corollary is a natural extension:

Corollary 3a:
If consumer valuations and team sizes have the same distribution, then sharing and not sharing are equally profitable on expectation in the limit as team size increases.

Clearly, it also follows from Proposition 3 that, if the distribution of team sizes is more favorable (e.g. more concentrated) than the distribution of individual valuations, sharing will be beneficial in the limit. Similarly, when the distribution of team sizes is less favorable, sharing will

tend to reduce profit. The analysis of constant team sizes in part II can now be seen to be simply a special case of this more general result. A constant-value distribution, after all, will always be more favorable than any other distribution of values.

How large can these profit-enhancing and profit-reducing effects be?   There is actually no theoretical limit to the amount by which sharing can increase or decrease profit. For instance, there are some distributions of individual consumer valuations for which the maximum profit extracted by the seller is an arbitrarily small fraction of the total area under the demand curve.[36] In these cases, if sharing were allowed, its beneficial effects could be arbitrarily large. Indeed, were sharing to occur in large teams of constant size, the seller would capture nearly the full value of his good instead of capturing just that negligible fraction. At the other extreme, if initial consumer valuations are distributed in a very favorable manner, but team sizes are distributed in a strongly unfavorable manner, then sharing can dissipate essentially all of the seller's profit.

In practice, of course, neither of these extreme cases is likely to result. To get a sense of the likely significance of the reshaping of demand due to sharing, we have simulated the outcomes for a variety of different situations, using a selection of the most commonly used demand curves. The simulations are also useful because Proposition 3 and its corollary were derived for large teams and are not very precise for smaller teams. By simulating a broad array of team size patterns and initial demand distributions, we have found that the theorem is nevertheless a valuable guide even in the small team setting.

A sampling of our simulation results is included in Table 1.

Relative Profit Under Sharing Regime

Sample Distributions of Individual Consumer Valuations / Sample Team Size Distributions	Identical Valuations of 0.5	Uniform Distribution on [0,1]	Gamma Distribution (3,1/6)	Power Distribution (-1)	Exponential Distribution (2)	Normal Distribution (1/2,1/12)
All Teams Size 3	**1.0000**	1.1477	1.2507	1.4625	1.2416	1.1347
Uniform Distribution on [0,6]	0.5714	**0.9597**	1.0449	1.5775	1.1905	0.9340
Gamma Distribution (3,1)	0.5409	0.8817	**0.9534**	1.4704	1.1017	0.8593
Power Distribution (-1)	0.3612	0.5331	0.5247	**1.5209**	0.8349	0.4920
Exponential Distribution (1/3)	0.4367	0.7474	0.8120	1.5495	**0.9999**	0.7267
Normal Distribution (3,3)	0.6100	0.9556	1.0347	1.5156	1.1654	**0.9316**

Table 1: The ratio of profit with sharing to profit without sharing, for various alternative assumptions about the original distribution of consumer valuations (the horizontal axis) and the distribution of team sizes (the vertical axis). Values greater than 1 indicate situations for which profit increased due to sharing. See appendix II for details on the simulation methods.

There are several notable features in the simulation results. The first row represents the relative profit when all teams have exactly 3 members. Reading across, we see that sharing has no effect on profit when individual valuations also are constant, but increases profit substantially for all other distributions in this row. For instance, if individual valuations are distributed uniformly between 0 and 1, then sharing in teams of size 3 will increase profit by over 14%.[37] This row can be interpreted as reflecting the pure "aggregation effect" of sharing; in these cases, after all, there is no team diversity. Similarly, the first column can be interpreted as capturing the pure "team diversity effect" -- in each case, consumers with identical valuations are assumed to team together in various patterns and, hence, buyer diversity can only increase. As one might expect, going from a demand curve based on constant individual valuations to demand curves in which valuations are dispersed because of non-constant team sizes always works to lower profit.

Each of the other cells represents a different combination of the aggregation and team diversity effects. Net profit sometimes increases and sometimes decreases. For instance, when

individual valuations are distributed according to a power distribution (also commonly known as a constant elasticity demand function), it can be particularly difficult to extract profit. Hence, allowing such goods to be shared and thereby reshaping the demand curve will usually increase profit, often by 50% or more. Conversely, if *team* sizes follow a power distribution, sharing will almost always reduce profit by reshaping the demand curve in an unfavorable direction.

The diagonal (in bold text) represents cases in which both the individual valuations and the distribution of team sizes follow the same distribution. As predicted by Corollary 3a, the profit ratio tends to be close to 1 in most cases. The ratios do differ from 1 somewhat because the Proposition relies on the law of large numbers and here we use teams with a mean size of only 3.[38]

The profit implications of small-scale sharing will of course vary depending on the actual distributions of individual valuations and team sizes relevant to the good in question. However, it is useful to consider the effect on profit assuming that team sizes follow the actual distribution of household sizes in the United States, and using some empirical estimate for the distribution of consumer valuations for information goods. We obtained the distribution of household sizes from the 1997 Statistical Abstract of the United States.[39] We are not aware of any estimates for the demand elasticity of information goods, but Brynjolfsson[40] did find that the demand for all types of information technology could be accurately approximated by a log-linear function with an exponent of approximately -1. Using these two data sources, we find that sharing increases profit by 13.5% compared to a baseline case with no sharing.

c. Correlated Team Member Valuations

In the Introduction, we noted that, while in parts II and III we would assume team member valuations to be independent, it was possible that, in practice, valuations were somehow correlated. For instance, it is possible that friends tend to have similar taste in music, so music is therefore shared among individuals with highly correlated valuations. Alternatively, households may consist of adults and children with very different willingness-to-pay for word processing software.

The topic of correlations has been extensively analyzed in the bundling literature, and many of those results carry over to the social sharing setting.[41] For instance, a negative correlation in team member valuations will tend to decrease buyer diversity and increase profit, much as a negative correlation between a consumer's valuation for two bundled goods typically makes bundling more attractive for the seller.[42] This implies, for example, that although information providers may find it unprofitable to price low enough to sell significant quantities individually to children or teenagers,

they may nevertheless find it profitable to allow parents to share information goods with the rest of the household. Sharing might also increase profit in the presence of a positive correlation between team member valuations, as long as the correlation is not perfect; however, akin to positive correlations in the commodity bundling context,[43] the beneficial effect will be markedly smaller when the correlation is high.

It is also possible that the propensity to share will be correlated with consumer willingness-to-pay. For instance, financially-strapped teenagers may be more likely to tape one another's CDs than are successful businessmen. It is easy to see that diversity in team valuations will be reduced if, *ceteris paribus*, team size is negatively correlated with the average valuation of team members. In this case, sharing will be more profitable than it would be were team size unrelated to individual valuations. By charging a single price for the original copy, the seller effectively charges a lower per-copy price to members of larger teams. Such differential pricing will increase profit if it reflects the distribution of individual valuations.

While it might seem to be a rare and fortunate coincidence for the seller to face a negative correlation between valuations and team size, there is reason to believe that this pattern may frequently arise as a natural equilibrium. Suppose that sharing imposes some inconvenience or congestion cost on team members and that consumers with high valuations are less willing to incur the inconvenience cost of sharing. This seems quite plausible; if nothing else, foregoing access to the good for any given amount of time will be more costly to an individual who values it highly. In this case, individuals will naturally sort themselves into team sizes based on valuations. High value consumers will choose to pay a higher effective price by being in smaller teams, including "teams" of size one.[44] Low-valuers will tolerate the inconveniences of larger teams. In such cases, aggregating demand into teams will be even more effective at reducing buyer diversity.

IV. Degradation of Shared Goods

We have assumed thus far that consumers would value a shared information good just as much as they would value an unshared original. This is why we were able to claim that teams of consumers would be willing to pay up to the sum of what each team member was willing to pay individually. Of course, this is a vulnerable assumption. Copying technologies are at times imperfect; and, even if that were not a problem, sharing itself often imposes non-zero coordination and transaction costs. These effects can make shared information goods less valuable than originals.

The extent of the decrease probably varies with the type of good. Copies of educational software might be close substitutes for legitimate versions; the copying is inexpensive, and any missing packaging or manuals are often either of little value, or are easily and inexpensively replaced. A passed-along magazine similarly retains much of its value, unless difficulties in coordination mean a significant delay in the sharing. On the other hand, a tape recording of the latest compact disc is often of diminished value, both because of its sound quality and the costs of duplication.

To whatever extent these factors cause shared information goods to be of lesser value, our model does overstate total demand. This overstatement is not so severe as might at first be expected; after all, in every team, at least one consumer enjoys access to the original version. Other consumers, however, would compensate for transaction costs, coordination costs, and any diminution in quality by decreasing their willingness-to-pay. That decrease would, in turn, translate into lower direct profit.

Offsetting this factor, however, is a corresponding decrease in the costs incurred by the producer of the original information good. This is the point first made by Besen and Kirby.[45] In a world without sharing, producers must create and distribute one copy of the good to each ultimate consumer. With sharing, by contrast, the producer need only deliver one copy to each team. The teams then duplicate and/or distribute the good to members.[46] Every unit that is potentially undervalued by the factors discussed above, then, is also a unit that the producer brought to market at zero cost. In some cases, this balance likely tips in favor of the producer; in others, it represents a decrease in profit.[47]

In addition to the unavoidable product degradation discussed above, we find that a producer might benefit from *intentionally* designing his product so that copies are less valuable than originals. Imagine, for example, that the maker of a popular word processing program believes it can maximize profit by discouraging adults from sharing with other adults, but encouraging parents to share with their children, as discussed in part III.c above. By implementing the appropriate degradation strategy -- say, refusing to sell customer support to users other than registered, legitimate purchasers -- such a seller might induce the desired behavior. Youngsters (who tend not to value customer support) would continue to share with their parents unabated, but teams of adults would be more likely to purchase multiple legitimate copies.[48]

We illustrate how intentional product degradation can beneficially reshape demand in the following simplified setting:[49]

1. Consumers have two possible valuations, which we refer to as "high" and "low" and denote respectively by v_H and v_L ($v_H > v_L$). A fraction α of consumers value the information good at v_H, leaving a fraction $1-\alpha$ that value it at v_L.

2. Consumers form teams of size 2 (e.g., they pair with an associate). These pairs are exogenously formed, as they would be if based on pre-existing social relationships. We assume that the valuations v_1 and v_2 of the two members in each team are not correlated.

3. The seller sets a price, p, and a degradation level, d, for shared goods, e.g., by deciding how much service to deny non-original copies. If a good is shared, the user of the original enjoys its full utility, while a maximum utility level d can be enjoyed by the user of the copy.[50]

We continue to maintain the assumption that originals can be produced at zero cost, which rules out the possibility of any production-cost savings from sharing.

Given the price, p, and the degradation level, d, each consumer team will fall into one of three categories: (a) those that purchase two originals and thus do not share the information good; (b) those that do not purchase the good; and (c) those that purchase a single original and then share the information good.

A team will fall into category (a) when the member of the team with the lower utility prefers to purchase an original at price p rather than get utility d from a shared copy. This will be the case if $\min(v_1, v_2) - d \geq p$, which implies $v_1 \geq p + d$ and $v_2 \geq p + d$. A team belongs in category (b) when its total valuation is less than the purchase price: $\max(v_1, v_2) + \min(v_1, v_2, d) < p$. Finally, a team will be in category (c) when both $\min(v_1, v_2) - d < p$ and $\max(v_1, v_2) + \min(v_1, v_2, d) \geq p$.

Figure 3 illustrates these decisions in two extreme cases: (1) when there is no degradation of the shared copy ($d \geq v_H$), which is equivalent to the seller allowing full sharing; and (2) when there is complete degradation of the shared copy ($d = 0$), which corresponds to the seller not allowing sharing at all.

Teams Purchase:

Two originals (not shared) One original (shared if allowed) No purchase

Full sharing allowed
(no degradation of copy)

No sharing allowed
(complete degradation of copy)
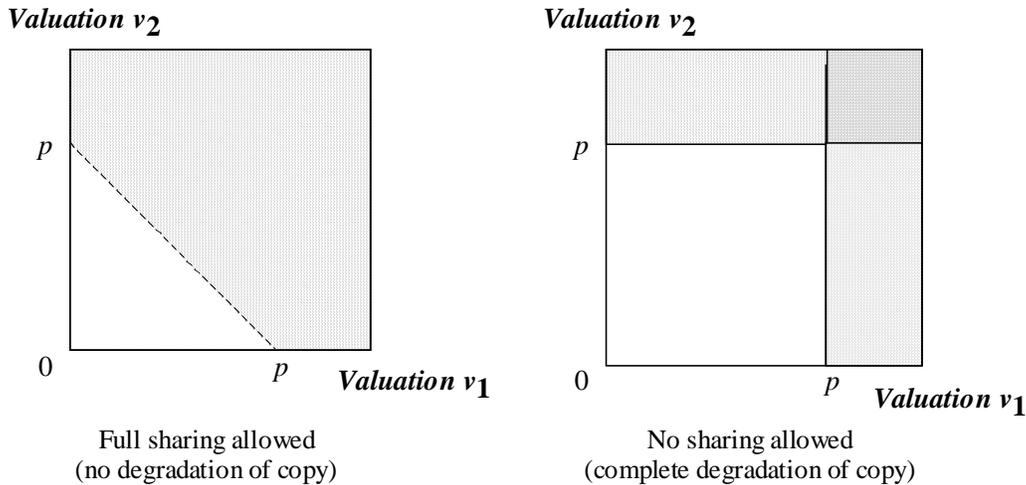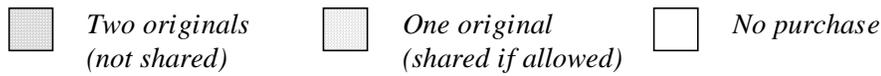
Figure 3: Team decisions when seller sets a price p and allows either full sharing or no sharing. The optimal p will depend on the distribution of valuations.

In the above setting, Proposition 4 demonstrates that, depending on the initial distribution of valuations, it may be possible for the seller to obtain higher profit by strategically degrading copies of the information good, i.e., by setting $d > 0$ but less than v_H.

Proposition 4:
For some distributions of valuations, the seller of an information good can increase profit by strategically degrading copies to a value less than that of originals, but more than zero.

While Proposition 4 shows that sharing can become more profitable if copies are made less valuable than originals, it should be noted that it is not always profitable to degrade copies. The benefits derive from the possibility of reshaping a demand curve that is initially unfavorable to extracting profit (i.e. one in which some teams have valuations that are significantly higher than others), into one more conducive to profit-extraction.

V. Limitations

The aggregation and team diversity effects modeled in this paper will reshape demand and thus influence profitability whenever three conditions hold: (1) by virtue of some form of

intellectual property protection, regulatory structure, or first-mover advantage, a seller has market power and thus can set his price instead of being forced by competitive pressures to charge a price equal to marginal cost;[51] (2) consumption of the shared good is largely non-rival, such that purchasers and copier/borrowers can each enjoy the good at approximately its full value despite the sharing; and (3) social norms, a legal rule, or a product packaging decision cause the relevant sharing to occur among friends and family members, not as part of an efficient market for second-hand goods.

These conditions hold in many settings, but seem especially likely in cases where consumers casually share information goods. Consider the market for computer software. Copyright protection gives software producers some degree of market power by limiting competitors in their ability to sell an identical good. Consumers can pirate computer software without significant expense and hence can share in a largely non-rival fashion. And legal rules, while unable to stop piracy altogether, do serve to encourage consumers to pirate *quietly* -- sharing with friends and family members but not risking detection by sharing with acquaintances or strangers.

Cable television offers another example of a good to which the model likely applies. Cable providers enjoy significant market power, both because cable programming is protected by copyright and because, in most communities, cable itself is provided monopolistically. Cable can be shared among family members, at least to the extent they watch at different times or have similar viewing preferences. And, again, this information good will primarily be shared only within close social circles since, to share, viewers must have access to common television equipment.

Where any of the three conditions are violated, of course, the model will not apply. For example, the argument does not apply very well to homogenous goods sold in competitive markets. Sellers of these goods have little ability to set their prices above the level of marginal cost. Likewise, our model does not explain the types of sharing facilitated by used bookstores and video rental establishments -- two mechanisms that help strangers share information goods in a market setting.[52] Finally, our model is unlikely to be compelling in cases where the shared good is a tangible resource (e.g., a car or computer), since one consumer's use of such a good often precludes another's, and this competition causes both consumers to lower their valuations.

Furthermore, even where our model does apply and, hence, demand aggregation and team diversity both work to reshape demand, the overall profit implications will not depend on these effects alone. Instead, other factors -- including the changes in production costs first considered by Besen and Kirby -- also have relevance, and it is the combination of these many factors that will

determine profitability in any given case. For example, the sharing of a magazine is likely to implicate both the Besen-Kirby model and our own: sharing in this case is surely less expensive than original production; but sharing also likely significantly reshapes demand. The relative importance of each factor will vary with the type of good. The effects we discuss are likely to have their greatest empirical relevance when the seller's marginal costs of production and consumers' marginal costs of sharing are either both low, or simply roughly comparable. The effects emphasized by Besen and Kirby, by contrast, will have particular application in cases where the difference between those costs are significant.

Conclusion

Until now, scholars considering small-scale social sharing have argued that profitability turns largely on the relationship between consumers' marginal cost of sharing and sellers' marginal cost of producing original units. This is unquestionably an important factor. However, our work here suggests that the role sharing plays in reshaping demand will often be equally or more important. Sharing, we find, will typically increase profitability where demand aggregation serves to decrease buyer diversity, but decrease profitability where heterogeneity in team sizes is so severe that it outweighs the aggregation effect. As a result, public and private policies that simply focus on the total amount of sharing done, without regard to the types of teams that do the sharing, will often severely miscalculate the effect on profit. Depending on the nature of demand reshaping induced, a given amount of sharing will, in some cases, diminish seller profit, while in other cases that same amount of sharing will actually increase it.

Producers, of course, need not idly sit by while sharing transforms demand. As highlighted in Part IV, strategic degradation can be used to influence team size, formation, and behavior, thereby encouraging consumers to share in ways that increase profitability. Moreover, natural sharing patterns will themselves sometimes enhance profitability, as where high-valuing consumers tend to be less willing to share than their low-valuing counterparts. One purpose of our analysis has been to uncover these demand-shaping tools and situations.

Our main goal, however, has been to focus attention on a new set of issues related to small-scale sharing. Two decades ago, the realization that commodity bundling could not only affect production costs, but also reshape consumer demand, provided a valuable new perspective for both business strategy and social policy. By highlighting how sharing can reshape consumer demand in

information markets, we hope to introduce similar issues into the legal, economic, and strategic

analyses of information policy.

Appendix 1: Proofs of Propositions

Proposition 1

For a team $\{\omega_1, \omega_2, \cdots, \omega_n\}$ of size n, let $w_n = \sum_{k=1}^{n} v(\omega_k)$ represent the total willingness-to-pay for the team, and denote by f_n and F_n the probability density function and cumulative distribution for w_n. By assumption A2, w_1 is uniformly distributed in $[0,1]$. When selling to teams of constant size n, the seller faces demand corresponding to a distribution of valuations given by f_n; at a price np_n, i.e., p_n <u>per consumer</u>, the seller will sell to a fraction $q_n(np_n) = 1 - F_n(np_n)$ of teams (and consumers). By assumption A2, the resulting profits will be $\pi_n = p_n q_n(np_n)$. The seller will set his price at p_n^* per consumer, which results in the optimal profit π_n^*, and corresponding sales $q_n^* = q_n(np_n^*)$.

When selling to teams of size $n+1$ at price p per consumer, the seller will realize sales of $q_{n+1}((n+1)p) = 1 - F_{n+1}((n+1)p)$. Computing $F_{n+1}((n+1)p)$ we get:

$$
\begin{aligned}
F_{n+1}((n+1)p) &= \int_{x=-\infty}^{\infty} \left[\int_{y=-\infty}^{(n+1)p-x} f_1(x)f_n(y)dy \right] dx = \int_{x=0}^{1} \left[\int_{y=0}^{(n+1)p-x} f_n(y)dy \right] dx \\
&= \int_{x=0}^{1} \left[\int_{y=0}^{np} f_n(y)dy \right] dx + \int_{x=0}^{1} \left[\int_{y=np}^{(n+1)p-x} f_n(y)dy \right] dx \\
&= F_n(np) + \int_{x=0}^{1} \left[\int_{y=np}^{(n+1)p} f_n(y)dy + \int_{y=(n+1)p}^{(n+1)p-x} f_n(y)dy \right] dx \\
&= F_n(np) + \int_{y=np}^{(n+1)p} f_n(y)dy - \int_{x=0}^{1} \left[\int_{y=(n+1)p-x}^{(n+1)p} f_n(y)dy \right] dx \qquad (1)
\end{aligned}
$$

It follows from (1) that if $np < \mu_n$, i.e., if $p < 1/2$, then $F_{n+1}((n+1)p) < F_n(np)$, and thus $q_{n+1}((n+1)p) > q_n(np)$. Similarly, if $p > 1/2$, then $q_{n+1}((n+1)p) < q_n(np)$.

Thus, if $p_n^* < 1/2$, $q_{n+1}((n+1)p_n^*) > q_n(np_n^*)$, which implies that $\pi_{n+1}^* > \pi_n^*$. Furthermore, $p_{n+1}^* < 1/2$ as well, because if $p_{n+1}^* \geq 1/2$, then $q_n(np_{n+1}^*) \geq q_{n+1}((n+1)p_{n+1}^*)$, and thus, $\pi_n(np_{n+1}^*) = p_{n+1}^* q_n(np_{n+1}^*) \geq p_{n+1}^* q_{n+1}((n+1)p_{n+1}^*) = \pi_{n+1}^* > \pi_n^*$, contradicting the optimality of p_n^*.

It can be easily verified that $p_2^* < 1/2$ and $\pi_2^* > \pi_1^* = 1/4$. The Proposition thus holds for $\hat{n} = 2$. If the Proposition holds for some n, and if $p_n^* < 1/2$, then we showed above that the Proposition holds for $n+1$, as $\pi_{n+1}^* > \pi_n^*$, and moreover $p_{n+1}^* < 1/2$. Thus by induction, the Proposition is true for all values of $\hat{n}$.

Corollary 1a follows directly from the proof of Proposition 1, as $\pi_{n+1}^* > \pi_n^*$ for all n.

Proposition 2

Using the same notation as in the proof of Proposition 1, Assumption A5 implies that when the seller sets price p per consumer, if $p < \mu$, then $F_{n+1}((n+1)p) < F_n(np)$, and thus $q_{n+1}((n+1)p) > q_n(np)$. Similarly, if $p > \mu$, then $q_{n+1}((n+1)p) < q_n(np)$.

Thus, if $p_n^* < \mu$, $q_{n+1}((n+1)p_n^*) > q_n(np_n^*)$, which implies that $\pi_{n+1}^* > \pi_n^*$. Furthermore, $p_{n+1}^* < \mu$ as well, because if $p_{n+1}^* \geq \mu$, then $q_n(np_{n+1}^*) \geq q_{n+1}((n+1)p_{n+1}^*)$, and thus,

$\pi_n(np_{n+1}^*) = p_{n+1}^* q_n(np_{n+1}^*) \geq p_{n+1}^* q_{n+1}((n+1)p_{n+1}^*) = \pi_{n+1}^* > \pi_n^*$, contradicting the optimality of p_n^*.

Since according to assumption A6, $p_1^* < \mu$, the above reasoning can be applied inductively, proving the proposition for teams of any size $\hat{n} > 1$.

Corollary 2a follows directly from the proof of Proposition 2, as $\pi_{n+1}^* > \pi_n^*$ for all n.

Proposition 3

Let m denote the total number of consumers, and let h be the distribution of consumer valuations, with finite mean μ_h and finite variance σ_h^2. Then, m/μ_g is the expected total number of teams, $mg(n)/\mu_g$ is the expected number of teams of size n, and $m\mu_h$ is the total area under the demand curve.

The expected total valuation of a team of size n is $n\mu_h$, and thus the total willingness-to-pay of consumers in teams of size n as a fraction of the total area under the demand curve is $ng(n)/m$. Thus, on expectation, the demand curve faced by a seller selling to teams with team size distributed according to g is identical to the demand curve faced by a seller to individual consumers whose valuations are distributed according to g, with the demand curve for teams scaled by μ_g, so that the mean team valuation is $\mu_g\mu_h$. It follows that the profit maximizing quantity $\hat{q}_g^*$ (as a fraction of the total number of teams and consumers respectively) and profit $\hat{\pi}_g^*$ (as a fraction of the total area under the demand curve) when selling to consumers will be identical to the corresponding quantity and profit when selling to teams, or $\hat{q}_g^* = q_g^*$ and $\hat{\pi}_g^* = \pi_g^*$.

In the limit as $\mu_g \to \infty$, g approaches f, and thus $\lim_{\mu_g \to \infty} E[\hat{\pi}_g^*] = \pi_f^*$, which implies $\lim_{\mu_g \to \infty} E[\pi_g^*] = \pi_f^*$, which proves convergence on expectation.

A stronger version of Proposition 3 holds as well, with convergence in probability: As $\mu_g \to \infty$, $\pi_g^* \xrightarrow{P} \pi_f^*$. To show convergence in probability, we must show that for any $\varepsilon > 0$, there is a large enough $\hat{\mu}_g$ such that $\text{Prob}\left[\left|\pi_g^* - \pi_f^*\right| > \varepsilon\right] < \delta$ for $\mu_g > \hat{\mu}_g$. In the limit as $\mu_g \to \infty$, and for any fixed team size $\hat{n}$, the total valuation of groups of size less than or equal to $\hat{n}$ (as a fraction of the total area under the demand curve) converges in probability to zero. Thus we can assume that optimal price, quantity and profit is determined by teams of arbitrarily large size. This allows us to apply Proposition 1 in Bakos and Brynjolfsson[53], which implies that for teams of large enough size, the total team valuation (as a fraction of the total area under the demand curve) will converge in probability to its expected value. Thus, as $\mu_g \to \infty$, $\pi_g^* \xrightarrow{P} \hat{\pi}_g^*$, and $\pi_g^* \xrightarrow{P} \pi_f^*$.

Proposition 4

Proof by example: In the setting discussed in the text, if $\alpha \dfrac{2+\alpha}{2-2\alpha-\alpha^2} < \dfrac{v_L}{v_H} < \dfrac{1}{3}$, then intentionally degrading the value of shared copies by setting $d = v_L$ will increase seller profit. Specifically, a fraction α^2 of teams will value the information good at $2v_H$, a fraction $(1-\alpha)^2$ will value it at $2v_L$, and the remaining fraction 2α of teams will have total valuation of $v_L + v_H$.

If no sharing is allowed, the seller can set his price either to v_H, realizing profits αv_H, or to v_L, selling to all consumers and realizing profits v_L. $\dfrac{v_L}{v_H} > \alpha \dfrac{2+\alpha}{2-2\alpha-\alpha^2}$ implies that $\dfrac{v_L}{v_H} > \alpha$, i.e., $v_L > \alpha v_H$, and thus the seller maximizes profits by setting his price to v_L. If sharing is freely allowed, the seller can set his price either to $2v_H$, $v_H + v_L$, or $2v_L$. The corresponding profits are $\alpha^2 v_H$, $(\alpha + \dfrac{1}{2}\alpha^2)(v_H + v_L)$, and v_L respectively. The assumption $\dfrac{v_L}{v_H} > \alpha \dfrac{2+\alpha}{2-2\alpha-\alpha^2}$ implies that $(\alpha + \dfrac{1}{2}\alpha^2)(v_H + v_L) < v_L$, i.e., the seller maximizes profits by pricing at $2v_L$. In the absence of degradation, the seller is indifferent between allowing and permitting sharing.

By degrading the value of copies to a maximum of v_L while keeping the price of originals at $2v_L$, the seller will not affect the purchasing behavior of teams with valuations of $2v_L$ and $v_L + v_H$. Teams with valuations of $2v_H$, however, can purchase two originals and realize a surplus of $2v_H - 4v_L$, or share a single original for a surplus of $v_H - v_L$. Since $\dfrac{v_L}{v_H} < \dfrac{1}{3}$, the former choice is

preferred, and thus degradation allows the seller to realize additional sales of $\frac{1}{2}\alpha^2$ units at price

$2v_L$, leading to additional profit of $\alpha^2 v_L$.

Appendix 2: Simulation methodology.

The simulations are implemented by Matlab (Version 5.1.0.421) on a Unix machine running Solaris 2.5.1. To calculate a value for the table, we first determine the frequency of teams of size k, which is given by g(k), where g is the density function of the team size distribution. The simulation is run one million times for each density function, hence the total number of teams of size k is 10^6*g(k). We restrict k to be an integer. For distributions with no upper bound in the domain, we generate teams of size k up to the point where 10^6*g(k) is less than 1 (i.e. less than one team in a million is this large). The exception is the power(-1) distribution. In this case, we truncate the domain from 1 to 20 in order to establish a finite mean which is equal to $\dfrac{\ln 20}{1 - 20^{-1}} \cong 3.15$ [54]

Given a team of size k, we can calculate the total value of the team by independently drawing k values according to the appropriate valuation distribution F(v), and then taking the sum. In the table, each column denotes a different value distribution for individuals.

For each distribution, we calculate profit by dividing the relevant quantities into a grid with a step size of one ten-thousandth of a percentile (i.e. 1.0×10^{-6}). For each quantity, we determine the corresponding price, based on the estimated demand function, and then calculate the corresponding profit. Finally, we choose the price and quantity that maximize profit for the function.

To obtain the ratios in the table, we first adjust profit by team size: an item sold to a single team of size k could have been sold to k individuals instead. Thus we divide the total profit for that team by k to give the profit per consumer from team sales. Next, we divide the adjusted profit by the profit that could be obtained by selling to individuals with the same valuation distribution but without any sharing. Therefore, a ratio larger than 1 can be directly interpreted as indicating that allowing sharing, given the specific distributions, is more profitable than not allowing sharing.

The distribution and density functions that are used are as follows:

Distributions for Individual Valuations:

Constant demand:

$q = 1_{\{p \leq 1/2\}}$, where $1_{\{p \leq 1/2\}}$ is the characteristic function of $\{p \leq 1/2\}$

Uniform distribution (linear demand):

$q = (1 - p)$

Gamma distribution:

$q = \int_p^\infty \frac{\lambda^\alpha}{\Gamma(\alpha)} e^{-\lambda x} x^{\alpha-1} dx$, where $\alpha = 3, \lambda = \frac{1}{6}$

Power distribution with an exponent of -1 (log-linear demand):

$\ln q = a - \ln p$

Exponential distribution with parameter 2 (semi-log demand):

$\ln q = a - 2p$

Normal distribution with mean 1/2 and variance 1/12 (Gaussian demand):

$q = \int_p^\infty \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) dx$, where $\mu = \frac{1}{2}, \sigma^2 = \frac{1}{12}$

Distributions for Team Sizes

Constant team size:

$G_C(k) = \begin{cases} 0, & k < 3 \\ 1, & k \geq 3 \end{cases}$, $g_C(k) = $ Dirac density function with point mass at 3

Uniformly distributed team size:

$$G_U(k) = \frac{k}{20}, \quad g_U(k) = \frac{1}{20}$$

Gamma distributed team size:

$$G_\Gamma(k) = \int_0^k \frac{\lambda^\alpha}{\Gamma(\alpha)} e^{-\lambda x} x^{\alpha-1} dx, \quad g_\Gamma(k) = \frac{\lambda^\alpha}{\Gamma(\alpha)} e^{-\lambda k} k^{\alpha-1}, \text{ where } \alpha = 3, \lambda = 1$$

Power distribution of team size with an exponent of -1:

$$G_P(k) = \frac{1 - k^{-1}}{1 - b^{-1}}, \quad g_P(k) = \frac{k^{-2}}{1 - b^{-1}}, \ b = 20$$

Exponential distribution team size with parameter 3:

$$G_E(k) = 1 - \exp\left(-\frac{k}{\mu}\right), \quad g_E(k) = \frac{1}{\mu} \exp\left(-\frac{k}{\mu}\right), \ \mu = 3$$

Normally distributed team size with mean 3 and variance 3:

$$G_N(k) = \int_{-\infty}^k \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) dx, \quad g_N(k) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(k-\mu)^2}{2\sigma^2}\right), \ \mu = 3, \sigma^2 = 3$$

Summary of parameters used for the simulation

Valuation distribution	μ	σ^2
Constant	$\dfrac{1}{2}$	0
Uniform(0,1)	$\dfrac{1}{2}$	$\dfrac{1}{12}$
$\Gamma\left(3,\dfrac{1}{6}\right)$	$\dfrac{1}{2}$	$\dfrac{1}{12}$
Power(-1)	∞	
exp(1)	$\dfrac{1}{2}$	$\dfrac{1}{4}$
$N\left(\dfrac{1}{2},\dfrac{1}{12}\right)$	$\dfrac{1}{2}$	$\dfrac{1}{12}$

Team size distribution	μ	σ^2
Constant	3	0
Uniform(0,6)	3	3
$\Gamma(3,1)$	3	3
Power(-1)	$\dfrac{\ln 20}{1-20^{-1}}\cong 3.15$	$20-\left(\dfrac{\ln 20}{1-20^{-1}}\right)^2\cong 10.06$
$\exp\left(\dfrac{1}{3}\right)$	3	9
$N(3,3)$	3	3

---

[1] Stanley M. Besen & Sheila N. Kirby, *Private Copying, Appropriability, and Optimal Copyright Royalties*, 32 J. L. & ECON. 255 (1989).

[2] Besen and Kirby provide what history reveals to be a somewhat ironic example: the argument put forward by Hollywood production companies that the VCR would severely diminish industry profit. *See id.* at 255, n.1.

[3] One of their examples here is the Home Recording Act of 1982, legislation that proposed to compensate copyright holders for consumer copying by levying a tax on "home recording devices and media." *See id*. at 255, n.2.

[4] *See* Daniel K. Benjamin & Roger C. Kormendi, *The Interrelationship Between Markets for New and Used Durable Goods*, 17 J. L. & ECON. 381 (1974); Stanley M. Besen, *Private Copying, Reproduction Costs, and the Supply of Intellectual Property* (The RAND Corporation, N-2207-NSF, 1984); Stanley M. Besen, *Private Copying, Reproduction Costs, and the Supply of Intellectual Property*, 2 INFO. ECON. & POL'Y 5 (1986); William R. Johnson, *The Economics of Copying*, 93 J. POL. ECON. 158 (1985); Stan J. Liebowitz, *The Impact of Reprography on the Copyright System*, Consumer and Corporate Affairs Canada (1981); Stan J. Liebowitz, *Copying and Indirect Appropriability: Photocopying of Journals*, 93 J. POL. ECON. 945 (1985); Ian E. Novos & Michael Waldman, *The Effects of Increased Copyright Protection: An Analytic Approach*, 92 J. POL. ECON. 236 (1984); Janusz A. Ordover & Robert D. Willig, *On the Optimal Provision of Journals qua Sometimes Shared Goods*, 68 AM. ECON. REV. 324 (1978).

[5] *See, e.g.*, Hal Varian, *Buying, Sharing, and Renting Information Goods* (mimeo, University of California, Berkeley, School of Information Management and Systems, May 1997); Ram D. Gopal & G. Lawrence Sanders, *Preventive and Deterrent Controls for Software Piracy*, 13 J. MGMT. INFO. SYS. 29 (1997); Fernando Nascimento & Wilfried R. VanHonacker, *Optimal Strategic Pricing of Reproducible Consumer Products*, 34 MGMT. SCI. 921 (1988).

A separate line of scholarship focuses on potential network benefits. *See, e.g.,* Kathleen R. Conner & Richard P. Rumelt, *Software Piracy: An Analysis of Protection Strategies*, 37 MGMT. SCI. 125 (1991); Michael Katz & Carl Shapiro, *Network Externalities, Competition, and Compatibility*, 75 AM. ECON. REV. 424 (1985).

[6] *See, e.g.*, Besen & Kirby, *supra* note 1, at 271 ("[The effect] on producer . . . welfare depend[s] on the relative costs of forming clubs and producing originals. If cheaper copies can be substituted for more expensive originals, . . . producers . . . are better off.").

[7] Obviously, Besen and Kirby's specific results were presented in a more narrow, nuanced, and technical fashion. We summarize the core ideas here, and refer readers to the original article for a more carefully delimited presentation.

[8] For example, the Software Publisher's Association publishes an annual report estimating the profit implications of software piracy. The most recent study -- available at <www.spa.org> -- claims that piracy cost the industry over $11 billion in 1997 alone.

[9] *See, e.g.*, National Information Infrastructure Copyright Protection Act, S. 1284, H.R. 2441, 104th Cong. (1995) (proposing new anti-piracy copyright provisions); No Electronic Theft Act, Pub. L. 105-147, Dec. 16, 1997, 111 Stat.

2678 (West 1998) (imposing for the first time criminal penalties on consumers who engage in non-commercial, consumer-to-consumer software piracy).

[10] Several examples are given *supra* note 5.

[11] Earlier papers use the term "indirect appropriability" to refer to this increase in consumer willingness-to-pay. The term is often contrasted to "direct appropriability" -- the assumption that sellers can only capture the value direct purchasers place on information goods. Leibowitz was first to carefully explore the case in favor of indirect appropriability, *see* Leibowitz, *supra* note 4.

[12] *See* Besen & Kirby, *supra* note 1, at 261 n.16 ("The idea that copies are sold [in a market] may be metaphorical."). *See also* Varian, *supra* note 5, at 2-3 (adopting market model); Gopal & Sanders, *supra* note 5, at 32-33 (same).

[13] All of the papers listed *supra* note 5, for example, implicitly or explicitly adopt the market metaphor.

[14] *See, e.g.*, Besen & Kirby, *supra* note 1, at 257 ("In analyzing the determinants of the size of these clubs, we explore the balance between the decreasing average cost of acquiring an original as club size increases and the associated increase in 'congestion costs.'"); Varian, *supra* note 5, at 9-12; Gopal & Sanders, *supra* note 5, at 32.

[15] *See, e.g.*, Besen & Kirby, *supra* note 1, at 265 ("All clubs have the same optimal size."); Varian, *supra* note 5, at 2; Gopal & Sanders, *supra* note 5, at 33.

[16] Varian, *supra* note 5, at 3.

[17] We discuss the issue of correlated valuations *infra* part II and again *infra* part III.C.

[18] Richard Schmalensee, *Gaussian Demand and Commodity Bundling*, 8 J. BUS. S211 (1984).

[19] Articles analyzing bundling in this way include: George J. Stigler, United States v. Loew's, Inc.: *A Note on Block Booking*, 1963 SUP. CT. REV. 152, which first demonstrated the effect with a simple example of two negatively correlated goods; Schmalensee, *supra* note 19, which showed that even two positively correlated goods could become more profitable by bundling; McAfee et al., *Multiproduct Monopoly, Commodity Bundling, and Correlation of Values*, 114 Q. J. ECON. 371 (1989), which provided a set of general conditions for the profitability of mixed bundling of two goods; and Yannis Bakos & Erik Brynjolfsson, *Bundling Information Goods: Pricing, Profits, and Efficiency*, Working Paper, Sloan School of Mgmt., Massachusetts Institute of Technology (1997), which introduced a set of statistical tools for analyzing large bundles of information goods.

[20] There may be other benefits and costs to bundling, including economies or diseconomies in production, distribution and consumption, (*see, e.g.*, Eppen *et al.*, *Bundling-New Products, New Markets, Low Risk*, 32 SLOAN MGMT. REV. 7 (1991)), just as there are other benefits and costs to sharing, like network externalities, production savings, and coordination costs. We discuss several of the latter effects *infra* part V.

[21] A related comparison of bundling, site licensing, and subscriptions can be found in Yannis Bakos & Erik Brynjolfsson, *Aggregating and Disaggregating Information Goods*, in PROCEEDINGS OF INTERNET PUBLISHING AND BEYOND: THE ECONOMICS OF DIGITAL INFORMATION AND INTELLECTUAL PROPERTY, (1997).

[22] For an interesting contrasting view, see Michael J. Meurer, *Price Discrimination, Personal Use and Piracy: Copyright Protection of Digital Works*, 45 BUFF. L. REV. 845, 851 (1997) ("The sharing of videotapes by friends and family works against price discrimination because it makes it harder for sellers to sort customers into homogeneous groups.").

[23] In contrast, earlier papers assume that each team member makes an equal contribution toward the team's purchase and so a team's valuation equals *n* times the lowest valuation of any team member, where *n* represents team size. The market metaphor naturally leads to this conclusion since it allows individuals to shop for the team with the most attractive terms. *See, e.g.*, Besen & Kirby, *supra* note 1 at 261, 264-65; Varian, *supra* note 5, at 3. As we point out in the Introduction, this process works to deny low-valuation consumers access to the good. Interestingly, this process also yields an allocation under a sharing assumption that is quite similar to that which would obtain under a no-sharing assumption.

[24] Although, the aggregation effect *would* obtain if an organizer for each team could act as a perfectly discriminating mini-monopolist, charging each team member a price equal to that consumer's specific valuation. However, this begs two important questions: (1) the question of how the team leader would have such detailed knowledge of team members' valuations; and (2) the question of how the mini-monopolist would avoid competition and hence maintain monopoly power. One answer to both questions might be that the team members are all part of a common social group like a family, an answer that brings us back full-circle to the social sharing assumption.

[25] In Besen & Kirby, *supra* note 1, for example, all teams have access to an identical technology for making copies and, in equilibrium, all teams are the same size because access to copies is the raison d'etre for the teams. A similar conclusion is reached by, *e.g.*, Varian, *supra* note 5, at 2, and also Gopal & Sanders, *supra* note 5, at 33.

[26] In fact, the prior papers (Varian, *supra* note 5; Besen & Kirby, *supra* note 1) have shown how sharing is even more likely to increase profit where the marginal costs of production are positive. For now, we assume this benefit away so as to better focus on the direct revenue effects of demand aggregation. We consider the role of marginal cost *infra* part IV.

[27] Earlier models, by contrast, assumed that consumers with low valuations would be excluded from teams, or that consumers team together exclusively with other consumers like themselves. *See, e.g.*, Varian, *supra* note 5, at 3.

[28] Sharing might similarly increase the value of the good that is shared. For instance, in our software hypothetical, members of the group might assist one another in their use of the good or induce any of a number of network externalities that come from having several consumers using an identical program. These advantages, however, are often not benefits of sharing *per se*; they could be achieved simply by purchasing multiple copies of the good. Thus, while we do not dispute their existence, we do not focus on these beneficial effects of group consumption.

[29] *But see, e.g.*, Besen & Kirby, *supra* note 1, at 265 (teams are all the same size in equilibrium); Varian, *supra* note 5, at 2 (same); Gopal & Sanders, *supra* note 5, at 32 (same). A few prior papers mention that constant team sizes are unlikely, but none directly pursue the issue. *See, e.g.*, Leibowitz, *supra* note 4, at 949 ("I do not believe that anyone would deny that the number of copies made from each original varies.").

[30] *See* Bakos & Brynjolfsson, *supra* note 20, at 5.

[31] *See supra* notes 1, 4, 5 & surrounding text.

[32] More precisely, as long as the profit function is continuous in the neighborhood of the optimum, sharing can increase profit even where the sharing itself is technologically inefficient.

[33] *See* sources cited *supra* note 20.

[34] As mentioned *supra* note 29, most prior work never reaches this interesting case. Those papers implicitly or explicitly assume that each information good is enjoyed by the same number of consumers. Such an assumption would be plausible if sharing were a market phenomenon, but seems to contradict common experience with respect to small-scale sharing.

[35] This is a valid probability mass function, as it is easy to verify that $g(n) > 0$ for all n, and $\sum_{-\infty}^{\infty} g(n) = 1$.

[36] For instance, consider the constant elasticity demand curve $\log q = a - \log p$, over some relevant range of support, as is commonly used in empirical work. For a small enough value of a or large enough support, the ratio of profit to total demand can be made arbitrarily close to zero.

[37] In part II, we showed that sharing in teams of size two would increase profit by just under 9%, so our results above illustrate Corollary 1a: given a constant team size distribution, sharing in larger teams is generally more profitable than sharing in smaller teams.

[38] Simulations with larger mean team sizes fairly rapidly approach 1 for most distributions, as theory would predict. For instance, with mean team size equal to 10, the profit ratio for sharing a good with uniformly distributed consumer valuations in teams that are also uniformly distributed rises to 0.9848 and in teams with mean size 20 it rises to 0.9920 (versus 0. 9597 in table 1, where mean team size is 3).

[39] STATISTICAL ABSTRACT OF THE UNITED STATES, Table 69, *Household Characteristics, by Type of Household* (using data from the March 1995 Current Population Survey).

[40] Erik Brynjolfsson, *The Contribution of Information Technology to Consumer Welfare*, 7 INFO. SYS. RES. 281, 291 (1996).

[41] The relevant resources are listed, and their conclusions summarized, *supra* note 20.

[42] *See* McAfee et al., *supra* note 20.

[43] *Id*. However, this is not always true. For instance, if the density of buyer valuations has two peaks, it can be profitable to segregate the high- and low-value consumers and sell only to the high-valuers. In this case, a positive correlation can increase profit by facilitating this sort of price discrimination.

[44] This phenomenon has some similarities to "ski lift pricing" as analyzed by Robert J. Barro & Paul M. Romer, *Ski-Lift Pricing, with Applications to Labor and Other Markets*, 77 AM. ECON. REV. 875, 882 (1987). There, the authors point out that differences in queue lengths and the resulting congestion costs will often lead consumers to naturally sort themselves into nearly-optimal groupings and consumption levels, even when prices are relatively insensitive to consumer valuations.

[45] *See supra* note 1 and accompanying text.

[46] This contrasts with the role marginal costs typically play in commodity bundling. A bundler incurs additional costs every time he sells a bundle to a consumer who would have purchased only a subset of the goods. Sharing, by contrast, does not impose an additional cost on the seller (it costs a book publisher nothing to make his book sharable) and often decreases total cost by decreasing the number of copies of the good required.

[47] Phrased another way, in certain contexts consumers can produce/distribute information goods more cheaply than can producers. The savings is likely shared by producers and consumers alike. As we have pointed out, Besen and Kirby explicitly model this effect; so does Varian, *supra* note 5, at 6-9. *But see* Johnson, *supra* note 4, at 161 (arguing that copy costs would almost always exceed the costs of legitimate production).

[48] Internet service provider America Online (AOL) uses a similar degradation strategy to influence team formation, in this case specifically trying to encourage sharing within families while discouraging sharing more generally. AOL does this by selling Internet access in the following unique way: for the same total fee, users can purchase anywhere from one to five separate AOL identities, the only limitation being that AOL will not allow any two of those identities to be in use simultaneously. For families, this limitation is often of little import. Even without it, multiple family members would rarely be able to access the Internet concurrently since few families have the requisite number of phone lines and computers. The limitation does discourage sharing more generally, however, since friends and neighbors can otherwise access the Internet simultaneously and hence would find the limitation restrictive.

An argument related to our degradation argument is presented in R. J. Deneckere & R. P. McAfee, *Damaged Goods*, 5 J. ECON. & MGMT. STRATEGY 149 (1996), where the authors suggest that a producer can beneficially price discriminate between high- and low-valuing consumers by selling two versions of a product, one in perfect condition, the other intentionally damaged. They cite the example of two IBM printers with identical components except that one has an additional chip which *slows down* printing, enabling IBM to segment the market by demand for print speed.

[49] This setting could be extended to allow for multiple or continuous consumer valuations.

[50] This could be due, for example, to a lack of access to customer service or the original media. Different assumptions would be reasonable as well. For instance, the purchaser of the original, typically the high valuation user, may suffer most of the utility loss, e.g., because of fear of prosecution; or both users may suffer the utility loss, as where they both need to coordinate and administer the sharing arrangement. It is also possible that the degradation is a fixed amount subtracted from consumers' valuations, or even a fractional loss. For instance, a fixed loss might better represent the case where missing components, such as the original documentation, can be separately purchased.

[51] This condition must hold in almost any market for information. Without market power, a seller would only be able to charge marginal cost. Since information goods are often produced at substantial fixed cost but near-zero marginal cost, information sellers would leave any such purely competitive market. For further discussion, see Hal Varian, *Pricing Information Goods*, in SCHOLARSHIP IN THE NEW INFORMATION ENVIRONMENT, PROCEEDINGS FROM HARVARD LAW SCHOOL'S RESEARCH LIBRARIES GROUP SYMPOSIUM (May 1995); *but see* Besen & Kirby, *supra* note 1, at 265 n.21 (noting that, in their setting, sellers can profit from sharing even in the absence of market power).

[52] In these cases, the models developed Varian, Benjamin & Kormendi, Besen & Kirby, and others seem more apt. See *supra* notes 1, 4, & 5.

[53] *See* Bakos & Brynjolfsson, *supra* note 20, at 8.

[54] In particular, our choice sets the mean of the power function approximately equal to the mean of all the other functions. Note that the truncated distribution function is $F(v) = \dfrac{1 - v^{-1}}{1 - b^{-1}}$, where b is the upper bound and 1 is the lower bound. $E(v) = \displaystyle\int_{1}^{b} v \dfrac{v^{-2}}{1 - b^{-1}} \, dv = \dfrac{\ln b}{1 - b^{-1}}$.