

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/239932557>

Discovering Digital Content in the Age of Peer-to-Peer

ARTICLE

READS

14

2 AUTHORS, INCLUDING:



Jeff Ubois

John D. and Catherine T. MacArthur Foundati...

12 PUBLICATIONS 23 CITATIONS

SEE PROFILE

Discovering Digital Content in the Age of Peer-to-Peer

Paul Laskowski
UC Berkeley
paul <at>
ischool.berkeley.edu

Jeff Ubois
Fujitsu Labs of America
jeff <at>
ubois.com

David Marvit
Fujitsu Labs of America
dave <at>
marvit.org

John Chuang
UC Berkeley
chuang <at>
ischool.berkeley.edu

ABSTRACT

The rise of peer-to-peer networks has dramatically altered the way consumers interact with digital content. Most visibly, pirate downloads substitute for legal purchases, eroding profits and calling into question the very viability of content production. This characterization has motivated a detect-and-punish model that aims to make pirate networks less attractive. Less apparent is the way peer-to-peer transforms the purpose of each download and the way consumers *discover* the content they want. Moreover, recent empirical evidence suggests that discovery processes play a commercially significant role, even providing a boost to sales that can sometimes overwhelm the direct substitution effect. While such effects bear strongly on the design of next-generation business models, they are poorly understood on a microeconomic level. We investigate how consumers explore a space of digital content with a series of analytical models. Our work leverages standard techniques in economics that represent content as points in a metric topology. We find that the structure of this content space strongly influences social outcomes, including the variety and price of content. Moreover, in an environment of content discovery, conventional assumptions of behavior may be reversed. Lower prices may decrease revenues, but simultaneously increase content production, as the marginal good attracts greater interest. We discuss how peer-to-peer may influence the long-tail distribution of content, and the implications for producers into the future.

Categories and Subject Descriptors

C.2.4 [Computer-Communication Networks]: Distributed Systems; J.4 [Social And Behavioral Sciences]: Economics

General Terms

Economics, Human Factors, Theory, Legal Aspects.

Keywords

Discovery, Digital Content, File-sharing, Peer-to-peer, Sampling, Learning, Exploration, Content Production.

1. INTRODUCTION

In the few years since their conception, peer-to-peer networks have dramatically altered the way consumers acquire, conceptualize, and interact with digital content. This transformation has been felt most keenly by content

producers, who now find that perfect copies of the products they sell are available on file-sharing networks. Many consumers who would otherwise have purchased these goods will instead opt to download the “pirated” versions for free, cutting into revenues. This *substitution effect* is a central focus of today’s public debate, and has called into question the very viability of content production in the broadband era [10].

A focus on how pirated copies substitute for legal content also serves to motivate a detect-and-punish strategy on the part of producers. This is exemplified by the Recording Industry Association of America (RIAA), which has issued lawsuits on behalf of copyright holders against 20,000 individual music sharers. Most of those targeted have settled for sums ranging from \$2,500 to \$7,500. More recently, the music industry has announced plans to end these mass lawsuits, and instead rely on Internet Service Providers (ISPs) to disconnect users who share copyrighted music from the internet.

One criticism of these detect-and-punish strategies is that many illegal downloads represent content that the user would never have purchased – or perhaps even known about – were it not for the file-sharing network [6]. As this suggests, to focus exclusively on the substitution effect is to oversimplify the impact of peer-to-peer technologies. A less visible, but highly relevant consideration is the way peer-to-peer transforms the purpose of each download and the way consumers *discover* the content they want.

As many commentators have noted, consumers may use file-sharing networks to sample a piece of content before legally purchasing it. They may also use the content to learn about an artist, before acquiring a larger album or collection. More generally, the user may download files in order to explore the larger content space, perhaps ending up somewhere rather different than the starting point. These activities, which we will refer to as *discovery processes*, are certainly not new, but file-sharing networks have enabled them on a much larger scale. Moreover, recent empirical evidence suggests that discovery processes can play a commercially significant role.

In one study, Brynjolfsson et al. compare sales data between a paper catalogue and a retail website [2]. Both channels offer identical goods, but the website features a variety of tools to help consumers navigate the product

space. The authors find that online sales tend to favor more specialized products, resulting in a “long-tail distribution.” Content discovery does not just affect the distribution of sales, however, it may also affect sales volume. Oberholzer and Strumpf examine the effect of file-sharing on legitimate music sales [7]. By examining data on the individual song level, the authors argue that any effect on sales due to file-sharing is statistically insignificant. The substitution effect, they speculate, is cancelled out by the beneficial effects of sampling and learning. In a third study, Smith and Telang examine how much DVD sales are depressed when a movie is broadcast on television [13]. On the face of it, seeing a free movie might be expected to depress sales much more than hearing a free song, since people rarely watch movies multiple times. It is all the more surprising then, that broadcasting a movie *enhances* sales – providing a boost of over 350% – even though piracy of the film also increases.

The latter examples suggest that discovery processes can stimulate sales under certain conditions, even to the point of overcoming the direct substitution effect. This makes content discovery highly relevant to a discussion of next-generation business models. While detect-and-punish has done much to antagonize consumers, it has so far failed to stamp out piracy. Wong et al. note that discouraging interest in piracy often does little to affect the availability of pirated content [15]. Simulations by Altman et al. suggest that music labels may maximize profits by abandoning all investment in punishment [1].

Meanwhile, numerous ideas have been floated as to how to fund content production in the future. Several music companies, including Popcuts [9], rVibe, and Grooveshark, reward music buyers by giving them a portion of future sales revenue [4]. Recent startup Sellaband allows users to pledge money towards an artist after listening to his or her demo materials [12]. If contributions exceed the \$50,000 threshold, they allow the artist to record an album. Frost argues for a music industry structure without record labels, in which lower prices are the key to fighting piracy [3]. He estimates that an album of music should cost about \$3.

The success of these business models may turn on how they interact with the way consumers discover content. Unfortunately, this process of exploration remains poorly understood on a microeconomic level. Analytic models tend to feature consumers that know the value of all products, leaving no opportunity for exploration [14]. In the closest study we are aware of, Peitz and Waelbroeck argue that free sampling may enhance producer profits by matching users to more preferred content [8]. Sampling, however, is modeled as all-or-nothing, and consumers are only allowed to buy one product, limiting visibility into the mechanics of the discovery process.

We believe that a more complete analytic understanding of how consumers discover digital content may yield a variety

of potential benefits. Such an understanding can form a useful complement to empirical work, shedding light on observed revenue effects. Analytic modeling may also highlight the social benefits of exploration that have been enabled by file-sharing networks, and provide a useful perspective to balance the public debate. Finally, we hope that a more detailed understanding of discovery can one day guide the design of next-generation business models, identifying those that align the way consumers discover content with producer incentives.

Motivated by these objectives, this paper will construct two economic models to examine different aspects of discovery behavior. Our analysis will leverage popular tools in economics that represent content and user preferences as points in a metric topology. Comparing two separate models will allow us to highlight the important role played by the structure of the content space. We will also identify ways in which content exploration can influence traditional assumptions of behavior, and find relevant insights for today’s public debate.

2. SAMPLING AND AN UNSTRUCTURED CONTENT SPACE

In order to describe content discovery, we need a notion of consumer tastes, and how these are related to different products. Using a technique pioneered by Hotelling, we visualize a topological space of products [5]. Goods are represented as points within a metric topology, M , with distance function d representing how different two products are. For example, M might represent music within a particular style, or movies within a given genre. A consumer’s personal taste may also be represented by a point in M , so that the consumer’s utility from consuming a good is a decreasing function of the distance to that good.

Following the well-known model of Salop, we first let M be the circle of unit circumference, S^1 , equivalent to the real interval $[0,1]$, in which the endpoints have been topologically identified [11]. Each good and user’s taste is represented by a point in S^1 , and we’ll assume the simplest linear valuation function. As in the model of Peitz and Waelbroeck, let the valuation of a good located at v for a user of type t be $m - |t - v|$, where $m \leq 1/2$ represents the maximum value a user can gain from a perfectly aligned product [8]. This is represented graphically in Figure 1. After a user downloads a set of goods, $\{v_i\}$, she enjoys utility equal to the value of the most preferred good (or zero if no good has a positive valuation): $u = \max(0, v_i)$. This form emphasizes the fact that users have limited time to consume digital content, and may only enjoy a fraction of the collection they own.

Our intention is to leave this content space rather unstructured. We will assume that the location of each product is drawn independently from a uniform distribution.

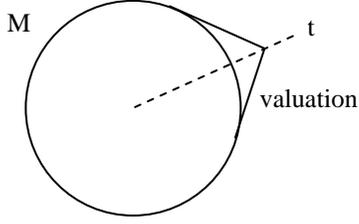


Figure 1: Valuation on the Circle Topology

Furthermore, a user does not know where in M a good lies until she obtains it. This structure is chosen so that consumers have very limited opportunities for learning. Specifically, a user that purchases a good only learns the value of that specific good. In this way, this model isolates a particular aspect of discovery that corresponds to a common notion of sampling.

We assume that there are k products, and a user must pay a cost, s , in order to obtain each one. This cost includes any price that is paid to the producer in the case of a legitimate purchase. It may also include any subjective costs of the purchase, such as the time and effort invested in using a peer-to-peer network.

2.1 Consumer Behavior

Given these essentials, we may compute the behavior of a rational consumer. Suppose that the user currently has a selection of goods with maximum valuation u_0 . In order for a new good to improve the user's valuation, it must fall within the range, $(t - (m - u_0), t + (m - u_0))$, which occurs with probability $2(m - u_0)$. If this does happen, the expected improvement to the user's valuation is half of the maximum improvement: $(m - u_0)/2$. Thus, the expected utility gain from obtaining this good is $(m - u_0)^2 - s$. This is positive as long as $u_0 < m - \sqrt{s}$. The user can always decline future goods, so as long as this single good provides positive expected utility, the possibility of further purchases can only improve the expected payoff from purchasing it, and it will be rational to obtain it. On the other hand, if $(m - u_0)^2 - s < 0$, we know that u_0 can only increase in each round, so every potential future purchase will yield negative expected payoff. It can then only be rational not to obtain the good.¹

We see that users will continue to obtain goods until $u_0 > m - \sqrt{s}$, or a good falls within the range, $R = (t - \sqrt{s}, t + \sqrt{s})$. The chance of any one good falling in this range is $2\sqrt{s}$. The user will obtain the n th good if

¹ We assume that $m - \sqrt{s} > 0$, which ensures that the user wants to obtain the first good.

none of the previous $n - 1$ goods fall in R , and the probability of this occurring is $(1 - 2\sqrt{s})^{n-1}$. The expected number of purchases for each user may be computed as the sum of these probabilities:

$$1 + (1 - 2\sqrt{s}) + \dots + (1 - 2\sqrt{s})^{k-1} = \frac{1}{2\sqrt{s}} - \frac{(1 - 2\sqrt{s})^k}{2\sqrt{s}} \quad (1)$$

As the cost s decreases, a user will tend to sample more, and therefore end up with a product that offers greater value. Specifically, when $s = 0$, the user will purchase all k products. If we assume that the entire cost s is comprised of a payment to the content producer, then social welfare is determined by the value each consumer derives from their favorite product. Thus, decreasing s increases social welfare. It may also be checked that decreasing s enhances consumer surplus.

We can make these observations more precise by considering what happens if there is an inexhaustible supply of products. In the limit as k grows without bound, the expected number of purchases per user simplifies to, $1/2\sqrt{s}$. The expected valuation of the last good a user obtains is now $m - \sqrt{s}/2$. This is also the social welfare if all of the cost is transferred to the producer as payment. Finally, the expected payoff to the consumer is $m - \sqrt{s}/2 - s(1/2\sqrt{s}) = m - \sqrt{s}$.

As the cost decreases, users earn greater surplus. Furthermore, the second derivative of consumer surplus with respect to s is positive. Roughly speaking, this means that each successive drop in price benefits consumers more than the last one.

2.2 Producer Behavior

While decreasing s is beneficial to welfare, producers may have an opposing interest, as the following claim shows. For this claim, we will assume that the entire cost s is in the form of payment.

Claim 1. A monopolistic producer will always select the maximum price such that consumers purchase any products, $s = m^2$.

Proof. The monopolist's payment is the price times the expected number of purchases,

$$s \left(\frac{1}{2\sqrt{s}} - \frac{(1 - 2\sqrt{s})^k}{2\sqrt{s}} \right) = \frac{1 - (1 - 2\sqrt{s})^k}{2} \sqrt{s} \quad (2)$$

This increases in s , so the monopolist will choose the maximum price. \square

In general, a monopolist will also have control over how many products to supply. Suppose that each product costs c to produce and the number of consumers is n . Then the following claim holds,

Claim 2. Given a fixed price, s , the number of goods produced by a monopolist is

$$k = \left\lceil 1 + \frac{\ln(ns/c)}{\ln\left(\frac{1}{1-2\sqrt{s}}\right)} \right\rceil \quad (3)$$

Proof. We may number the goods that are produced, $1, 2, \dots, k$, and assume without loss of generality that all consumers obtain the goods in this order. Then the producer expects profits from good k equal to

$$ns(1-2\sqrt{s})^{k-1}. \quad (4)$$

If each good costs c to produce, the producer will pay for the k th good as long as $s(1-2\sqrt{s})^{k-1} > c/n$. The left hand side is an exponentially decaying function. Solving for equality gives the required expression. \square

In general, the monopolist will tend to undersupply the market. A particularly dramatic demonstration of this occurs if we let $m=1/2$. This is a natural value to consider, as it implies that exactly one point in the product space offers zero value to a user – the antipodal point diametrically across the circle.

Corollary. For $m=1/2$, a monopolist will produce just a single good, and charge the maximum viable price, $s=1/4$.

This may be verified by substituting into equation (3). Our negative results stand in contrast to the circle model of Peitz and Waelbroeck, in which a producer may sometimes encourage consumers to freely sample products before buying [8]. For this result, however, the authors bound their parameters such that a user expects negative utility from a random free download. For more moderate parameter values, both our models predict that a producer will want to discourage sampling.

2.3 Restricting the Price

We may also consider what happens to production if the price s is restricted to a value below the maximum price. This may be the result of regulation, for example. Alternately, we may imagine that the price must stay below a certain level or consumers will switch to using free file-sharing networks. Finally, under today's industry structure, third parties, including major distributors like iTunes, may wield considerable influence over the selection of prices.

Interestingly, decreasing s below the maximum level always lowers the monopolist's profits, but the rational level of production may actually increase. Examining equation (4), we see that decreases in s always reduce the profit from the first good produced. On the other hand, the decay of profit

with each successive good also decreases as consumers sample more. This means that while the monopolist makes less profit on the whole, the *marginal* product may become more profitable. As a result, lowering s may increase production. This effect will only occur up to a point: As s decreases to c/n , production will drop to a single good.

This feature is clearly illustrated when $m=1/2$. In this case, we saw that an unregulated monopolist will only produce a single good. Restricting s to a lower level will generally increase the level of production. In the limit as n/c grows large, an arbitrary level of production may be induced, and social welfare may be doubled.

These observations lend some support to Frost's proposal for selling music more cheaply [3]. As he points out, profits from every album sale have stayed relatively constant in the transition from physical formats to digital downloads. Our model allows us to speculate that moderate decreases in price might not diminish the supply of music, as many fear. Instead, the low prices suggested by Frost might conceivably *increase* content production. Meanwhile, the same strategy will mitigate the price advantage of illegal file-sharing networks, reducing the incentive for piracy. Further research is needed to calibrate our model and estimate what prices optimize content production.

2.4 Subscription Charges

Rather than charging consumers for each product downloaded, some business models envision charging consumers a fixed subscription fee for unlimited content. As long as consumers are identical (differing only in type), this will be socially beneficial in our model. A monopolist may extract all available surplus, by charging the expected value each user obtains from the best possible product. This aligns the monopolist's incentives with those of society in general, so that social welfare will be maximized.

Of course, a more realistic model would feature heterogeneous consumers, with a distribution of valuations for digital content. Any positive subscription fee charged by the monopolist will tend to price some users out of the market, resulting in deadweight loss and reducing social welfare. On the other hand, those users who buy the subscription will sample aggressively to maximize their utility, providing a countervailing boost to welfare. The net effect depends on exactly how consumers are distributed. Further research may characterize consumers in greater detail, shedding light on when each pricing strategy is advantageous.

3. MOVING THROUGH A CONTENT LATTICE

The content space in the previous section was deliberately unstructured in order to highlight sampling behavior. In this section, we will be interested in some other, more

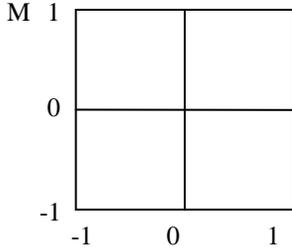


Figure 2: Representation of Lattice Topology

sophisticated aspects of content discovery. We will want to describe how users might learn about the greater content space, and how previous purchases can influence the selection of future content. To do this, we will leverage a more highly structured topological space.

Let M be a lattice in \mathbb{R}^N , consisting of the points, $\{x = (x_1, \dots, x_N) \mid x_i \in \{-1, 0, 1\}\}$. Two points are connected in this space if they differ only in one dimension, and by exactly 1. Each dimension of this space is meant to represent a separate characteristic of the good that consumers care about. For example, one dimension could model whether a song is simplistic or intricate. Another, whether it is energetic or downbeat.

This topology is actually a simplified version of the one used by the Music Genome Project to represent songs. That topology, which simply includes more points than ours, has found commercial application, being used by music service Pandora to identify music similar to a specific song.

We will assume that consumers have a preference over each dimension in our topology, so a consumer type is randomly drawn from the corner points, $\{v = (v_1, \dots, v_N) \mid v_i \in \{-1, 1\}\}$. We assume a consumer of type v values a good of type x by

$$m - d^2(v, x), \quad (5)$$

where distance function d is the length of the shortest path along the lattice.

This structure incorporates a notion of mainstream and fringe goods. Notice that all users value the product at the origin by exactly $m - N^2$. If a good, x , differs from the origin in exactly j dimensions, the average value over all users is $m - N^2 - j$. Thus, as we move along the lattice from the origin out to a corner, products become less appealing to the population on average, even though there is a particular type of consumer that likes that particular region of the content space.

3.1 Consumer Behavior

We will assume that a user's view of this topology is limited. Suppose that only the most mainstream product can be identified by users at the beginning (we may imagine that it is heavily promoted). After this point, we will

assume that a user can only identify goods that are adjacent to those she already owns. The user may not know that other goods exist, or may not have information about where these other goods are located in the topology.

We may alternately assume that a user has perfect knowledge of where goods fit into the topology, but does not know her own tastes. Each time the user purchases a good, she becomes aware of one axis that she would prefer to move along to enhance her valuation.

In either case, a user will move from the origin towards a corner, bringing one dimension of the good into line with her preference at each step. The product at the origin provides value, $m - N^2$, and obtaining the n th good after this enhances the value by $(N - n)^2 - (N - n - 1)^2 = 2(N - n) - 1$. This is a decreasing linear function. Thus, a user will purchase the n th good as long as,

$$2(N - n) - 1 \geq s, \quad (6)$$

or,

$$n \leq N - \frac{s+1}{2} \quad (7)$$

As s decreases, we can see that users will move from the mainstream content farther out to the fringe content at the corners, and end up liking the product they end up with more as a result. If we assume, as before, that the entire cost s is transferred to a producer as payment, decreasing s will increase welfare. This closely mirrors our observations from Section 2.

Because the lattice model includes a notion of mainstream and fringe content, we can observe the movement of users visually. To do this, we slightly modify the consumer valuation in equation (5) to $m - Zd^2(v, x)$, where Z is a random variable that adds variation. We may then graph the fraction of users that purchases a good a distance n from the origin. This is shown in Figure 3 for a 10-dimensional lattice when Z is drawn from a normal distribution with standard deviation 0.2. As the price is decreased from 18, to 12, and to 6, more users buy the fringe products to the right. This result depicts the formation of a so-called long-tail distribution.

3.2 Producer Behavior

While the lattice and circle models were similar in terms of consumer behavior, differences do appear when we consider the producer's perspective. Assuming that the entire cost s is absorbed by a monopolist producer as payment, the producer faces a linear demand curve, and earns revenue per user,

$$s \left[N - \frac{s+1}{2} \right] \quad (8)$$

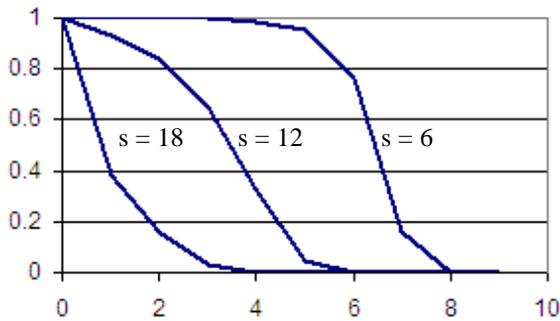


Figure 3: Distribution of Consumers in Lattice

If we assume for the moment that content is costless to produce, and ignoring the whole integer problem, the monopolist will opt to maximize revenues by setting $s = N - 1/2$. This is a moderate price that encourages users to explore about halfway from the mainstream to the fringe. Because user utility is based on the square of distance, however, this traverse actually captures about $3/4$ of the value available to the user. – a marked improvement over the unstructured circle topology.

4. DISCUSSION

By describing content discovery across two different topological models, we have witnessed the important role that the structure of the content space plays in determining behavior. We saw in the circle topology that a monopolistic producer would always choose the maximum viable price. Furthermore, under natural specifications, the producer only supplies a single product, leaving no opportunity for further discovery. In the lattice model, by contrast, the monopolist tends to set a moderate price that encourages some exploration, and captures a substantial fraction of the utility available to consumers. More research is needed to see how realistic content spaces are structured, and how analytic models may be calibrated for added realism.

Our analysis only considers profits resulting directly from selling goods to each user. A large part of the file-sharing debate, however, concerns the degree to which related revenue streams can make up for the loss in direct sales. For example, musicians can supplement their sales income through concert sales, merchandizing, and promotions. Even when lower prices yield less revenue in our models, they still result in users that like the music they end up with more. This feature can be expected to enhance profits from auxiliary revenue channels. More ardent fans may pay more for concerts, or be more interested in buying t-shirts. In practice, modeling the tradeoff between different revenue sources is difficult and will likely require a more data-driven approach than we have adopted here.

We have suggested that discovery can shift interest from mainstream content to more specialized content, driving the formation of a long-tail distribution. In turn, such a

distribution can be an enabling factor for a variety of potential business models. Smaller, more enthusiastic audiences may be more likely to invest in musical artists, or recommend content to friends. They may experience a sense of community that discourages piracy. Revenue sharing strategies may be more robust when fewer people are required to cooperate. Niche content is also harder to find on file-sharing networks than mainstream hits, making it less of an attractive alternative.

In order to provide a clear description of discovery processes, we have aimed to keep our models simple, focusing on a single distribution system. A lot of interesting aspects of today’s industry, however, emerge when multiple distribution systems interact. Of course, pirate file-sharing networks may draw demand away from legal distribution channels. Future revisions of our models may be geared to describing discovery across these two competitive mediums.

In another direction, today’s industry also features multiple *legal* distribution channels, which may interact and compete with each other. As we saw in our lattice model, a monopolist may want to set a moderate price to encourage exploration through a structured content space. Given recent advances in tools for content navigation, a music label may be motivated to encourage greater exploration through its digital distribution channels. Such a strategy may interfere with the label’s traditional, non-digital retail channels, however. Unlike digital downloads, traditional retail is constrained by limited shelf space, limited promotion time on radio, and so forth. This scarcity motivates a strategy of mega-hits, which may be in tension with the features of digital distribution. As online sales continue to increase in importance, however, we may expect a greater transition away from megahits towards niche products. Further research is needed to examine this possibility.

Throughout our analysis, we have seen that discovery processes can enhance social welfare. By moving consumers away from the mainstream towards niche content, discovery can also stimulate a more diverse content industry. We believe that these considerations bear strongly upon today’s public debate, and future policies should be considered in light of their effect on content discovery.

One important aspect of today’s debates concerns the pricing of digital content. File-sharing exerts downward pressure on today’s prices, which is often portrayed as a threat to content production. We were surprised to find, however, that less money for labels does not universally imply less content. Instead, lower prices can make marginal content more profitable, as it attracts greater interest from consumers, and the result can be a greater diversity of digital products.

5. REFERENCES

- [1] Altman, E., Wong, S., & Rojas-Mora, J. P2P Business and Legal Models for Increasing Accessibility to Popular Culture. The International Conference on Digital Business. London, UK. June 2009.
- [2] Brynjolfsson, E., Hu, Y., Simester, D. Goodbye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales. Working Paper Nov 2007. Available at SSRN: <http://ssrn.com/abstract=953587>.
- [3] Frost, R.L. Rearchitecting the music business: Mitigating music piracy by cutting out the record companies. *First Monday*, 12 (8). 2007.
- [4] Hosanagar, K., Tan, Y., and Han, P. Dynamic Referrals in Peer-to-Peer Media Distribution. Working Paper June 1, 2008. Available at SSRN: <http://ssrn.com/abstract=926915>
- [5] Hotelling, H. Stability in Competition, *Economic Journal*, vol. 39, no. 153 (March 1929), 41-57
- [6] Huygen, A., Rutten, P., Huveneers, S., Limonard, S., Poort, J., Leenheer, J., Janssen, K.S., van Eijk, N., Helberger, N. Ups and downs: Economic and cultural effects of file sharing on music, film and games. Ministries of Education, Culture and Science, Economic Affairs and Justice, The Netherlands. 2009.
- [7] Oberholzer, F. & Strumpf, K. The Effect of File Sharing on Record Sales: An Empirical Analysis. *Journal of Political Economy*, 2007.
- http://www.unc.edu/~cigar/papers/FileSharing_March2004.pdf
- [8] Peitz, M. & Waelbroeck, P. Why the Music Industry may Gain from Free Downloading - The Role of Sampling. *International Journal of Industrial Organization*, vol. 24, no. 5 (Sep 2006), 907-913.
- [9] Popcuts. <http://www.popcuts.com/>
- [10] Pouwelse, J., Garbacki, P., Epema, D., Sips, H. Pirates and Samaritans: A Decade of Measurements on Peer Production and their Implications for Net Neutrality and Copyright. *Telecommunications Policy* 32 (11), 701-712. 2008.
- [11] Salop, S. Monopolistic Competition with Outside Goods. *The Bell Journal of Economics*, Vol. 10, No. 1 (Spring, 1979), pp. 141-156
- [12] Sellaband. <http://www.sellaband.com/>
- [13] Smith, M. & Telang, R. Competing with Free: The Impact of Movie Broadcasts on DVD Sales and Internet Piracy. *MIS Quarterly*, (33: 2) pp.321-338.
- [14] Varian, H. The Social Cost of Sharing. Workshop on Economics of Peer-to-Peer Systems, Berkeley, 2003.
- [15] Wong, S., Altman, E., & Ibrahim, M. P2P Networks: The Interplay between Legislation and Information Technology. Networking and Electronic Commerce Research Conference (NAEC). Lake Garda, Italy, October, 2009.