



P. Jean-Jacques Herings,
Ronald Peeters, Michael S. Yang


**Piracy on the internet:
Accommodate it or fight it?
A dynamic approach**

RM/09/034

METEOR

Faculty of Economics and Business Administration
Maastricht Research School of Economics
of Technology and Organization

P.O. Box 616
NL - 6200 MD Maastricht
The Netherlands



Piracy on the internet: Accommodate it or fight it?

A dynamic approach

P. Jean-Jacques Herings* Ronald Peeters* Michael S. Yang*

August 3, 2009

Abstract

This paper uses a dynamic stochastic model to solve for the optimal pricing policy of the music recording companies in the presence of P2P file-sharing networks eroding their CD sales. We employ a policy iteration algorithm on a discretized state space to numerically compute the optimal price policy. The realistically calibrated model reflects the real-world figures we observe and provides estimates of figures we can not observe, such as changes in total welfare. The results suggest that, thanks to the existence of P2P networks, total welfare in 2008 in the U.S. is about \$25.6 billion more per annum than in 1999 before P2P was introduced. Moreover, the results predict that the current trend of decreasing CD sales will continue until around the year 2020 when it will stabilize at around 231.2 million copies per year, comparing to the industry all-time high of 938.9 million in 1999. The comparative static analysis shows that full enforcement of intellectual property rights, although helpful for the industrial profit, may have adverse effect on total welfare.

JEL classification: L11; L82; L86.

Keywords: dynamic price competition; music industry; piracy; P2P networks.

1 Introduction

Ever since Napster emerged in 1999, P2P (peer-to-peer) file-sharing networks have been the center of piracy on the internet. Over the past ten years, we have seen different generations of P2P technology come and go thanks to the legal pursuits launched by various authorities and interest groups alike and the development of new technologies that help evade them. The content providers try to defend their creations from free riding; whilst the online communities try to preserve the true free and sharing spirit of the internet. The battle goes on and on, but we can not ignore the impact it has had on the markets of information goods, particularly the music industry.

Numerous studies, especially empirical ones, over the years have tried to explain the relationship between the decline in music sales and the rise of P2P networks. Blackburn (2004),

*Department of Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands. E-mail: {P.Herings,R.Peeters,S.Yang}@maastrichtuniversity.nl. P. Jean-Jacques Herings and Ronald Peeters would like to thank the Netherlands Organization for Scientific Research (NWO) for financial support. Corresponding author: P. Jean-Jacques Herings.

Liebowitz (2004), Rob and Waldfogel (2006), and Zentner (2006) find that downloading (via P2P file-sharing networks) are at least partially responsible for the recent decrease in CD sales. Liebowitz (2006) concludes that file-sharing has clearly brought significant harm to the recording industry. In contrast, Oberholzer-Gee and Strumpf (2007) and Peitz and Waelbroek (2004) find little robust evidence that P2P has caused the decrease in CD sales for the recent years.

The classical theoretical literature on piracy and network effects include, among others, Economides (1996), Katz and Shapiro (1985, 1986), Johnson (1985), and Takeyama (1994). Recent papers that study P2P in specific include, among others, Gayer and Shy (2003), Bae and Choi (2006), Peitz and Waelbroek (2006) and Herings, Peeters, and Yang (2008). Gayer and Shy (2003) shows how publishers of digitally-stored products, including music, can utilize P2P to enhance sales of their product sold in the store. This result is mainly attributable to the positive consumptive externality ingredient in their model. Peitz and Waelbroek (2006) shows that under sufficient taste heterogeneity and product diversity, the positive effect of downloading on sales due to sampling may compensate the direct negative effect. This result is mainly driven by the information asymmetry between the buyers and the sellers of the product, and downloaded files can help buyers identify their favorite products more easily and hence encourage sales. Herings et al. (2008) analyzes the different market structures that may form. The paper concludes that, depending on different circumstances, the firm will employ pricing strategies either to deter the entry of a network or to accommodate it. Moreover, music industry profits decrease when the generic cost factor of downloading declines, i.e. when the society is more downloader-friendly, whereas total welfare increases.

However, all the theoretical studies on this subject have been of a static nature. This implies that in the presence of demand side externalities there may be multiple equilibria, which ultimately leads to an equilibrium selection of some sort. Consequently, coordination of the network formation is often assumed. The authors of these models always focus on the maximally achievable network size. They lack insights on how such networks actually form, step by step, from zero to a steady-state network size; how the firm prices strategically to compete with the network every step of the way; and how society is affected during the process.

Only a dynamic model can bring such insights, which is what this paper is designed to achieve. Dynamic stochastic models with network externalities have been recently studied in the IO literature, mainly using numerical methods. Examples are Markovich (2008) and Markovich and Moenius (forthcoming) which study the dynamics caused by the iterations between hardware and software; Jenkins et al. (2004) studies a stylized version of the browser war between Netscape and Microsoft, where the entrant may have “grabbed” market shares from the incumbent and thereby tipping the market; Arie and Grieco (2009) investigate the effect of switching costs on market dominance and equilibrium prices; and Chen, Doraszelski, and Harrington (forthcoming) studies competing firms’ incentives to make their products

compatible and the possible effects that may prevent market dominance. A framework for numerically analyzing dynamic interactions in imperfectly competitive industries is proposed by Doraszelski and Pakes (2007), which provides an excellent summary of the main approach for models of this kind.

In this paper, we use a dynamic model to solve for the optimal pricing strategy of a firm that releases music CDs and sells them in the store while being exposed to a competing P2P file-sharing network on the internet. In the model, there is a firm who sets the price of its CDs every period and a continuum of consumers who decide whether to purchase the CD from the store, download the music from P2P, or not acquire the music at all. The timing involves discrete periods with infinite horizon. The firm is forward looking and strives to maximize the present value of all future profits by choosing a state-dependent pricing policy; the state being the market shares of the store, the P2P network, and the unserved market. The consumers make one of the three choices based on the price and the state in order to maximize their utility. Consumers are *ex ante* identical, but receive random utility shocks prior to their purchasing decision every period anew. Switching costs are imposed on the consumers who switch to a product, from previously consuming the other product or not consuming at all. The firm's optimal pricing policy is derived numerically and provides insight in the market share dynamics. The parameters of the model have been calibrated in such a way that the model outcomes (such as prices and sales quantities) match the real-world data in the years 1999 and 2003.

An interesting point of this approach is that we can study how networks develop in a dynamic process without having to make any assumptions on consumer coordination as in the conventional static models of network economics. We can also observe exactly how the firm sets its price conditioned on its market share and that of the network. Sometimes it sets the price very low in order to win vital market shares to fight against the P2P network, while other times accommodates the network by setting a high price to reap the profit from its own installed customer base.

We use the model to predict future sales quantities and to compute the changes in consumers' surplus and total welfare. For instance, it predicts the sales quantity of the year 2008 rather accurately to be around 374.3 million copies, while the real figure is 384.7 million. It also predicts that the CD sales will have dropped to around 231.2 million copies per year around the year 2014, when according to the model estimations the long-run steady state is approximately reached. More importantly, this model is able to estimate the welfare gain over the years brought by the P2P file-sharing networks. In the year 2003, 4 years after P2P was introduced, while the music industry was suffering from a \$1.6 billion forgone profit per year, the total welfare was up by \$13.1 billion per year. By the year 2008, the industry profit has dropped by a further \$1.4 billion, and welfare soared by a further \$12.4 billion per year, making it \$25.6 billion more than in 1999. In the long-run steady state (around 2020), the industry profit will have dropped to a mere \$865.5 million per year, and total welfare are set

to have improved since 1999 by a whopping \$29.7 billion per year.

In the comparative static analysis of the model, one of our most important findings is that total welfare is negatively related to the generic cost factor of downloading. This result coincides with the findings from papers such as Rob and Waldfogel (2006) and Herings et al. (2008) in the sense that the existence of P2P actually enhances total welfare. This implies that by making file-sharing more difficult for the consumers, the government is effectively curbing the society from enjoying a high welfare level.

The remainder of the paper is organized as follows. In Section 2 the dynamic model is described in detail. Next, Section 3 explains how the numerical computations are conducted. In Section 4 the model is calibrated and the parameter values for the benchmark scenarios are set. The model outcomes and predictions are then presented in Section 5, and comparative statics are analyzed in Section 6. Section 7 concludes.

2 The model

Each period, music albums are being offered on CDs by the firm and online via P2P networks. A continuum of consumers decides whether to buy the album at the store (S), to download it via P2P networks (N), or not to acquire it at all (E). A consumer's decision in one period determines her type in the subsequent period. That is, at each period, depending on the decision in the previous period, a consumer is of one of the three possible types: θ_S (store), θ_N (network), or θ_E (empty). The state at a certain period is defined as the distribution of consumers over types. Setting the total mass of consumers to 1, the state space is given by

$$\Omega = \{ (s, n, e) \in \mathbb{R}_+^3 \mid s + n + e = 1 \},$$

where s , n , and e represent the share of consumers of type θ_S , θ_N , and θ_E , respectively. A typical state in Ω is denoted by ω .

Every period, given the current state ω , the firm sets a price for its CD in the store $p(\omega)$.¹ This generates an immediate profit of

$$\pi(\omega, p(\omega)) = (p(\omega) - \mu) \cdot s(\omega, p(\omega)),$$

where μ represents the cost of producing one unit and $s(\omega, p(\omega))$ represents the resulting sales quantity, which equals the number of consumers choosing to buy the album in store given the current state ω and price $p(\omega)$. Hence, in the next period the process will be in a state with $s(\omega, p(\omega))$ consumers of type θ_S . When designing an optimal pricing scheme, the firm realizes that the price chosen in the current state does not only affect the immediate profit, but also the state transition and thus potential profits in the future. We assume the firm to

¹We assume that the firm exhibits stationary pricing behavior, which means that the price depends only on the state. It is well known that our framework possesses an optimal pricing policy in stationary strategies.

be rational and farsighted. That is, in any period, it sets the price as to maximize the present value of the stream of profits discounted by a factor of δ each period.

Given the current market state ω and the firm's price at this state $p(\omega)$, the state transition is completely specified by the consumers' decisions. We assume consumers to maximize their utility in the current state. In her decision, a consumer takes into account the price at the store, the expected downloading costs, and possible switching costs τ . The purpose of the switching costs τ in the model is to promote consumer loyalties to the platforms, reflecting a possible lock-in effect. In our model the expected downloading cost is negatively related to the network size and the switching cost is incurred only when a consumer acquires the product while changing type. Hence, a consumer's decision is determined by her type, the CD price at the store and the expected network size. Regarding the latter, we assume that consumers use last period's network size as a prediction for the network size in the current period. Formally, in state $\omega = (s, n, e)$ a consumer of type θ_S , θ_N , and θ_E maximizes her utility:

$$\begin{aligned}
 u(\theta_S, d) &= \begin{cases} \beta - p(\omega) + \varepsilon_S & \text{if } d = S \text{ (buy in store)} \\ \gamma - c(n) - \tau + \varepsilon_N & \text{if } d = N \text{ (download from network)} \\ 0 + \varepsilon_E & \text{if } d = E \text{ (no acquisition at all)} \end{cases} \\
 u(\theta_N, d) &= \begin{cases} \beta - p(\omega) - \tau + \varepsilon_S & \text{if } d = S \text{ (buy in store)} \\ \gamma - c(n) + \varepsilon_N & \text{if } d = N \text{ (download from network)} \\ 0 + \varepsilon_E & \text{if } d = E \text{ (no acquisition at all)} \end{cases} \\
 u(\theta_E, d) &= \begin{cases} \beta - p(\omega) - \tau + \varepsilon_S & \text{if } d = S \text{ (buy in store)} \\ \gamma - c(n) - \tau + \varepsilon_N & \text{if } d = N \text{ (download from network)} \\ 0 + \varepsilon_E & \text{if } d = E \text{ (no acquisition at all)}. \end{cases}
 \end{aligned}$$

Here, $\beta > 0$ and $\gamma > 0$ represent the basic utility of the physical and the digital form respectively. The costs of downloading when the resulting network is of size $n \in [0, 1]$ are represented by $c(n)$. A natural shape of this cost function is for it to be decreasing in the network size at a diminishing rate. This is due to the fact that the more users are sharing this file, the easier it is to acquire it from the P2P network.² In our model we implement the following convex decreasing cost function:

$$c(n) = \frac{\sigma}{\rho + n}, \quad n \in [0, 1].$$

Here $\sigma > 0$ represents the generic cost factor of downloading, incorporating a collection of factors that may affect downloading costs, for instance, population computer literacy, the availability of broadband internet infrastructure, and most importantly, the degree of legal enforcement of intellectual property rights. The parameter $\rho > 0$ influences the curvature of

²Notice that unlike local area networks, contemporary P2P networks are less prone to network congestions.

the cost function. The smaller ρ is, the steeper $c(n)$ is around the region where n is close to 0. Note that ρ and σ are identical for all consumers and are independent of the network size.

Finally, the terms ε_S , ε_N , and ε_E are random terms leading to heterogeneous consumer behavior. We assume that these terms are independently drawn according to a Gumbel extreme value distribution with location parameter 0 and scale parameter $\lambda \geq 0$ for each consumer and each of the three alternative choices every period anew. The scale parameter λ is inversely related to the level of heterogeneity among consumers' preferences.

Given the current state ω and the store's price at this state $p(\omega)$, the probability that a consumer with type θ is of type θ_d next period is then given by:

$$q(\theta_d | \theta)[\omega, p(\omega)] = \frac{\exp(\lambda u(\theta, d))}{\exp(\lambda u(\theta, S)) + \exp(\lambda u(\theta, N)) + \exp(\lambda u(\theta, E))},$$

for $\theta \in \{\theta_S, \theta_N, \theta_E\}$ and $d \in \{S, N, E\}$. Hence, from state $\omega = (s, n, e)$ with price $p(\omega)$, the process resumes in state $\omega' = (s', n', e')$ next period, where

$$\begin{aligned} s' &= s \cdot q(\theta_S | \theta_S) + n \cdot q(\theta_S | \theta_N) + e \cdot q(\theta_S | \theta_E), \\ n' &= s \cdot q(\theta_N | \theta_S) + n \cdot q(\theta_N | \theta_N) + e \cdot q(\theta_N | \theta_E), \\ e' &= s \cdot q(\theta_E | \theta_S) + n \cdot q(\theta_E | \theta_N) + e \cdot q(\theta_E | \theta_E). \end{aligned}$$

We denote this process of state transitions by $Q : \Omega \times \mathbb{R}_+ \rightarrow \Omega$.

The firm maximizes the present value of all future profits by implementing the price policy $p : \Omega \rightarrow \mathbb{R}_+$, that for all $\omega \in \Omega$ maximizes the value of

$$V(\omega, p) = \pi(\omega, p(\omega)) + \delta \cdot V(Q(\omega, p(\omega)), p),$$

the Bellman equation for the firm's profit maximization problem.

The firm's search for the optimal pricing policy constitutes a Markov decision problem. To solve this problem, we turn to numerical methods. More precise, we discretize the state space and apply a policy iteration algorithm to find the optimal pricing strategy. Details on the numerical method are further explained in Section 3. The discretized version of our model is guaranteed to possess a stationary optimal pricing policy. The process of state transitions that is induced by this pricing policy provides insight in the implied market share dynamics. In our simulations presented in later sections, the state transition process leads to a unique invariant distribution. The set of states with positive probabilities in the invariant distribution is called the absorbing set. We refer to it as the long-run steady state or just the steady state in the non-technical parts of the paper.

3 Numerical method

In this section, we explain how the model described in Section 2 is solved numerically. Firstly, we discretize the state space and adapt the firm's problem accordingly. Next, we apply the

policy iteration algorithm³ on the discretized state space. Finally, we apply a bracketing algorithm⁴ to determine improvements in each state of the price policy.

3.1 Discretization

Given a natural number $k \geq 1$, we define the discretized state space by

$$\widehat{\Omega} = \{(\widehat{s}, \widehat{n}, \widehat{e}) \in K^3 \mid \widehat{s} + \widehat{n} + \widehat{e} = 1\},$$

where $K = \{0, 1/k, 2/k, \dots, 1\}$. Figure 1 illustrates graphically (for $k = 4$) the state space and the discretized state space. The large triangle represents the state space. The corner points of this triangle represent the extreme states with full consumer mass at either Store, Network, or Empty. The interior points refer to states where the full mass of consumers is divided over the three platforms, the actual division being proportional to the proximity to the corner points. The triangulation of the triangle represents the discretized state space. The vertices are precisely the states in $\widehat{\Omega}$.




Figure 1: An illustration of the discretized state space and the discretized transition mapping.

Next, we define a transition mapping $\widehat{Q} : \widehat{\Omega} \times \mathbb{R}_+ \rightarrow \Delta(\widehat{\Omega})$ on this discretized state space. Given current state $\widehat{\omega} \in \widehat{\Omega}$ and price policy $p : \widehat{\Omega} \rightarrow \mathbb{R}_+$ the process resumes in state $\omega' = Q(\widehat{\omega}, p(\widehat{\omega}))$. Typically the new state $\omega' = (s', n', e')$ is not an element of $\widehat{\Omega}$. In such a case we allocate probabilities to the nearest states in $\widehat{\Omega}$ with probabilities being proportional to the proximity to these states.

In order to provide a precise formulation of the discretized transition mapping it is convenient to define for a state $\omega' \in \Omega$ the sets:

$$\Delta(\omega') = ch\{(\lceil s' \rceil, \lfloor n' \rfloor, \lfloor e' \rfloor), (\lfloor s' \rfloor, \lceil n' \rceil, \lfloor e' \rfloor), (\lfloor s' \rfloor, \lfloor n' \rfloor, \lceil e' \rceil)\} \cap \Omega$$

and

$$\nabla(\omega') = ch\{(\lfloor s' \rfloor, \lceil n' \rceil, \lceil e' \rceil), (\lceil s' \rceil, \lfloor n' \rfloor, \lfloor e' \rfloor), (\lceil s' \rceil, \lceil n' \rceil, \lfloor e' \rfloor)\} \cap \Omega,$$

³See Howard (1960); Blackwell (1962); see also Judd (1990) p. 416.

⁴Cf. Judd (1990).

where $\lfloor x \rfloor$ ($\lceil x \rceil$) refers to the nearest element in K less (larger) than or equal to x . Depending on the location of ω' , four possible situations can occur. If ω' is a vertex of the triangulated state space, then $\Delta(\omega') = \nabla(\omega') = \omega'$. If ω' is on an edge, then $\Delta(\omega') = \nabla(\omega')$ describes this edge. In case ω' is in the interior of one of the subtriangles, ω' is an element of either $\Delta(\omega')$ or $\nabla(\omega')$. When ω' is located in an upward pointed triangle, $\omega' \in \Delta(\omega')$ and $\nabla(\omega') = \emptyset$. Vice versa, when ω' is located in a downward pointed triangle, $\omega' \in \nabla(\omega')$ and $\Delta(\omega') = \emptyset$.

Contingent on ω' being in $\Delta(\omega')$ or $\nabla(\omega')$, we specify the discretized transition probabilities. If $\omega' = Q(\hat{\omega}, p(\hat{\omega}))$ is an element of $\Delta(\omega')$, then

$$\hat{Q}(\hat{\omega}' | \hat{\omega}, p(\hat{\omega})) = \begin{cases} k \cdot (s' - \lfloor s' \rfloor) & \text{if } \hat{\omega}' = (\lceil s' \rceil, \lfloor n' \rfloor, \lfloor e' \rfloor) \\ k \cdot (n' - \lfloor n' \rfloor) & \text{if } \hat{\omega}' = (\lfloor s' \rfloor, \lceil n' \rceil, \lfloor e' \rfloor) \\ k \cdot (e' - \lfloor e' \rfloor) & \text{if } \hat{\omega}' = (\lfloor s' \rfloor, \lfloor n' \rfloor, \lceil e' \rceil) \\ 0 & \text{otherwise.} \end{cases}$$

If $\omega' = Q(\hat{\omega}, p(\hat{\omega}))$ is an element of $\nabla(\omega')$, then

$$\hat{Q}(\hat{\omega}' | \hat{\omega}, p(\hat{\omega})) = \begin{cases} k \cdot (s' - \lceil s' \rceil) & \text{if } \hat{\omega}' = (\lfloor s' \rfloor, \lceil n' \rceil, \lceil e' \rceil) \\ k \cdot (n' - \lceil n' \rceil) & \text{if } \hat{\omega}' = (\lceil s' \rceil, \lfloor n' \rfloor, \lceil e' \rceil) \\ k \cdot (e' - \lceil e' \rceil) & \text{if } \hat{\omega}' = (\lceil s' \rceil, \lceil n' \rceil, \lfloor e' \rfloor) \\ 0 & \text{otherwise.} \end{cases}$$

For the first of the two contingencies, the situation is graphically illustrated in Figure 1.

3.2 Policy iteration algorithm

We derive the optimal stationary price policy for the discretized state space, $p : \hat{\Omega} \rightarrow \mathbb{R}_+$. We use the policy iteration algorithm to numerically approximate the optimal stationary price policy for the store.

Policy Iteration Algorithm

initialize Choose stopping criterion ε and an arbitrary starting policy $p_0 : \hat{\Omega} \rightarrow \mathbb{R}_+$.

Set $i = 1$.

loop 1. Compute the present value from p_{i-1} for each possible starting state in $\hat{\Omega}$.

That is, solve the system of linear equations

$$V(\hat{\omega}, p_{i-1}) = \pi(\hat{\omega}, p_{i-1}(\hat{\omega})) + \delta \cdot \sum_{\hat{\omega}' \in \hat{\Omega}} \hat{Q}(\hat{\omega}' | \hat{\omega}, p_{i-1}(\hat{\omega})) \cdot V(\hat{\omega}', p_{i-1}).$$

2. Improve the policy for each state in $\hat{\Omega}$.

That is, solve for each state $\hat{\omega}$ the problem

$$p_i(\hat{\omega}) := \operatorname{argmax}_z \pi(\hat{\omega}, z) + \delta \cdot \sum_{\hat{\omega}' \in \hat{\Omega}} \hat{Q}(\hat{\omega}' | \hat{\omega}, z) \cdot V(\hat{\omega}', p_{i-1}).$$

3. Terminate loop if the improvement is negligible:

if $\|p_i - p_{i-1}\|_\infty < \varepsilon$:

then Terminate loop and return p_i as the optimal price policy.

else Increase i with one and resume at step 1.

In step 2 of the loop, given a certain policy, the optimal one-shot deviation with respect to this policy is determined for each possible state. The stationary price policy that is composed from the state prices that induce optimal one-shot improvements guarantees at least the same present value for each possible starting state and, hence, is a better price policy relative to the previous one. This is based on two arguments. First, a one-shot improvement in a state leads to a stationary improvement whenever this state occurs. Second, any improvement implemented in one state implies a weak improvement in all other states. From the one-deviation principle it follows that once no (non-negligible) improvement can be found, an (almost) optimal policy is reached. To sum up, any iteration throughout the running of the policy iteration algorithm guarantees an improvement and, once the algorithm terminates, it returns a (nearly) optimal policy.

In our numerical derivations, we used a bracketing algorithm to numerically solve the maximization problems in step 2 at every iteration.⁵ The procedure underlying this algorithm is analogue to the bisection algorithm for root-solving. Whereas the bisection algorithm starts with two initial points with unequal sign, the bracketing algorithm starts with three initial points (say, x , y , and z with $x < y < z$) with the property that the middle point (y) has the largest function value. Next, at every iteration, the midpoint of the largest of the two intervals that are determined by the three points (i.e., $[x, y]$ and $[y, z]$) is taken. If the value of the midpoint (m) is less than the value of the middle point (y), then the larger interval is bisected (i.e., $(x, y, z) := (m, y, z)$ or $(x, y, z) := (x, y, m)$ depending on whether the first or the second interval is the largest). Otherwise, the smaller interval is dropped (i.e., $(x, y, z) := (x, m, y)$ or $(x, y, z) := (y, m, z)$ depending on whether the first or the second interval is the largest). This procedure is iterated until the search area has shrunk to a size less than a predefined tolerance level.

One challenge towards implementing the bracketing algorithm is to find three starting points satisfying the property that the middle point implies the largest function value. One starting point (x) we fix at the minimum price of zero. As a second point (z) we take any price sufficiently high for the resulting present value to be either less than the one that results from taking a price of zero or to be below the present value of half that price. Such a price is guaranteed to exist. Finally, we are to search for the middle point (y). Starting in the middle of x and z , we half the price until either the present value is above the one that result from taking a price of zero or when the price gets below the predetermined tolerance level for the bracketing algorithm. In the first case, a triple with the desired properties is found and the bracketing algorithm is started; in the second case, a price of zero is optimal in the given state at the given iteration.

⁵See Judd (1990) p. 94.

4 Model calibration and benchmark outcomes

In order for the model to make insightful real-life predictions, we need to calibrate the parameters in such a way that it yields static outcomes consistent with real-life data. The model contains two types of parameters. The first type are those whose values are directly retrievable from consumer surveys and financial statements. The second type are not directly retrievable and are calibrated in order for the model outcomes to fit real-life data on market outcomes. We adopt two benchmark scenarios: the situation before P2P was made available to the public and a few years thereafter. More precisely, we take the years 1999 and 2003 for calibration. The year 2003 is chosen because of the abundance of data and the fact that it leaves room for the model to predict market outcomes for later years for which we have the data to verify. The careful calibration of parameters allows us to analyse, with a reasonable amount of accuracy, the order of magnitude of the societal impacts of the P2P networks on firms' profits and consumers' welfare. Thus, in this model, we speak of numbers in dollar terms.

4.1 Data and evidences

First of all, we start with the CD prices over the years. According to the aggregate data from RIAA (2007), the average price of a CD in the U.S. has been kept at around \$14 over the years. However, the real prices, if one takes inflation into account, have been slightly decreasing over the years. In our model, therefore, we should allow the past prices to be slightly higher than the present prices, as long as the average CD price is around \$14.

Secondly, CD sales. According to the year-end report from RIAA (2009), the record companies shipped to the consumers, in the year 1999, about 938.9 million copies of CDs. That figure has dropped to around 746 million in 2003. This suggests that the CD sales has decreased by about 20% in 4 years. This implies that in our model, the sales in the store in the fifth period should be about 80% of those in the first period.

Finally, the P2P network size. When Napster was introduced to the general public in the year 1999, the usage of P2P became instantly wide spread. For simplicity though, we shall take the network size in the year 1999 to be 0, and assume the network started in the beginning of the year 2000. Now the question is: how big was the network in 2003? In other words, how many albums are being downloaded via P2P networks in 2003 and beyond? Blackburn (2004) obtained P2P downloading volume per album from BigChampagne, a company that measures P2P traffic. In the descriptive analysis of the data, it appears that, the average albums were downloaded about 7 to 10 times more often than they were sold. So the ratio of downloads versus store sales (n/s) should be around 7 to 10.

4.2 Parameter calibration

In the study by Rob and Waldfogel (2006), surveys suggest that the average valuation of purchased CDs was around \$16. We therefore take $\beta = 16$ in our benchmark calibration. In the same survey, the average valuation of a downloaded album was revealed to be around \$11. This will lead us to take $\gamma = 11$. For the scenario of the year 1999, however, we will set $\gamma = -1000$, since P2P networks were not available. The discount factor δ can be taken safely as 0.97 according to the convention of 3% interest rate. In a study done by Peitz and Waelbroeck (2005), the authors show a report from IFPI that explains the break-down of the average cost of a CD. This can be estimated to be about €13 including taxes. However, since in our model we look at all figures in dollar terms, and we know that the CD prices are on average \$14 in the U.S., we can extrapolate that the per unit cost μ in our model should be around \$10.

The generic cost factor of downloading σ and the parameter that captures the curvature of the downloading cost function ρ are model specific and we try to calibrate them in such a way to reflect realistically the downloading costs. As the cost function is $c(n) = \frac{\sigma}{n+\rho}$, we use a calibration with $\rho = 0.025$ and $\sigma = 0.5$. This implies that when no one is using the network, it costs \$20 to download an album, which is more expensive than buying it in the store. However, when 7.5% of the total market downloads an album, the downloading cost drops significantly to \$5, which might be tempting for many consumers, though not all. As around half of the market download, the cost drops further towards \$1, which makes it seriously attractive for everyone. Finally, when the whole population downloads, the cost drops to the minimum of about \$0.5.

The consumer heterogeneity parameter λ is determined by calibration, and is chosen such that both the store and the network achieve realistic amount of quantities in both benchmark scenarios. This leads to a value for λ of 0.20.

The switching cost τ influences the number of periods it takes for any starting state to converge to the long-run steady state. It takes 4 years (periods) to go from the state of 1999 to the state of 2003. We will therefore choose a value for τ that yields outcomes that confirm both the situation of 1999 for $\gamma = -1000$ and that of 2003 for $\gamma = 11$, and the fact that the transition takes 4 periods. An appropriate value for τ appears to be 11. Table 1 summarizes the chosen parameter values with all prices being in dollar terms.

	β	γ	σ	ρ	τ	λ	δ	μ
before P2P	16	-1000	0.5	0.025	11	0.2	0.97	10
after P2P	16	11	0.5	0.025	11	0.2	0.97	10

Table 1: Calibrated parameter values.

5 Model outcomes and analysis

The numerical method presented in Section 3 is programmed into a series of Matlab routines.⁶ We ran the resulting program using the parameter values of Table 1. For the discretization, we opt for a grid size of $1/20$, i.e. $k = 20$; the tolerance level for the price policy to convergence is set to 4 digits after the decimal for the maximum state-wise difference between consecutive iterations.⁷ We obtain outcomes such as price policy, profits, state dynamics, (long-run) probability distribution over states, and welfare.

The triangle of Figure 2 in the appendix displays the optimal price policy in 1999; that is, before the introduction of P2P. The triangulation of the triangle represents the discretized state space for $k = 20$. The corner points of this triangle represent the extreme states with full consumer mass at either Store, Network, or Empty. The interior vertices refer to states where the full mass of consumers is divided over the three platforms, the actual division being proportional to the proximity to the corner points. The numbers in the triangle indicate the prices charged by the firm in the different states (231 in total) in the year 1999 before P2P was available. The firm charges higher prices in the states where it has higher market shares, and vice versa, in a monotonic fashion. This seems natural given the presence of the switching cost. Since this case corresponds to $\gamma = -1000$, the value of s is a sufficient state variable for the firm; accordingly, the price policy is a constant on the diagonals from north-west to south-east which represent a particular value of s .

Regardless of which distribution of states we start the Markov process from, given the optimal price policy and the induced process of state-transition, we always end up in a unique steady state distribution. This is the one depicted in Figure 4, where the numbers denote the probability to be in the respective state in the long-run. In this steady state the firm has an expected market share of 23%, the network size is 0 by construction, and the rest of the market (77%) is unserved. The expected steady-state price charged by the firm is \$14.60.

Figure 3 presents the firm's optimal pricing policy after the introduction of P2P. With the exception of the situation when the network size is 0, the firm still charges higher prices in states where its market share is higher. Moreover, for the same market share of the firm, the firm charges in general a higher price when the network size is larger. This is perhaps due to the fact that it is more difficult and less profitable to attract consumers from the network than from the unserved market, hence the pricing strategy is less aggressive. The prices charged in the presence of P2P are in general lower than before P2P was introduced.

Figures 4–10 show the evolution and convergence of market shares resulting from the introduction of P2P. These market share dynamics are obtained from following the Markov process induced by the optimal price policy of the situation after the introduction of P2P (Figure 3) starting from the steady state distribution of the situation before the introduction

⁶The Matlab routines are available upon request.

⁷The tolerance level of the bracketing algorithm for policy improvement is set to 6 digits after the decimal.

of P2P (Figure 4). Every iteration represents precisely one state transition and, therefore, one year in reality. The figures indicate that the mass of consumers first stirs slightly at the starting state and then gradually shifts towards a new steady state (Figure 10). It takes 20 years (periods) to converge to within 1% of the steady state. However, the market share division of 2008 (Figure 9) is in fact already close to the steady state. The price in the steady state is about \$13.73 and the firm’s market share drops to 6%, while the network size increases to about 75%: more than 10 times the CD sales quantity; only 19% of the market remains unserved.

Table 2 provides some additional insights into the dynamics of the static model outcomes. The columns are arranged in years, where “Std.St.” stands for the steady state. The rows show the model outcomes, where p , π , cs , and w stand for price, profit, consumers’ surplus, and welfare, respectively. Notice that the numbers refer to weighted averages of the state outcomes, where the weight on a particular state is the probability that this state is visited in the respective year. All numbers except for the prices are obtained by multiplying the values of the model outcomes with a factor of 4.1 billion. This scaler is estimated by taking the ratio between the total CD sales in the U.S., according to the RIAA’s year end report, and the CD sales quantity s in our model.

	1999	2000	2001	2002	2003	2008	Std.St.
p (\$)	14.60	14.02	14.17	14.20	14.19	13.88	13.73
s (billions)	0.9377	1.0262	1.0025	0.9049	0.7843	0.3743	0.2312
n (billions)	0	0.0554	0.2517	0.6408	1.0951	2.5728	3.0787
π (\$ billions)	4.3206	3.9844	3.7134	3.2480	2.7679	1.3370	0.8655
cs (\$ billions)	16.6120	18.8030	22.0781	26.6213	31.2834	45.1558	49.7650
w (\$ billions)	20.9330	22.7874	25.7915	29.8693	34.0513	46.4928	50.6305

Table 2: Evolution of market outcomes over years.

The years 1999 and 2003 were used to calibrate some of the model’s parameters; so the match with real-life figures in these two periods is quite accurate. The remaining numbers in the table are outcomes of the model, but nevertheless provide quite a good fit to the real-life figures (as far as they are available). Perhaps surprisingly, the outcomes not only reflect the general decreasing trend of CD prices and sales as in reality very well, but capture the anomalies of the CD sales increase in 2000 and the CD price increase in 2001 and 2002 as well.

Now, let us scrutinize some of the key numbers across the years. The firm’s profit decreases steadily over the years. By the year 2003, this decrease had led to a huge loss of \$1.6 billion in annual profits, according to the model’s estimation. As RIAA (2009) suggests that the CD sales revenue dropped from around \$12.8 billion in the highest year 1999 to around \$11.2 billion in 2003, this profit estimate seems to be in the right range. Consumers’ surplus, however, increases in the same period by some \$14.7 billion per annum. As a result of that,

total welfare benefits by around \$13.1 billion per annum.

So, what does the model predict for the market outcomes for the year 2008 then? This is as far as the data is available, which enables us to verify the accuracy of the predictions. According to the model estimations, the price drops further, CD sales too drop to 374.3 million copies, and the network size continues to grow to around 2.6 billion copies. The official figure from RIAA being 384.7 million copies suggests that, despite various factors in the economy that could influence the CD sales, our model is doing a good job in predicting market outcomes beyond the year 2003. The figures in Table 2 suggest a further drop in annual profit by about \$1.4 billion compared to 2003, and that makes it nearly \$3 billion less than in 1999. Meanwhile, not surprisingly, there is a further increase in total welfare by about \$12.4 billion per annum, comparing to 2003. Finally, in the long-run steady state (around the year 2020), the price will stay around \$13.7, the CD sales will hoover around 231.2 million, the profit will settle down to around \$865.5 million, while the total welfare will be around \$29.7 billion per annum more than in 1999. For more predictions on market outcomes (between 2009-2020), please refer to Table5 in the appendix.

6 Comparative statics

The previous section focussed on the market outcomes across the years 1999 to 2008 and made predictions on market outcomes for the years beyond 2008 in the steady state. This section extends the analysis towards the effects of small changes in the parameters on the steady state market outcomes. Table 4 presents the comparative statics at the steady state. The numbers in the cells refer to the elasticities of market outcomes (in rows) with respect to the model's parameters (in columns). As a reference point, the comparative statics of the benchmark's starting state in 1999 without P2P are included Table 3. In this way, the impact of P2P on certain comparative statics becomes apparent.

	β	τ	λ	μ
p	0.2632	-0.1665	-0.3829	0.5293
s	2.7418	-0.9403	0.0868	-1.7097
e	-0.8131	0.2789	-0.0258	0.5070
π	3.5843	-1.4586	-1.1266	-2.1913
cs	0.9945	-0.1298	-1.1581	-0.6232
w	1.5291	-0.4041	-1.1528	-0.9469

Table 3: Comparative statics of steady state outcomes without P2P.

The intuition behind the numbers in Table 3 compares well to standard intuitions of a monopoly analysis. Meanwhile, according to Table 4, the model outcomes are generally most sensitive towards changes in the parameters β , γ , and μ . In particular, CD sales quantity and the firm's profit increase substantially in β and decrease in γ and μ . This seems fairly

	β	γ	σ	ρ	τ	λ	μ
p	-0.0815	0.1459	-0.0033	0.0002	-0.0717	-0.3707	0.8517
s	4.6613	-3.2408	0.1815	-0.0096	-0.4208	-1.3290	-3.0482
n	-0.3365	0.8674	-0.0517	0.0026	-0.2115	0.4559	0.2217
e	-0.0531	-2.4314	0.1483	-0.0073	0.9474	-1.3876	0.0285
π	4.4888	-2.8493	0.1742	-0.0089	-0.6790	-2.7032	-2.7093
cs	-0.1191	1.0267	-0.1314	-0.0496	-0.2162	-0.2354	0.0172
w	-0.0403	0.9604	-0.1262	-0.0489	-0.2241	-0.2775	-0.0294

Table 4: Comparative statics of steady state outcomes with P2P.

intuitive, as β and γ represent the valuation of the CD and the downloads, and μ the cost of producing one CD. Less intuitive is the fact that higher per unit cost μ increases consumers' surplus. This is due to the fact that a higher production cost induces a higher price, which encourages consumers to download and consequently implies a gain in consumers' surplus because of positive downloading externalities. Total welfare, however, does suffer from high costs, which implies that the gain in consumers' surplus is not sufficient to compensate for the loss in the firm's profit. As one would expect, this effect does not exist in the case before P2P was introduced (see Table 3).

The effect of β is exactly the opposite as the effect of γ on all outcomes except for the unserved market e : both a higher β and a higher γ make the unserved market smaller. The intuition for β and γ generally having an opposite effect is quite obvious, with the exception of the negative effect of the valuation of CD on consumers' surplus and total welfare. This effect can be explained by the fact that a higher valuation of CD induces consumers to switch from the P2P network to the store, thereby making it more costly to download for those downloaders who stay in the network (network effect). Again, without the presence of P2P, β naturally has a positive effect on the consumers' surplus. A more careful look at the effects of β and γ on the CD price reveals even more trickeries. A higher valuation of CD decreases the CD price while a higher valuation of downloads increases it. As we will explain in more detail little later on, this is due to the price elasticity of demand of CDs given the convex functional form of the cost of downloading.

Although the model outcomes are relatively inelastic with respect to the switching cost τ , with the exception of the size of the unserved market e , a higher τ induces a decrease in all of the predicted market outcomes. In particular, switching costs have a negative effect on the prices indicating that, for the calibrated values, our model predicts the 'investment effect' to dominate the 'harvesting effect'.⁸ A higher switching cost deters consumers from consuming any product and forces the firm to charge a lower price, resulting in less consumers' surplus and profit. This holds true for the situation with P2P as well as the situation without P2P.

Consumers' taste heterogeneity is captured by the parameter λ : a larger λ implies less

⁸See Cabral (forthcoming).

heterogeneity. Both tables reveal that a decrease in consumers' heterogeneity λ leads to a lower price and a smaller size of the unserved market e . In the situation without P2P, despite the fact that these extra consumers go to the store, the profit suffers. This is caused by the large decrease in price that is needed to accommodate the lower heterogeneity. In the situation with P2P, despite the lower price charged by the firm, the extra consumers go to the network, together with some consumers that switch from the store. Both negative effects on prices and sales lead to a decreased profit. Apparently, more heterogeneity complicates the formation of the P2P network and hence is of benefit to the firm. Welfare is in both situations (with and without P2P) negatively affected by a decrease in heterogeneity, though the effect of λ on consumers' surplus should be interpreted with caution, since λ determines the sizes of the shocks in utility of all consumers, and as such directly affects consumers' surplus.

The generic cost factor of downloading σ has different properties than the switching cost τ . In fact, it has the exact opposite effects as γ , the valuation of downloads. Higher σ implies higher downloading cost for everyone, making P2P less attractive to consumers, similar to the effects of lower γ . However, the interpretation of σ and γ are quite different, which makes the implications different as well. The quality parameter γ measures the value of a download to a consumer regardless of how the download is acquired. Thus, γ only reflects the intrinsic quality of the download copy, e.g. in the case of music, the sound quality. It does not reflect the fact that the copy has to be illegally downloaded and that it costs time and effort. That is what the cost factor σ captures. Enforcement of intellectual property rights relates to an increased value of σ (consumers face a high risk of punishment when using P2P). Table 4 shows that a higher σ improves the profit of the firm and decreases consumers' surplus and welfare. The total welfare suffers due to the increased size of the unserved market e and the sheer downloading costs that are incurred for those who do download. In other words, a social planner who aims at maximizing welfare should be careful in enforcing intellectual property rights in the context of P2P file sharing, considering the welfare enhancing effect of P2P shown in this crude yet vivid demonstration.

Furthermore, an increase in σ does not lead to an increase in the price of CDs as intuition would suggest, since high σ harms the attractiveness of the P2P network, with which the firm competes in price. In fact, if anything, the price decreases slightly in σ . This is attributable to the price elasticity of demand for CDs, just as for the effects of β and γ on the price. To put it simply, when σ increases, the P2P network becomes less attractive and there are a lot more potential consumers out there for the firm to win over, and hence it charges a slightly lower price to attract these switching consumers. Conversely, when σ decreases, P2P becomes more attractive and a big chunk of consumers will go to the network, making the remaining consumers in the store relatively inelastic to prices, and hence the higher price. Ultimately, this pricing scheme is caused by the convex form of the downloading cost function with respect to the network size, $c(n) = \frac{\sigma}{\rho+n}$. This functional form induces strong network externalities when the network is relatively small and weak network externalities when the

network is large.

7 Concluding remarks

This paper uses a dynamic model to solve the optimal dynamic and state-dependent pricing strategy for a firm that releases music CDs and sells them in the store to compete with the P2P file-sharing networks on the internet.

An interesting point of this model is that we can observe the dynamics of network formation. We can also observe exactly how the firm sets its price conditional on its market share and the network size. Sometimes it sets the price very low in order to win vital market shares to fight against the P2P network, while other times accommodates the network by setting a high price to reap the profit from its own customers.

Based on data from 1999 and 2003, the model predicts the sales quantity of the year 2008 rather accurately to be around 374.3 million copies, while the real figure is 384.7 million. It also predicts that the CD sales will have dropped to around 231.2 million copies per year around the year 2020, when according to the model estimations the long-run steady state is approximately reached. More importantly, our model is able to estimate the welfare gain over the years brought by the P2P file-sharing networks. In the year 2003, four years after P2P was introduced, while the music industry was suffering from a \$1.6 billion forgone profit per year, the total welfare was up by \$13.1 billion per year. By the year 2008, the industry profit has dropped by a further \$1.3 billion, and welfare soared by a further \$12.4 billion per year, making it \$25.6 billion more than in 1999. In the long-run steady state (around year 2020), the industry profit will have dropped to a mere \$865.5 million per year, and total welfare are set to have improved since 1999 by a whopping \$29.7 billion per year.

In the comparative static analysis of the model, one of our most important findings is that total welfare is negatively related to the generic cost factor of downloading. This result coincides with the findings by Rob and Waldfogel (2006) and Herings et al. (2008) in the sense that the existence of P2P actually enhances total welfare. This implies that by making file-sharing more difficult for the consumers, the government is effectively curbing the society from enjoying a high welfare level that P2P technology would generate.

Other findings include an interesting pricing scheme of the firm in the sense that the steady state price decreases when its product becomes more attractive (higher valuation of the CD or higher generic cost factor of downloading), and increases in the valuation of the competing product (higher valuation of the downloads and lower generic cost factor of downloading). Like in Herings et al. (2008), this is due to the convex functional form of the cost of downloads, which affects the price elasticity of demand of the CDs.

Finally, it must be clarified that the results of this paper by no means support the infringement of intellectual property rights or piracy in general. Piracy of intellectual properties such as patents or copyrighted materials for the purpose of regenerating illegitimate profit for

the copier undoubtedly hurts the creative incentives of the content provider. In the case of music CDs against P2P file-sharing, the monumental drop in sales revenue of CDs over the last ten years is astonishing. On the other hand, thanks to the colossal amount of pirated music that serves the otherwise unserved consumers, the immense gain on total welfare is just as impossible to ignore. The real question is then where to draw the line.

Although it is beyond the analysis of this paper, it is worth discussing whether indeed a new form of IPR mechanism (especially copyrights law) is needed in this fast-paced information-based society. Perhaps one that is much more relaxed with regard to the fair use of copyrighted materials. After all, the Web2.0 concept (community-based web environments such as Facebook, Youtube, and Myspace) encourages new content creation by the average public on such a massive scale that the traditional media could never have dreamed of. Yet, the existing copyright laws to a large extent stand in the way of such creative activities. Moreover, copyright laws are meant to encourage and protect creative incentives. When are creators hurt so much that their creative incentives are hindered by file-sharing? One should not forget that the content creators and the intellectual property rights owners are very often different entities – the former being the artists and the latter being the publishing firms. Damages incurred to IPR owners do not necessarily imply damages incurred to content creators, and it is the latter’s creative incentives that matter. Is it not a good thing to reach an otherwise impossibly large audience even if that means one has to give away some music for free? In order to answer these questions, one has to consider the development of quantity, quality and variety of the music products over the years and search for signs of improvement or deterioration.

References

- [1] Arie G and P Grieco (2009). Who pays for switching costs. Working paper: Northwestern University.
- [2] Bae S and J Choi (2006). A model of piracy. *Information Economics and Policy* 18: 303-320.
- [3] Blackburn D (2004). On-line piracy and recorded music sales. Working Paper: Harvard University.
- [4] Blackwell D (1962). Discrete dynamic programming. *The Annals of Mathematical Statistics* 33: 719-726.
- [5] Cabral L (forthcoming). Small switching costs lead to lower equilibrium prices. *Journal of Marketing Research*.
- [6] Chen J, U Doraszelski and J Harrington (forthcoming). Avoiding market dominance: Product compatibility in markets with network effects. *Rand Journal of Economics*.

- [7] Doraszelski U and A Pakes (2007). A Framework for Applied Dynamic Analysis in IO. In: M Armstrong and R Porter (eds.). *Handbook of Industrial Organization*, Volume 3, Chapter 30, pp. 1887-1966. Amsterdam: North-Holland.
- [8] Economides N (1996). The economics of networks. *International Journal of Industrial Organization* 14: 673-699.
- [9] Gayer A and O Shy (2003). Internet and peer-to-peer distributions in markets for digital products. *Economics Letters* 81: 197-203.
- [10] Herings P, R Peeters and M Yang (2008) Competition against peer-to-peer networks. METEOR Research Memorandum RM/08/020, Maastricht University.
- [11] Howard R (1960). *Dynamic Programming and Markov Processes*. Cambridge MA: Technology Press-Wiley.
- [12] Jenkins M, P Liu, R Matzkin and D McFadden (2004). The browser war: Econometric analysis of Markov perfect equilibrium in markets with network effects. Working paper, Stanford University, Stanford.
- [13] Johnson W (1985). The economics of copying. *Journal of Political Economy* 93: 158-174.
- [14] Judd K (1998). *Numerical Methods in Economics*. Cambridge MA: MIT Press.
- [15] Katz M and C Shapiro (1985). Network externalities, competition and compatibility. *The American Economic Review* 75: 424-440.
- [16] Katz M and C Shapiro (1986). Technology adoption in the presence of network externalities. *Journal of Political Economy* 94: 822-841.
- [17] Liebowitz S (2004). Will MP3 downloads annihilate the record industry? The evidence so far. *Advances in the Study of Entrepreneurship, Innovation, and Economic Growth* 15: 229-260.
- [18] Liebowitz S (2006). File sharing: Creative destruction or just plain destruction? *Journal of Law and Economics* 49: 1-28.
- [19] Markovich S (2008). Snowball: A dynamic oligopoly model with indirect network externalities. *Journal of Economic Dynamics & Control* 32: 909-938.
- [20] Markovich S and J Moenius (forthcoming). Winning while losing: Competition dynamics in the presence of indirect network effects. *International Journal of Industrial Organization*.

- [21] Oberholzer-Gee F and K Strumpf (2007). The effect of file sharing on record sales: An empirical analysis. *Journal of Political Economy* 115: 1-42.
- [22] Peitz M and P Waelbroeck (2004). The effect of internet piracy on CD sales: Cross-section evidence. *Review of Economic Research on Copyright Issues* 1: 71-79.
- [23] Peitz M and P Waelbroeck (2005). An economist's guide to digital music. *CESifo Economic Studies* 51. 2-3/2005, 359-428.
- [24] Peitz M and P Waelbroeck (2006). Why the music industry may gain from free downloading – The role of sampling. *International Journal of Industrial Organization* 24: 907-913.
- [25] RIAA (2007). The CD: A better value than ever. An RIAA report prepared by the Communications and Strategic Analysis Department Of the Recording Industry Association of America.
- [26] RIAA (2009). 2008 Year-End Shipment Statistics, RIAA, 202-775-0101.
- [27] Rob R and J Waldfogel (2006). Piracy on the high C's: Music downloading, sales displacement, and social welfare in a sample of college students. *Journal of Law and Economics* 49: 29-62.
- [28] Takeyama L (1994). The welfare implications of unauthorized reproduction of intellectual property in the presence of network externalities. *Journal of Industrial Economics* 42: 155-166.
- [29] Zentner A (2006). Measuring the effect of file sharing on music purchases. *Journal of Law and Economics* 49: 63-90.

A Optimal price policies and market share dynamics




Figure 2: Stationary optimal pricing policy before P2P.

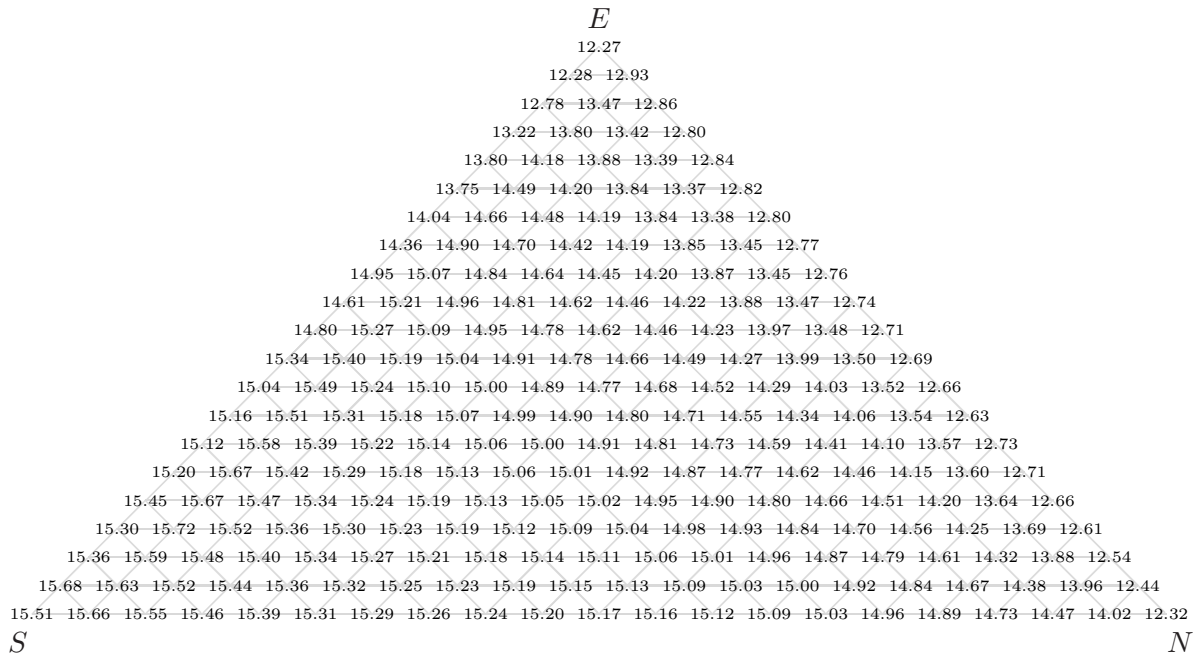


Figure 3: Stationary optimal pricing policy with P2P.




Figure 4: Probability distribution over states in steady state just before the introduction of P2P (1999).




Figure 5: Probability distribution over states one year after the introduction of P2P (2000).




Figure 6: Probability distribution over states two years after the introduction of P2P (2001).




Figure 7: Probability distribution over states three years after the introduction of P2P (2002).




Figure 8: Probability distribution over states four years after the introduction of P2P (2003).




Figure 9: Probability distribution over states nine years after the introduction of P2P (2008).

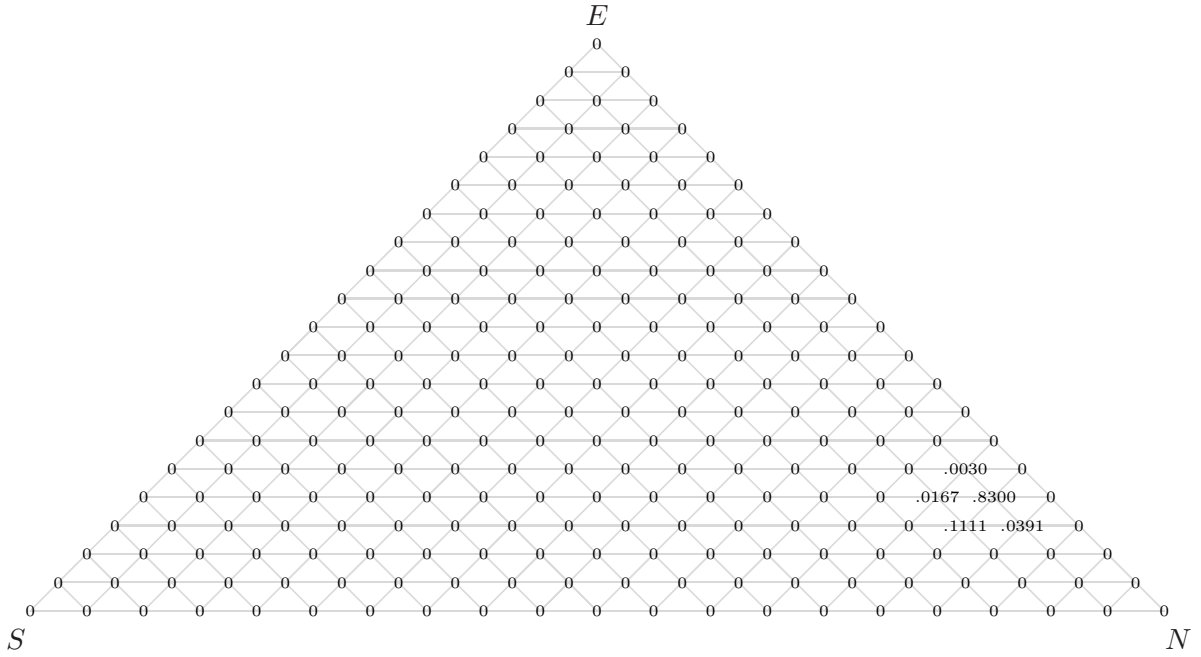


Figure 10: Probability distribution over states in steady state after the introduction of P2P (approximately 2020).

B Future predictions

	2009	2010	2011	2012	2013	2014
p (\$)	13.84	13.81	13.79	13.77	13.76	13.75
s (billions)	0.3374	0.3100	0.2895	0.2743	0.2628	0.2546
n (billions)	2.7031	2.8007	2.8733	2.9270	2.9668	2.9963
	2015	2016	2017	2018	2019	2020
p (\$)	13.75	13.74	13.74	13.73	13.73	13.73
s (billions)	0.2485	0.2440	0.2407	0.2382	0.2362	0.2349
n (billions)	3.0180	3.0340	3.0459	3.0545	3.0611	3.0660

Table 5: Some predictions of market outcomes in the future.