Pricing information goods with piracy and heterogeneous consumers

Waters, James

12 May 2013

Pricing information goods with piracy and heterogeneous consumers

James Waters *†
Nottingham University Business School

Working paper - comments welcome

May 12, 2013

Abstract

We present an information good pricing model with persistently heterogeneous consumers and a rising marginal propensity for them to pirate. Three offsetting pricing mechanisms occur: skimming, compressing price changes, and delaying product launch. We identify a novel trade off in piracy's effect on welfare. We find that piracy quickens sales times and raises welfare in fixed capacity markets, and does the opposite in growing markets. In our model, consumers benefit from piracy except at very high rates in rapidly expanding markets, legal sellers always dislike it, and pirate providers like high but not very high rates. Purchase delay, transient heterogeneity, inelastic demand, and network externalities reduce piracy's effect, but demand uncertainty doesn't.

1 Introduction

Information goods are typically easy to replicate at low cost, and so are susceptible to piracy. Piracy generates no direct revenue for legal sellers, and may be considered harmful to their interests. However, a number of papers (Givon ***et al.***, 1995, 1997; Prasad and Mahajan, 2003; Haruvy ***et al.***, 2004; Liu ***et al.***, 2011) have suggested that piracy can benefit legal sellers. A common theme is that piracy can act as a control on diffusion, either to reach a certain diffusion rate or network size.

These papers include assumptions to stop piracy getting out of hand. In Givon ***et al.*** (1995), Givon ***et al.*** (1997), Haruvy ***et al.*** (2004), and Liu ***et al.*** (2011) consumers are no more likely to pirate in preference to buying when there are many owners or pirates than when there is just one. In Prasad and Mahajan (2003) piracy can be varied over time by

*University of Nottingham Business School. Correspondence: Nottingham University Business School, Jubilee Campus, Nottingham, NG8 1BB, United Kingdom. E-mail: james.waters@nottingham.ac.uk

†Thanks to my PhD supervisors Professor Guiliana Battisti and Professor Paul Fenn. All mistakes are my own.

using a piracy protection parameter; Haruvy *et al.* (2004) allow for a one-time choice of a piracy control parameter. These assumptions are important in explaining the benign nature of piracy in the models' dynamic analyses.

A second frequent feature of these models is the absence or transience of consumer heterogeneity. In Prasad and Mahajan (2003) and Liu *et al.* (2011) the aggregate consumer demand function does not change in response to prior prices and the valuations of previous buyers. Haruvy *et al.* (2004) have identical consumers. An optimal strategy of a seller for extracting value from such consumers cannot be based on the intertemporal price discrimination observed in many markets for information goods (Nair, 2007; Liu, 2010). The dynamic consequences of such skimming is absent from the models.

In this paper, we present a theoretical model in which the marginal propensity to pirate can rise with the number of current users of an information good, while consumer heterogeneity is persistent over time. The model is formulated as a dynamic problem of profit maximisation by the legal seller, who controls market sales through the price. We solve by search the problem for different propensities to pirate and other influential variables.

Two competing impetuses occur when we consider pricing with fixed market size. The first is to extract all surplus from the market by intertemporal price discrimination (skimming) over the product lifetime. The second is to sell before piracy emerges, by reducing prices early and heavily. As the piracy propensity or rate rises, the mean sales times falls as the second impetus becomes more dominant and companies sell the good quicker. Total welfare rises with the piracy rate, with consumers preferring a higher piracy rate than pirate providers, while the pirate providers in turn prefer a much higher piracy rate than legal sellers who want no piracy.

When we consider pricing in a market with rising capacity, the competing benefits of skimming and piracy-evading sales acceleration remain. These behaviours occur later in the product life than in a fixed capacity market. As the piracy rate rises, sales acceleration again becomes more significant. However, its rise also entails more delay in product launch so that pirates are not active too long before the market reaches its peak size. The delay may be so protracted that the market demand is not fully met by the time the product becomes obsolescent. Mean sales time rises slightly with piracy as the launch date and post-launch sales time move in opposite directions. Welfare declines as piracy rises, with consumers liking moderate to quite high rates, pirates liking quite high rates, and companies preferring no piracy. Very high piracy rates are bad for everyone. We believe that our recognition of contradictory dynamic effects is a novel addition to the literature on welfare trade-offs due to piracy (Johnson, 1985; Novos and Waldman, 1984; Yoon, 2002; Belleflamme, 2002; Cremer and Pestieau, 2009).

We then turn to modifying influences. Purchase delay mitigates the effect of piracy, while rising elasticity increases it. Transient heterogeneity renders pricing immune to piracy's impact, as expected, while even low network demand externalities reduce the impact. Piracy continues to delay product launch if capacity growth is subject to uncertainty.

Our analysis shares with Khouja and Smith (2007) assumptions of market eroding piracy based on lagged sales and persistent heterogeneity. Like us, they demonstrate that piracy leads to departure from skimming. However, they use a linear demand function rather than our general constant elasticity function, and employ algebraic solutions for pricing, rather than our more flexible numerical solutions. They consider market contraction rather than our market expansion, and do not study the resulting dynamic trade-off we recognise, nor do they examine mean sales times or welfare.

Section 2 presents our model and section 3 describes the numerical analysis method. Section 4 looks at pricing, sales time, and welfare in the presence of piracy when capacity is fixed, while section 5 does the same when capacity is growing. Section 6 considers modifying influences on piracy's effect, and section 7 concludes.

2 Model

2.1 Introduction

In this section, we describe our model of information good pricing in the presence of piracy. Diffusion is divided into acquisitions from legal sellers and pirate providers. The split is decided by competition between the two groups. Potential buyers are heterogeneous in their valuation of the good, so that price acts as a control variable on diffusion. The number of pirate providers rises with the number of previous buyers. Aside from pirate entry, additional dynamics in the model are induced by emergence of market capacity and purchase delays. Profit maximising legal sellers perform dynamic optimisation over pricing, taking into account the dynamics within the model.

2.2 Specification

There are k_l legal producers of an information good. Legal producers are profit maximising and possess the ability to produce the information good developed from their own research and development. The good is assumed to have been developed by the start of the time period under consideration, so the number of legal producers is constant. In later numerical analysis we put $k_l = 1$. Legal producers can instantly produce copies of the good at a constant unit cost of c_l. As the cost of production can be absorbed into net price, we can without loss of generality set $c_l = 0$.

At time t, there are k_t pirate providers of the good. Pirate providers produce copies of the good innovated by a legal seller. They initially get the production technology by acquiring a copy of the good from a legal seller or pirate provider. The technology may be as little as computer software and a DVD burner. Pirate providers are a subset of current good users, so as the number of past acquisitions rises, the number of pirate providers may increase too. We assume that the number of pirate providers is proportional to the number of goods previously sold, so that $k_t = sS_t$ where S_t is all goods sold by time t and s is a coefficient of proportionality. Fractional numbers of pirate providers are allowed, representing providers who are less active than legal sellers. The proportionality may be rationalised as arising from the inherent tendency to pirate of a certain percentage of individuals who have the opportunity to do so. Alternatively, we could reason that the motivation to pirate rises with increased numbers of past sales. The unit cost of pirate production is constant. We absorb it into the net price charged by pirates, so again we can without loss of generality set production cost at zero.

The population who could have acquired the good by time t is denoted $N(t)$. $N(t)$ includes people who have already acquired the good and those who have not. We may term this population as market capacity, as they are people who will acquire the good if they are offered it at a price below their valuation (possibly with a purchase delay), but may not yet have been offered at such a price. $N(t)$ can vary over time, for example with a rise in the number of owners of a technology necessary for the information good's usage, such as DVD players or computers. Once the potential adopter has acquired the good, they will not acquire it for a second time.

We denote the number of new potential buyers at time t by $n(t) = N(t) - N(t-1)$. Potential buyers are heterogeneous in their willingness to pay for the good. We discriminate between their willingness to pay for the legally supplied and the pirate supplied good. For the legally supplied good, the new entrants have an aggregate demand function of $q = n(t)p^{-a}$ where $p \geq 1$ and $a > 1$. Similarly, the initial consumers in the first period have an aggregate demand function of $q = N(1)p^{-a}$. In numerical analysis, we discretise the demand function by dividing it into M equal parts $n(t)/M$, $2n(t)/M$, …, $n(t)$, with corresponding prices p_1, p_2, …, p_M. Zero demand can be obtained by setting prices arbitrarily high. We find price p_m by solving $mn(t)/M = n(t){p_m}^{-a}$ or $p_m = (M/m)^{1/a}$. We define $n_m(t)$ to be the number of new entrants with willingness to pay of p_m, for $m = 1, 2, \ldots, M-1$. Then $n_m(t) = int(n(t)/M)$ where $int(x)$ denotes the integer part of x. We take the lowest valuation when $m = M$ as additionally including all rounding errors, so $n_M(t) = n(t) - \sum_{m=1}^{M-1} n_m(t)$. The willingness to pay of the initial consumers at time $t = 1$ is similarly evenly spread. We define C_m as the number of current potential buyers (both from new entrants and previously entered non-users) with willingness to buy of p_m.

When the unique offer price in the market is from legal sellers and is $p(t)$, the number of individuals valuing the good at more than $p(t)$ is $\sum_{m \in m_p} C_m(t)$ where $m_p = \{m : (M/m)^{1/a} \geq p(t)\}$. The remainder of potential buyers is $C(t) - \sum_{m \in m_p} C_m(t)$. A proportion q of the potential buyers purchase the good in the current period for each valuation band exceeding $p(t)$. The number of buyers in each valuation band describes the distribution of aggregate demand and is given by

$$\begin{aligned} qC_m(t) \quad & \textbf{if}\,(M/m)^{1/a} \geq p(t) \\ 0 \quad & \textbf{if}\,(M/m)^{1/a} < p(t) \end{aligned} \tag{1}$$

and the total number of buyers is $\sum_{m \in m_p} C_m(t)$.

The remainder of potential buyers delay their purchases and remain in the market. Thus, the number of potential buyers in each valuation band who are carried over for another period is $(1-q)C_m(t)$ if $(M/m)^{1/a} \geq p(t)$ and $C_m(t)$ otherwise. There are many plausible explanations for a purchase delay. The good may not presently be required if an older version is in use. A potential buyer may have incomplete information about market prices, or may not process or act on them quickly. They may expect future declines in price, or at least consider it a possibility worth acting on.

Pirate suppliers provide their good at $p_{pirate}(t)$. In using the pirate supplied good, buyers choose to bring the quality up to the quality of a legally supplied good and in doing so must pay a proportion c of the legal price $p(t)$. We may optionally consider the price to be equal to the cost of quality restoration, so the $cp(t)$ is an additional deadweight loss of pirate production. The effective price of acquiring a pirate copy of the good is thus $p_{pirate}(t) + cp(t)$. Pirate suppliers price to remain competitive with legal suppliers, and set $p_{pirate}(t)$ such that $p_{pirate}(t) + cp(t) = p(t)$, or $p_{pirate}(t) = (1-c)p(t)$.

As for the legally supplied good, the pirate supplied good may be bought by a proportion q of potential buyers valuing the good at more than the offer price. The expression for the distribution of aggregate demand is the same as inequalities 1.

Sales at any price are divided evenly between legal sellers and pirate providers. The share of legal sales in total sales is $k_1/(k_1+k_t)$, and the share of pirate sales in total sales is $k_t/(k_1+k_t)$. Hence aggregate legal sales are

$$\frac{k_l}{k_1+k_t} p(t) q \sum_{m \in m_p} C_m(t) \tag{2}$$

where $m_p = \{m : (M/m)^{1/a} \geq p(t)\}$, and aggregate pirate sales are

$$\frac{k_t}{k_1+k_t} p(t) q \sum_{m \in m_p} C_m(t) \tag{3}$$

Legal sellers selling at the same price share sales equally, so that an individual legal seller's sales are

$$\frac{1}{k_1 + k_t} p(t) q \sum_{m \in m_p} C_m(t) \tag{4}$$

We assume that identical prices apply. The pricing decisions of a single legal seller are followed by all other legal and pirate sellers; if the leader is profit maximising, then the followers also maximise their own profits by symmetry, conditional on other sellers following their lead.

The leader legal seller acts to maximise discounted future profits by setting a price sequence $\{p(t)\}$ for $t = 1, \ldots, T$ for some upper time limit T. They take into account the exogenous dynamics in market capacity and purchase delay, and the endogenous dynamics in demand preference and pirate emergence. The legal seller is assumed to have perfect knowledge of all parameters and future dynamics. They face a dynamic programming problem whose objective function to be maximised follows from collecting the above expressions:

$$\sum_{t=1}^{T} \frac{d^t}{k_1 + k_t} p(t) q \sum_{m \in m_{p(t)}} C_m(t) \tag{5}$$

where d is the legal sellers' discount per period, $m_{p(t)} = \{m : (M/m)^{1/a} \geq p(t)\}$. The dynamics of the number of potential buyers in each valuation band ($C_m(t)$) are given by $C_m(t+1) = (1-q)C_m(t) + n_m(t)$ if $(M/m)^{1/a} \geq p(t)$ and $C_m(t+1) = C_m(t) + n_m(t)$ otherwise, where $n_m(t) = int(n(t)/(M+1))$ for $m = 0, 1, \ldots, M-1$ and $n_M(t) = n(t) - \sum_{m=0}^{M-1} n_m(t)$ for $m = M$. The dynamic in the number of pirates is described by $k_t = s \sum_{\tau=1}^{t-1} \sum_{m \in m_{p(\tau)}} C_m(\tau)$. Starting values are $C_m(0) = 0$ for all m and $k_{p,0} = 0$.

We note that $k_t = s \sum_{\tau=1}^{t-1} \sum_{m \in m_{p(\tau)}} C_m(\tau)$. The expression 5 for maximisation may be considered a weighted sum of the per period sales, where the weights are themselves endogenously defined as the legal sellers' current shares. The per period sales are constrained by the emergence of capacity. This interpretation helps to clarify the subsequent role of piracy, which reduces the weights over time and whose effect depends on the possible set of sales.

We conclude our model presentation by noting its relation to models of information diffusion by central broadcasting (Geroski, 2000). Under central broadcasting, a single source transmits information about the product to a fixed proportion of non-owners, who can then buy it. The relation between past and new acquisitions in central broadcasting is similar to that in our model through the assumption of delayed purchase. In any time period, the construction of links over a proportion of non-owners in the central broadcasting corresponds to partial use of links for sales in our model. In both models, the partial link employment controls the rate of sales and thereby influences diffusion pace. We could, at the cost of sacrificing our economic explanation of delays, use an assumption on linking consistent with broadcasting from all present users rather than a central source (as in the Givon ***et al.*** (1995) model which has both central and current user broadcasting). The alteration would not change the validity of the other model assumptions on reduction of legal shares by piracy and on dynamic programming of the objective function.

3 Numerical analysis

We seek to solve equation 5 for the decision sequence of the legal sellers. For a particular parameter set, the solution will give us a price sequence. We can then calculate sales patterns and welfare, and perform comparative analyses and alter assumptions. Prior work on diffusion in the presence of piracy has used various methods of solution. In Givon ***et al.*** (1995) and Givon ***et al.*** (1997) a full deterministic diffusion pattern can be generated by iteration on their difference equations. Liu ***et al.*** (2011) use exhaustive search to find optimal prices in their iterative model. Prasad and Mahajan (2003) form the Hamiltonian for a legal seller's problem, and proceed to partial analytical solution expressing recursive forms for price and piracy protection. Numerical solution and recursion is used to find value maximising sequences for various parameter values. Khouja and Smith (2007) give an explicit algebraic solution for price. None of these methods is perfect for our problem with heterogeneous consumers. Nair (2007) and Liu (2010) face dynamic pricing problems similar to ours with such consumers. However, they either specify infinite timescale problems and can use Bellman iteration (Nair, 2007) or work with what turns our to be a relatively low dimensional problem (Liu, 2010). Again, their methods are not entirely suitable for our problem.

3.1 Our estimation method

We investigate the theoretical properties of the model by solving the legal seller's problem through exhaustive search of pricing sequences with tracking of demand structure. We calculate the properties over a five year period with annual price setting and monthly emergence of capacity. These time parameters are set for practical and empirical reasons. The annual pricing is intended to reflect the persistence of prices due to contractual agreements, menu costs, and the increase in option values of delaying purchase (so dampening the effectiveness of a price change). When we have a five year horizon, the number of numerical calculations required by our solution algorithm is relatively limited and the solution is quick.

We take five years as the lifetime of our information good, following Liu ***et al.*** (2011) on software products. There are three parameters that vary in the model for the purposes of comparative theoretical analysis. The first is the piracy rate s expressing the sales captured by current pirates as a proportion of past sales. The second parameter is the proportion q of potential buyers who buy in the current month when the market has an offer price below their valuation price for the good. The third is the demand elasticity a. We can also vary the changes in market capacity.

For a given parameter set, the legal sellers' problem is solved to give a sequence of prices as the control variable and a sequence of sales as the response variable. We solve the problem by a dynamic programming algorithm over the five year period. The algorithm tracks the price sequences that lead to the highest discounted sales value. At each time period, the possible non-user distributions at the start of the period are found by taking the possible distributions at the end of the previous period and adding capacity newly emerging in the period. The capacity increase is spread evenly over all five valuation points, and increases at the specified speed. For each distribution at the start of the period, new distributions are generated by having the legal sellers offer to sell at each possible price. The possible prices at the start of each year are the valuation prices plus a higher price corresponding to no sales, and for all other months are constrained to be the previous month's price. At each price and each starting distribution, the specified proportion of non-users at each valuation level exceeding the price buys the good and leaves the non-user distribution. Each resulting distribution has an associated price sequence comprising the price sequence associated with the preceding distribution together with the present price, and a sales sequence generated similarly. We iterate over all periods to obtain final price and sales sequences after five years. The sequence of numbers of pirate sellers is calculated as the specified proportion s of the one period lagged cumulative sales sequence. We retain the pair with the highest discounted sales value, giving the optimal price choices of the legal sellers. Discounted sales values to the legal sellers are calculated as the discounted legal sales values, where legal sales are the legal share of total sales divided between legal and pirate sellers. A numerical device for obtaining the same end result with less memory requirements is used, whereby at the end of each period we only retain the sequence pairs with the highest discounted values for each non-user distribution, as any later sequences with these subsequences within them will have a higher discounted value than the same later sequences with the subsequences replaced by lower value subsequences with the same terminating distribution.

In all solutions we hold the monthly discount factor d constant at a rate of $1/1.01$, equivalent to annual discounting of 12.7 percent per year. The current period share of potentially surplus increasing purchases, q, is held at one until we examine its comparative effect. The capacity growth is specified to be zero or a positive constant, while the elasticity a is fixed at one until we examine its comparative effect. We use five valuation points on the discretised demand distribution ($M = 5$).

The estimation was implemented in the R programming language (R Development Core

Team, 2009). The code is available from the author's website [1].

4 Pricing with fixed capacity

4.1 Pricing

Figure 1: Price variation as a function of the piracy rate with fixed capacity. The rates are 0 (top left), 0.000025 (top right), 0.0001 (middle left), 0.0002 (middle right), and 0.0003 (bottom).

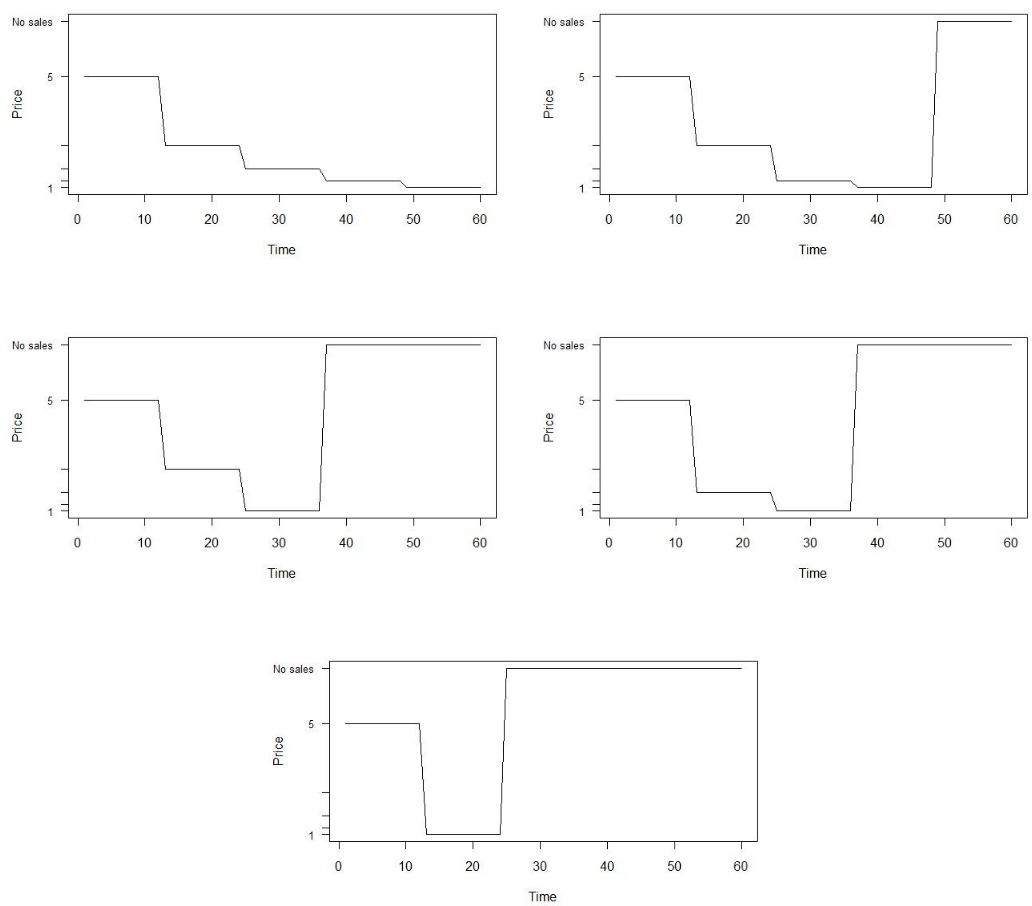

Figure 1 shows how dynamic prices change as the piracy rate rises, when market capacity is held constant at 100,000. In the top left graph, there is no piracy. Prices decline four

[1] http://ebasic.easily.co.uk/02E044/05304E/pricing_info_goods.html

times, between every period, as the legal seller sells to each valuation band in turn. Demand is fully met in the final period as prices reach the lowest valuation band. The top right graph shows pricing at a slightly higher rate of piracy, so that when all demand has been met in the market, piracy's effect is equivalent to sharing sales with $100000 \times 0.000025 = 2.5$ other identical legal sellers. There are three price declines, with the lowest valuation band and market exhaustion reached one year earlier. The initial price is the highest valuation band. In the final period, the company faces no demand and exits the market. Piracy is increased further in the middle left graph, reaching the equivalent of sharing with ten legal sellers. There are just two price declines at the end of the first two years, falling from the highest valuation band to the lowest in year three. The legal seller then exits the exhausted market. In the middle right graph, piracy's effect reaches the equivalent of dividing revenue with twenty other legal sellers when all market demand is met. The pricing strategy is again a two step reduction in prices from the highest to lowest valuation band. The intermediate price choice is lower, however, so that more of the market demand is captured by the end of the second period. Piracy reaches the equivalent of thirty legal sellers in the bottom graph. Prices fall from the highest to the lowest valuation band after the first period, followed by market exit from the third period onwards. Considering all the graphs together, we see that pricing always starts at the highest valuation band, it finishes in the lowest band with all demand satisfied, and price changes become more compressed near the start of the period as piracy increases.

We can rationalise the observations by reference to the legal seller's profit maximisation. In the absence of piracy, the company can extract all surplus from the market by charging each buyer their highest willingness to pay. The company can do so by practicing intertemporal price discrimination. If discounting is not excessive relative to the gains of intertemporal price discrimination - that is, relative to the gap between valuation bands - then skimming will be practiced, as shown in the first graph.

As the piracy rate increases, pirates will emerge as competitors in proportion to the number of owners. So sales made in months after earlier sales periods will be reduced in value. The value of strategies that spread sales over long periods, including intertemporal price discrimination, will decline relative to the value of strategies that concentrate sales over smaller timescales. Selecting an optimal pricing strategy balances the returns from intertemporal price discrimination against the losses due to piracy. We can see the selection in the graphs in figure 1. Price changes are increasingly concentrated over smaller periods and large price changes occur earlier. The timing of the compressed price changes is determined by the discounting making early sales more valuable to the company than late sales. The bottom graph (with the highest piracy rate) shows pricing consistent with maximising value from sales in the first year followed by capturing whatever residual value remains in the second year.

Our finding of increased price compression in response to piracy also occurs in Khouja and Smith (2007), where they work with a linear demand function and a similar skimming and piracy mechanism to ours. In Liu ***et al.*** (2011), prices reduce to capture more current sales when faced with increased piracy. In Prasad and Mahajan (2003), piracy tolerance endogenously controls the importance of piracy. Higher tolerance is associated with higher prices, with either piracy or low prices able to facilitate accelerated diffusion. The shape of the price sequence is unchanged by changing tolerance, with prices always first increasing then decreasing over time.

4.2 Mean sales time

Figure 2: Mean sales time with fixed capacity: time of launch (black serrated line), mean sales time after launch (red dotted line), and total mean sales time (green dot-dash line) as functions of the piracy rate

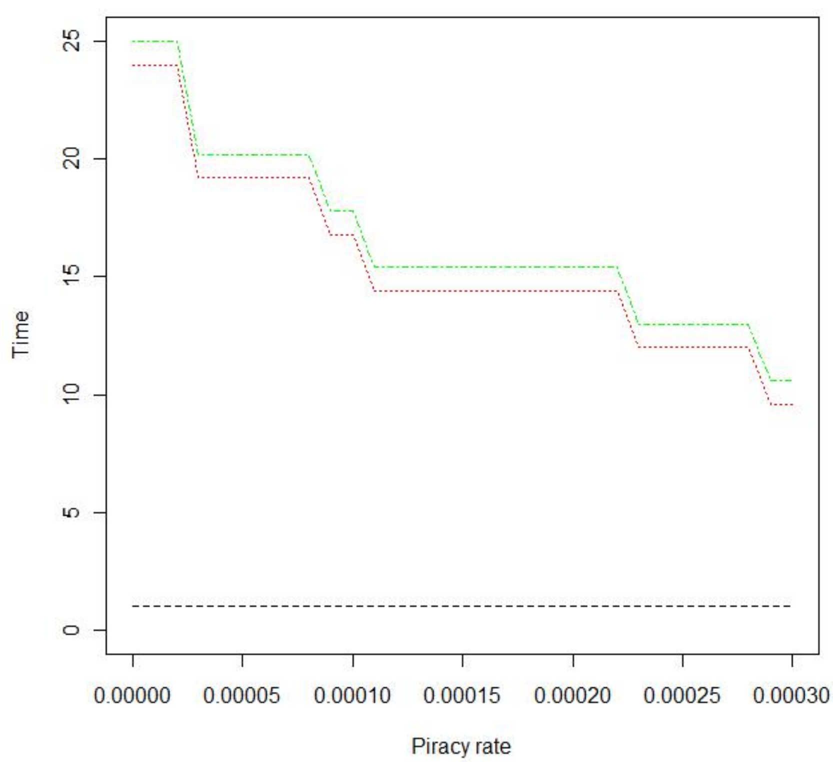

With fixed capacity and positive discounting, piracy leads to earlier and steeper price declines than under intertemporal price discrimination. Such pricing strategies accelerate sales. We can quantify the acceleration in response to piracy by examining the mean time for sales,

$$\sum_{t=1}^{T} tS_t / \sum_{t=1}^{T} S_t \tag{6}$$

where S_t are the total sales at time t. The formula only includes sales made during the product life. When sales satisfy all market demand, the mean times for sales and diffusion coincide. However, as we will shortly see instances in which diffusion is incomplete and the mean time for diffusion is infinite, we consider the mean time for sales instead.

The mean time for sales may be decomposed into time until product launch and mean sales time after launch, in the form

$$\sum_{t=1}^{T} (t-L)S_t / \sum_{t=1}^{T} S_t + L \tag{7}$$

where L is the launch time, that is, the first period of sales.

We calculate the launch time and post-launch mean sales time for piracy rates between 0 and 0.0003. The results are shown in figure 2. The launch time is always the first month, so the total and post launch mean sales times are just shifted by a single month. They reduce quickly as the piracy rate rises, with slightly faster decline at low rates.

The finding that piracy accelerates diffusion is common in the literature. For example, in Liu ***et al.*** (2011), higher piracy reduces the ability of legal sellers to extract profits by high prices because of switching. So prices reduce and sales accelerate. Givon ***et al.*** (1995) suggest that piracy can accelerate legal sales (and by extension, total sales) although this is a result of a relaxation of a parameter constraint making pirate and legal acquisitions perfect substitutes that applies in their main model. Prasad and Mahajan (2003) present a model similar to Givon ***et al.*** (1995), but with control over parameters by prices and anti-piracy measures, and piracy acquisitions strictly supplemental to legal sales rather than contemporaneously replacing them. When piracy is present, it accelerates sales in their monopolistic analysis with plausible parameter values (although it is not clear that sales are accelerated for all parameter values).

4.3 Welfare

We now examine welfare. A small number of researchers have looked at how piracy affects welfare in a dynamic framework. Khouja ***et al.*** (2008) observe that the total number of sales for any level of piracy is close to the total market capacity, so pirate sales compensate for restrictions on acquisitions due to legal prices. Although the authors do not note it, profit maximisation by legal sellers would tend to exhaust market capacity independent of pirate influence, as it is suboptimal to leave many goods unsold. The authors do not calculate discounted welfare although their model allows it, which would be informative about the effects of pirates. In Herings ***et al.*** (2009), the cost of pirate copying declines directly with the number of copiers. Discouraging piracy by increasing its cost reduces welfare. The authors do not present the dynamic patterns of emergence.

We saw in the last subsection that in our model, as sales are accelerated by piracy, the value people derive from them should be discounted less, and so welfare should increase. The changing pricing strategy may also be expected to change the division of welfare between company and buyer. We divide welfare into profits, consumer surplus, and pirate charges. As it is assumed that profits and charges are net of production costs, pirate charges are surplus captured by producers of the pirate good. We could include non-zero and differential production costs in the model, but the analysis would move us away from timing issues in welfare and towards other theoretical mechanisms (Besen, 1986; Johnson, 1985; Belleflamme, 2002; Bae and Choi, 2006).

We have seen that the formula for a company's discounted profits is given by equation 5,

$$\sum_{t=1}^{T} \frac{d^t}{k_1 + k_t} p(t) q \sum_{m \in m_{p(t)}} C_m(t) \tag{8}$$

Thus, the total discounted profits across all companies are

$$\sum_{t=1}^{T} d^t \frac{k_1}{k_1 + k_t} q \sum_{m \in m_{p(t)}} p(t) C_m(t) \tag{9}$$

The total charge of acquisition from pirate sources is

$$\sum_{t=1}^{T} d^t \frac{k_t}{k_1 + k_t} q \sum_{m \in m_{p(t)}} p(t) C_m(t) \tag{10}$$

Then the total cost of adoption from any source is the sum of equations 9 and 10, or

$$\sum_{t=1}^{T} d^t q \sum_{m \in m_{p(t)}} p(t) C_m(t) \tag{11}$$

The total gross value derived from adoption from either legal or pirate sources is the same as the total welfare. Sales occur for buyers with willingness to buy of $(M/m)^{1/a}$ if and only if $(M/m)^{1/a} \geq p(t)$. Thus, the formula for discounted total value derived from sales (and hence total welfare) is

$$\sum_{t=1}^{T} d^t q \sum_{m \in m_{p(t)}} (M/m)^{1/a} C_m(t) \tag{12}$$

Total consumer surplus is the difference between total welfare in equation 12 and total acquisition costs in equation 11, or

$$\sum_{t=1}^{T} d^t q \sum_{m \in m_{p(t)}} ((M/m)^{1/a} - p(t)) C_m(t) \tag{13}$$

Figure 3: Welfare with fixed capacity: consumer surplus (blue horizontal stripes), pirate charges (red vertical stripes), and legal seller profits (green slanting stripes) as functions of the piracy rate

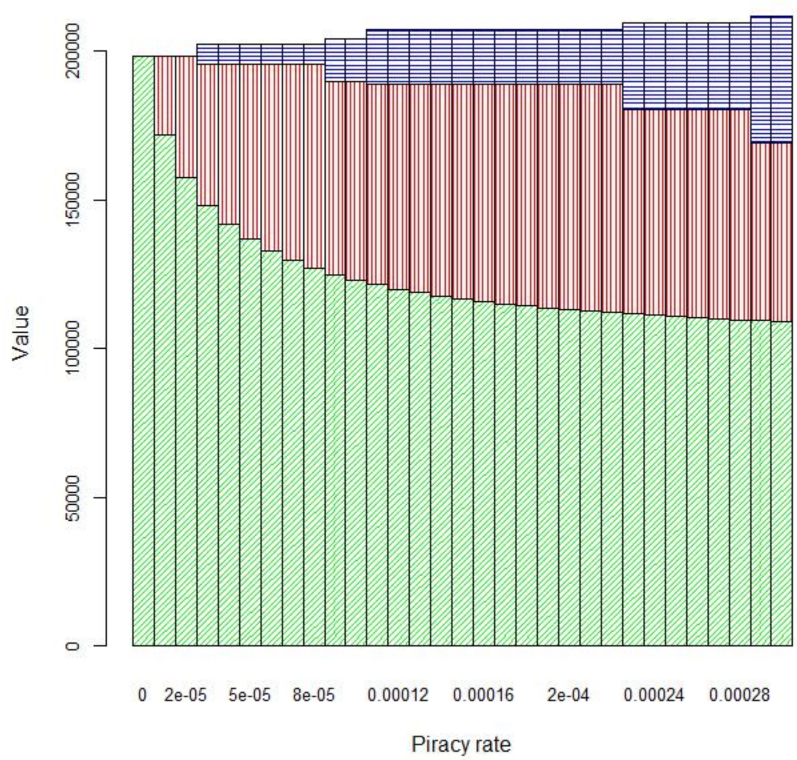

Figure 3 shows welfare and its components as a function of the piracy rate. As expected, total welfare rises as piracy does due to the reduced discounting. The welfare rise is much gentler than the decline in sales time, as most sales acceleration is to buyers with lower valuations. In the mean sales time calculation, only the number of sales matter, whereas in the welfare calculation the valuations are included. The rise in welfare is quite linear.

Profits account for the entire surplus when piracy is zero as perfect skimming is practiced. When the piracy rate increases a little, profits drop sharply as companies begin to share our their profits with small number of pirates over most of their extended sales strategy. As the piracy rate rises, the number of pirating agents who share the income increases linearly, so the legal shares decline inverse proportionally and by smaller amounts. Moreover, price changes are highly compressed in time at the higher rates of piracy, so that increases in piracy act over small periods of time.

As the piracy rate rises from zero, pirate charges rapidly increase their share of total welfare. The rate of increase declines over time, and the share reaches its maximum value at a rate of 0.00022. A general small trend to increase persists thereafter. However, there are large periodic downward corrections in the share of piracy as the rate rises, as these are substantial at higher rates. As a result, pirate charges fall at higher rates.

At low levels of piracy, pricing strategy mainly adjusts the prices later in the product life because earlier adjustment reduces the extent of intertemporal price discrimination excessively. So pirate charges are a share of sales at the time when prices are higher, during the early skimming, and so are quite a large part of total welfare. As piracy increases, it becomes optimal for legal sellers to reduce prices earlier on, and so pirate providers capture lower shares of total welfare. At the highest levels of piracy, the compression of price reductions is so heavy that total pirate charges fall as the pricing strategy changes.

Consumer surplus is zero at the lowest rates of piracy. It increases by steps as the rate rises, with the surplus on average following a roughly linear trend. At the higher rates of piracy it accounts for a little less than pirate charges, and around a third of profits.

When the piracy rate is low, perfect intertemporal price discrimination extracts all profits from consumers. The surplus extraction rises as the piracy rate increases and pricing strategy departs from skimming. The departures become more infrequent as the piracy rate rises, but the price reductions are steeper, so we see the general linear drift in welfare over time.

We summarise the preferences for piracy rates by the market participants when market capacity is fixed. Legal sellers prefer no piracy as they can extract the entire market surplus. Pirate providers prefer a moderately high rate where they are capturing as much of the market as possible without triggering legal sellers to reduce prices too steeply and early and leave the pirate providers with few and low valuing buyers. Consumers like piracy as high as possible because of the price reductions entailed by piracy prevention.

5 Pricing with capacity growth

5.1 Pricing

Figure 4: Price variation as a function of piracy in the presence of capacity drift. The rates are 0 (top left), 0.00006 (top right), 0.00007 (middle left), 0.00009 (middle right), and 0.0006 (bottom).

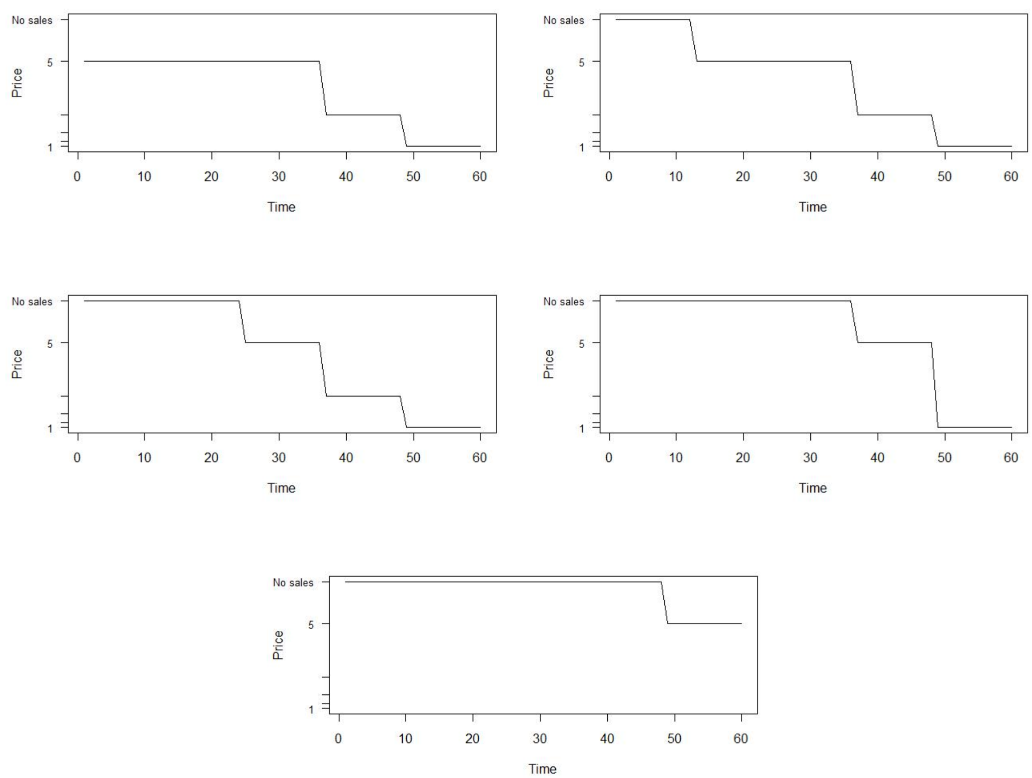

In this section, we examine pricing in the presence of piracy when market capacity starts at zero and increases by 1000 every month. Figure 4 shows the pricing strategy as the piracy rate rises. In the top left graph, piracy is zero, and prices are held at the highest consumer valuation until two years remain. Then prices are reduced twice, firstly to the second highest valuation and then to the lowest. In the top right graph, the piracy rate increases so that there are pirate providers equivalent to $60 \times 1000 \times 0.000006 = 3.6$ legal sellers by the end of the product life. The pricing adjusts in the first year so that the legal seller delays its market entry. In the middle left graph the piracy rate rises slightly to be equivalent to an ending 4.2 legal senders, and the product launch is delayed for another year. The middle right graph has an ending number of pirates equivalent to 6.3 legal sellers. A further launch delay is observed, and the final price drop is steeper from the highest valuation to the lowest valuation. In the bottom graph, piracy ends up equivalent to 36 legal sellers. The legal seller holds launch until the last year, and then keeps prices at the highest valuation.

When there is no piracy, the explanation for the pricing is that the legal seller attempts to extract as much available surplus as possible. In the first few years, they sell only at the top valuation price and so capture all surplus from the top valuers. As the market capacity is growing, the consumers in the top valuation band are replaced and surplus can continue being extracted from them. However, the number of consumers in lower valuation bands also grows and it enhances profits to sell to them after a time. The timing is influenced by the remaining product duration.

For non-zero piracy, the early sales increase the number of pirates who share later revenues. There is an increase in the value of strategies that compress price reductions relative to strategies that stagger them, which we again see in the graphs with the compression becoming acuter as the piracy rate rises. The timing of the compression is due to the capacity trend. Early sales would capture little of the total capacity emergent over the whole product lifetime, but would expose the legal seller to piracy when capacity is much larger. So it is optimal to delay the compression until later in the period, with the delay rising with the piracy rate. At the highest rate of piracy, entry is delayed until the last year and just one year of revenue is earned. The highest price is chosen, but the company's profits are actually unchanged by other prices given unitary demand elasticity and the single period selection.

We can examine the relation between piracy and the growth of capacity by referring again to the legal seller's optimisation problem. A change in the constant growth rate of capacity alters $n_m(t)$ to $an_m(t)$ for some constant a and all t. Since the valuation bands evolve under $C_m(t+1) = (1-q)C_m(t) + n_m(t)$ if $(M/m)^{1/a} \geq p(t)$ and $C_m(t+1) = C_m(t) + n_m(t)$ otherwise, $C_m(t)$ is also scaled up by a factor of a for all t and m. Further, as pirate sales are proportional to past sales, $k_t = s\sum_{\tau=1}^{t-1}\sum_{m\in m_{p(\tau)}} C_m(\tau)$ and so k_t is also increased by a factor of a. Inserting these adjusted functional forms in maximised expression 5, we have

$$\sum_{t=1}^{T} \frac{d^t}{k_1 + as\sum_{\tau=1}^{t-1}\sum_{m\in m_{p(\tau)}} C_m(\tau)} p(t)qa \sum_{m\in m_{p(t)}} C_m(t) \tag{14}$$

We can write $s' = as$ to describe a piracy rate rescaled by a. The a in the denominator can be factored out of the objective function entirely and so does not affect decision making. Thus, the effect of a change in the constant rate of capacity growth on diffusion is the same as a change in the piracy rate, up to a rescaling of the diffusion curve.

Another way of expressing the result is that piracy is a share of past sales, so the effect on piracy of a scaling in sales is equivalent to an increase in the share of pirate sales out of past sales. All income is directly rescaled by the same factor, so doesn't affect the legal seller's decision making. Thus, a rescaled capacity growth is equivalent to a piracy rate change for decision making, and an equivalent rescaling for the overall income. The result is shown graphically in figure 5 [2]. Curve $\boldsymbol{C1}$ traces total capacity at one rate of emergence and curve $\boldsymbol{C2}$ traces capacity emerging at a rate a times higher. The sales path $\boldsymbol{O1}$ is optimal out of many possible paths, with the sum of points on the path giving the sales path's value to the company. Rescaling piracy by $1/a$ and increasing capacity by a maps the set of possible sales paths for $\boldsymbol{C1}$ to the set for $\boldsymbol{C2}$. Thus, the optimal sales path $\boldsymbol{O2}$ is just the optimal sales path $\boldsymbol{O1}$ scaled upwards by a at the revised piracy rate.

Figure 5: Optimal sales paths for capacity emergence curves. Sales path $\boldsymbol{O1}$ corresponds to capacity emergence curve $\boldsymbol{C1}$, and $\boldsymbol{O2}$ $(= a \times \boldsymbol{O1})$ corresponds to capacity emergence curve $\boldsymbol{C2}$ $(= a \times \boldsymbol{C1})$. The piracy rate for the second curve is a times lower than for the first curve.

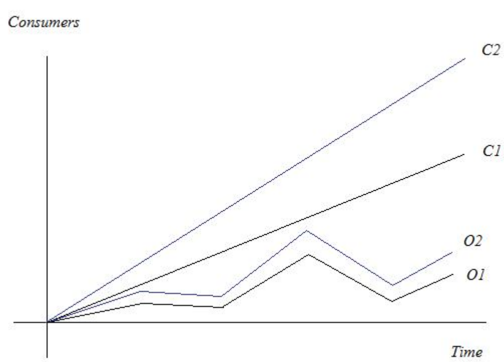

5.2 Mean sales time

Figure 6 shows the mean total sales time as a function of the piracy rate, divided into the time until launch and mean sales time after launch. The total sales time increases a little as the piracy rate increases, undergoing small jumps as the market entry date shifts backwards. The launch date is delayed heavily by rises in the piracy rate. Its movement alternates between large jumps and long periods of no change. The post-launch mean sales time displays the opposite movement. There is initially a long sales time, but the time falls to almost zero at the maximum piracy rate.

[2]Thanks to Paul Fenn for suggesting a graphical interpretation.

Figure 6: Mean sales time with rising capacity: time of launch (black serrated line), mean sales time after launch (red dotted line), and total mean sales time (green dot-dash line) as functions of the piracy rate

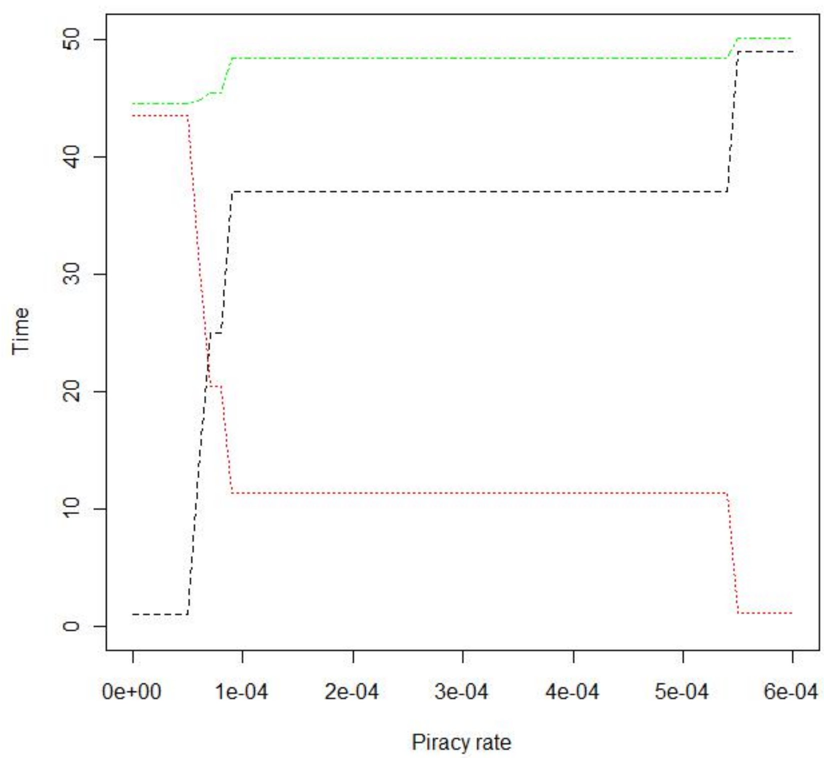

When we examined sales times for the fixed capacity market, the post-launch sales time also declined with rising piracy. However, the launch time remained the same for all piracy rates, in contrast to the increasing capacity case here. The differences in launch time explain the divergent findings on mean total sales time.

5.3 Welfare

Figure 7: Welfare with rising capacity: consumer surplus (blue horizontal stripes), pirate charges (red vertical stripes), and legal seller profits (green slanting stripes) as functions of the piracy rate

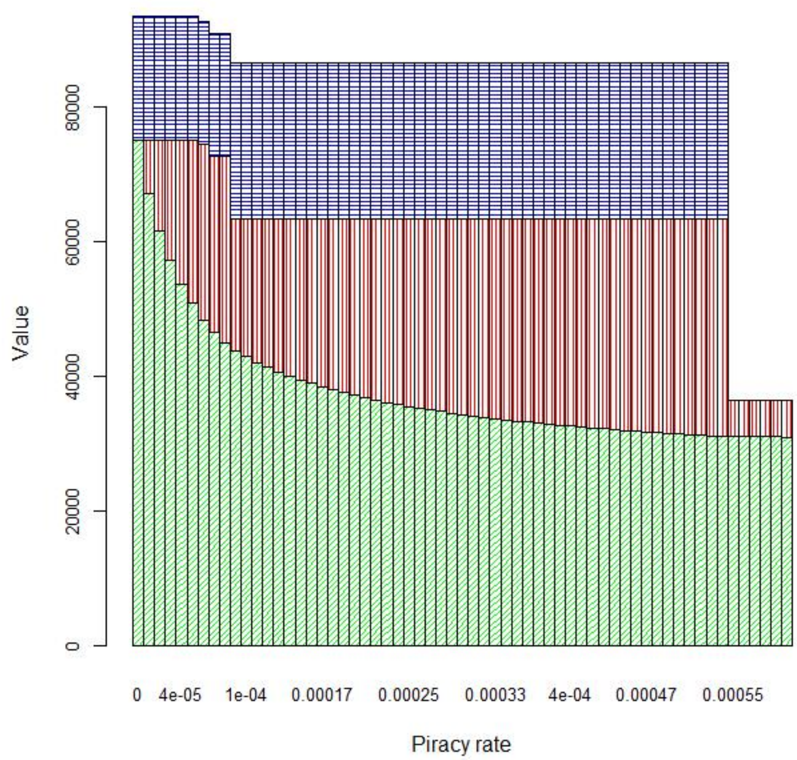

Figure 7 shows total welfare as a function of the piracy rate, decomposed into consumer surplus, pirate charges, and legal seller profits. The total welfare declines over time, with three minor declines in welfare and a single very large one. Welfare is unchanging for most rates. Company profits decline rapidly when the piracy rate first starts to increase and slows down subsequently, as profits are inversely related to the number of pirating agents.

Pirate charges are subject to a general increase that is initially rapid and then slows, as the rate of substitution for legal sales declines. However, pirate charges undergo four downward revisions of increasing magnitude as the piracy rate increases, with the final revision eradicating most of the pirate charges. The first three declines are due to compression, while the final delay is due to compression and truncation of sales at the end of the product life.

Consumer surplus reduces very slightly as the piracy rate rises through very low levels, with the product launch being delayed without the subsequent pricing changing. As the rate rises a little further, there is a large increase in welfare as price changes are compressed and more surplus is transferred from providers to buyers. The piracy rate then increases considerably without any change in surplus, before surplus drops suddenly as the truncation of sales at the end of the product life means that many consumers never buy in the market, even though they could do so at some price and create value for themselves, the company, and pirate providers.

We can again order the preferences for piracy of the market participants based on the welfare they derive. The legal seller prefers no piracy. Pirate providers prefer a high rate, but are worst off at very high rates. Piracy rates that extract the maximum surplus for them are just below the non-zero rates that give them minimal surplus. Consumers are indifferent to any rates that are not too low or too high. They are also worst off with very high rates.

Preferences for small changes in piracy rates depend on the current rate. At low rates, it is in the interest of pirate providers to increase rates slightly but against the interest of consumers. At slightly higher rates, consumers benefit from small increases in piracy rates but pirate providers suffer. Often consumers are indifferent but pirate providers benefit, while around one particular rate they both suffer considerably from a small increase.

Our analysis of welfare indicates the presence of a trade-off due to piracy. Previous work has examined other ambiguities in the effect of piracy on welfare, and the distributional consequences. Besen (1986) and Besen and Kirby (1989) consider circumstances in which different productive efficiencies of legal and pirate copies can result in welfare increases from piracy. Johnson (1985), Novos and Waldman (1984), and Yoon (2002) allow for the possibility that pirate copies may be more inefficient in production but change market access through lower purchase costs. Short term welfare increases from expanded use but reduced long term welfare from legal production disincentivisation is examined in Belleflamme (2002) and Cremer and Pestieau (2009). Takeyama (1994) consider welfare changes when piracy can expand user networks.

Our model identifies a dynamic trade-off. Total welfare increases with the piracy rate when capacity is constant, while it declines with the piracy rate when capacity is rising. The different welfare outcomes indicate the presence of two competing mechanisms that are differentially activated by the choice of capacity trend. On one hand, piracy induces compression of price changes that accelerates sales and increases welfare. On the other hand, piracy increases the product launch delay in the presence of a capacity trend, so decelerating sales. If the piracy rate or capacity growth are very high, product launch may be so delayed that some consumers do not buy the product before it becomes obsolescent.

6 Modifying influences

This section considers factors that modify piracy's effect on pricing. We vary either the parameters previously held constant, or one of the model assumptions. We start by looking at the impact of purchase delay, then demand elasticity, transient heterogeneity, network effects, and uncertainty.

6.1 Purchase delay

Figure 8: Price variation as a function of purchase delay in the presence of piracy. The piracy rate is 0.0001 and current acquisition shares are 1 (top left), 0.8 (top right), and 0.6 (bottom).

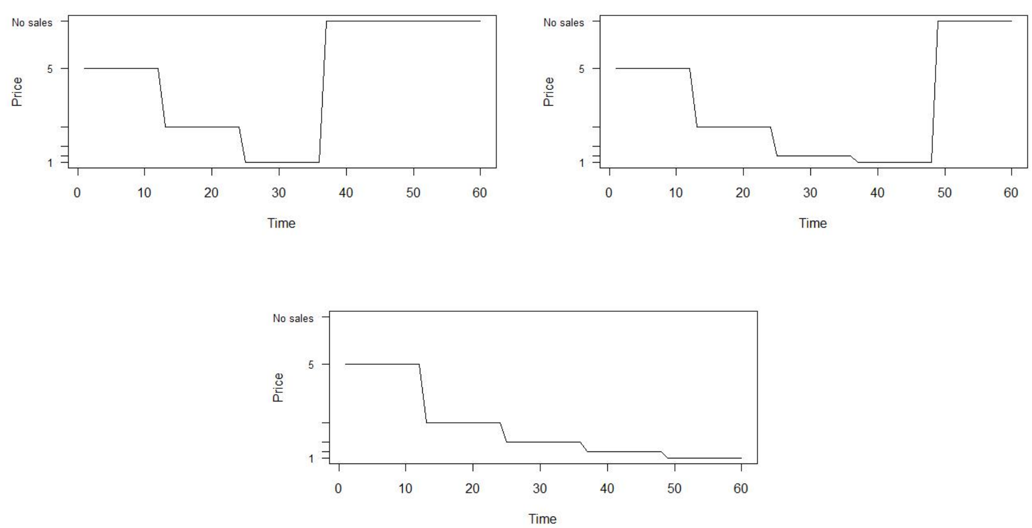

In this subsection, we introduce purchase delay into our model. With the earlier specification, the q parameter is reduced below one. Figure 8 shows q parameters of one (top left graph), 0.8 (top right graph), and 0.6 (bottom graph), with the piracy rate held constant at 0.0001 and capacity fixed. As the purchase delay rises, there is less departure from intertemporal price discrimination.

The reversion to skimming can be explained by considering how the relative values of current and future sales change when purchase delays are introduced. A proportion $1 - q$ of sales will be subject to competition from pirates, whereas it would not be if purchase delays were absent. Thus, the value of current sales relative to future sales is lower for any sales decision in the presence of piracy. As the size of the relative value of current sales motivates the compression of sales, staggered sales are used until a much higher level of piracy in the presence of purchase delay.

6.2 Demand elasticity

Figure 9: Price variation as a function of demand elasticity in the presence of piracy. The elasticities are 1.3 (top left), 1.1 (top right), 1 (middle left), 0.9 (middle right), and 0.7 (bottom).

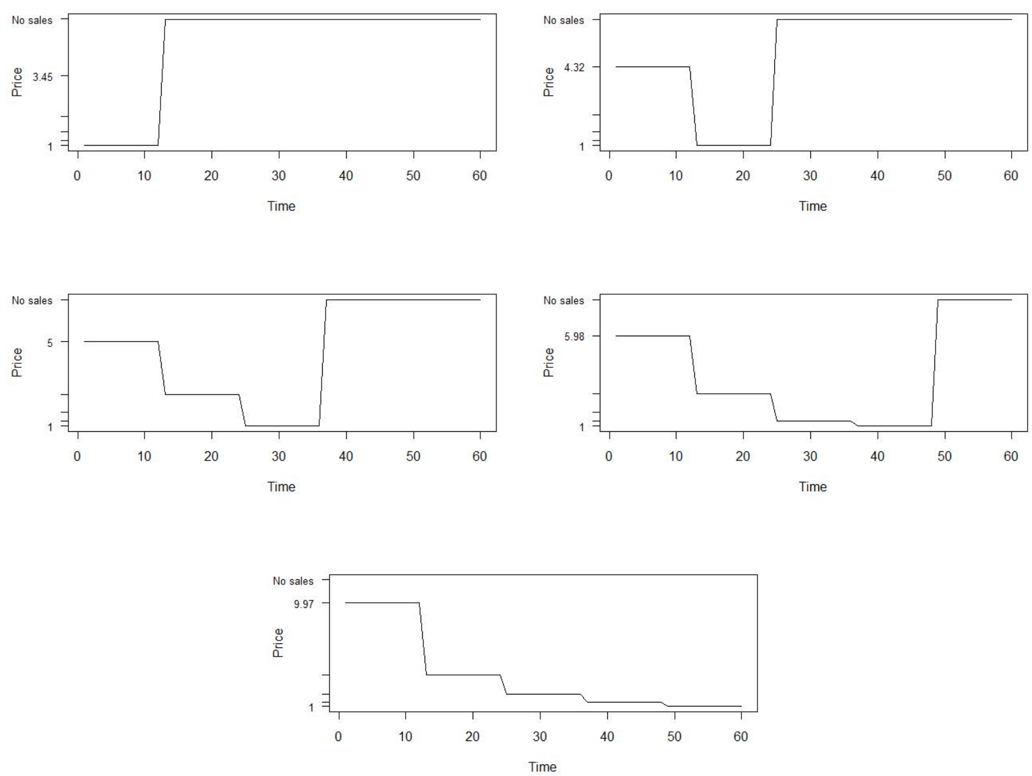

We next consider how piracy's effect on pricing is affected by demand elasticity. In our model specification, elasticity is represented by the parameter a which we set equal to 1.3, 1.1, 1, 0.9, and 0.7 in turn. The piracy rate is set at 0.0001 with fixed capacity. Figure 9 shows the optimal price sequences.

In the top left graph the elasticity is 1.3. The legal seller sets their price at the lowest buyer valuation and sells immediately to the whole market. In the top right graph, the elasticity is 1.1 and the prices start at the highest valuation in year one before dropping to the lowest valuation in year two. The remaining graphs show that as the elasticity drops further, increased staggering of price declines occurs with market exhaustion happening one period later in each graph.

The behaviour can be explained by looking at the prices where demand is concentrated. With elastic demand, demand is concentrated at low prices so the gap between the high and low valuation bands (which have equal density of consumers) is small. So the gains of skimming are limited relative to the cost of piracy, and selling as quickly as possible is most profitable. When demand is inelastic, demand is more concentrated at high prices and the gains of skimming are larger. So skimming is practiced despite piracy.

6.3 Transient heterogeneity

In the introduction, we noted that many papers (Haruvy ***et al.***, 2004; Liu ***et al.***, 2011; Prasad and Mahajan, 2003) in pirate diffusion analysis have assumed no or transient heterogeneity in consumer demand. In this section, we see what the effect is of removing persistence in heterogeneity from our model. To do so, we adjust our model so that after any sales, the remaining non-users are redistributed evenly in the next month over all valuations. We then run the model for various piracy rates and with fixed capacity.

We do not show any graphs, as a single simple behaviour was exhibited. We find that prices remain at the highest valuation whenever there are any non-users and for all rates of piracy. The reason is that in any period with a dispersion over every valuation in the demand function, lower prices do not increase revenue in that period because of unitary elasticity, and they lower the size of the future market (despite the piracy that will reduce the legal seller's access to that market). With transient heterogeneity the demand function is recycled every month, so the argument applies repeatedly and prices remain at the highest valuation.

6.4 Network demand externalities

In this subsection we examine how network demand externalities affect piracy's effect. In both the static and dynamic literature, it is recognised that when there are externalities, piracy can act as a means of reaching a market size offering network gains without having to use low prices to do so. The freedom to increase prices can increase legal seller profits (Takeyama, 1994; Haruvy ***et al.***, 2004).

We assume that the number of current users of the good increases the value of adoption to non-users, where the increase is described by a multiplier linear in the number of current

Figure 10: Price variation as a function of the network parameter in the presence of piracy. The network parameters are 0 (top left), 0.000005 (top right), and 0.00001 (bottom), and the piracy rate is 0.0001.

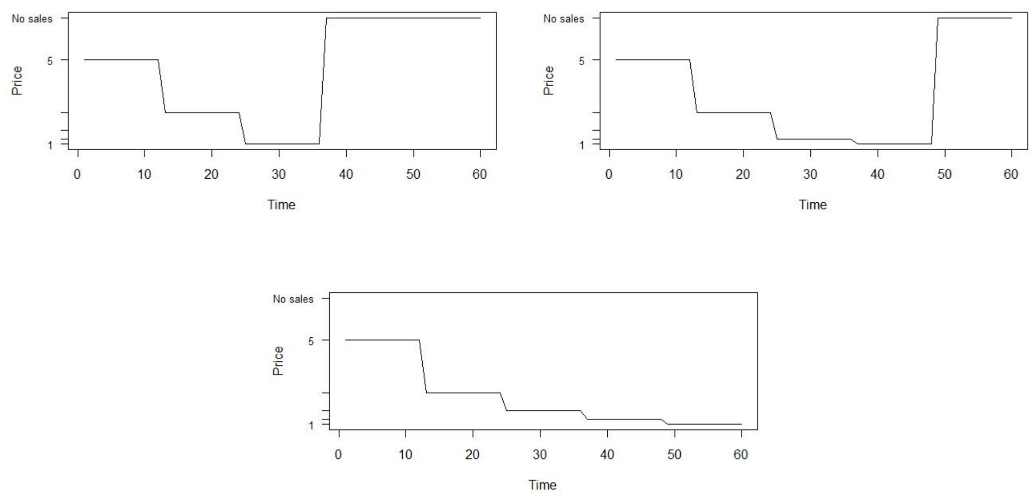

users. Using the notation of section 2.2, their valuations at time t are equal to their initial valuations times a factor of $1+fS_t$, where f is a constant and $S_t = \sum_{\tau=1}^{t-1} \sum_{m \in m_{p(\tau)}} C_m(\tau)$. We continue to discretise the prices, and the initial valuation bands with values of $(M/m)^{1/a}$ for some $m \in 1, \dots, M$ map to valuation bands at time t with values of $(M/m)^{1/a}(1+fS_t)$.

The legal seller continues to set prices at the start of each year. In between active price setting, prices grow to remain in the same relative valuation band as it inflates with the number of lagged sales. Legal sellers face the problem of setting prices $p(t)$ in a revised objective function of

$$\sum_{t=1}^{T} \frac{d^t}{1+k_t} p(t)(1+fS_t)q \sum_{m \in m_{p(t)}} C_m(t) \tag{15}$$

If the piracy rate s is equal to the network parameter f, all sales values are inflated by a network effect and all values are deflated by an exactly equal piracy effect. So sales values in response to pricing decisions expressed in the non-inflated prices are the same as if neither effect was present. Hence pricing decisions expressed in the non-inflated prices are the same as if both were set at zero.

When the piracy rate and network parameter differ, we can solve for the price sequence $p(t)$ using the method described in section 3. For ease of comparability of the pricing structures before and after network effects we show the choices of price before inflation for network effects. Figure 10 shows pricing when the piracy rate is held at 0.0001, capacity is fixed, and the network parameter takes the values of 0, 0.000005, and 0.00001.

The top left graph shows moderate compression of price changes in the absence of any network effects. As the network parameter rises to a twentieth of the piracy rate in the top right graph, skimming lasts a further period. In the bottom graph the network parameter reaches a tenth of the piracy rate, and there is no compression and pricing follows the pattern of intertemporal price discrimination.

The pricing effect of piracy is very sensitive to the presence of network externalities. We can see the reason by considering the joint adjustment factor at time t due to piracy and network externalities in equation 15, which is

$$\frac{1 + fS_t}{1 + sS_t} \tag{16}$$

after noting that $k_t = sS_t$ and that s is the piracy rate. We can solve for an equivalent piracy rate e in the absence of network externalities by equating expression 16 to $1/(1+e)$ to give

$$e = \frac{s - f}{1 + fS_t} \tag{17}$$

For $s = 0.0001$ and $f = 0.00001$, $e \approx s$ for $S_t = 0$ when the product is newly launched, while $e \approx s/2$ for $S_t = 100000$ when the market is exhausted. The equivalent piracy rate acting late in the diffusion is much reduced relative to that in the early diffusion, so the benefit of price change compression relative to skimming is much less. A network parameter only a tenth of the piracy rate can substantially reduce the compression due to piracy.

6.5 Uncertainty

Until now we have assumed that capacity growth is deterministic. In this subsection, we examine the effect of an anticipated shock on pricing. There is initially a capacity of 100000 that is only changed by a shock occurring at the end of year three. The shock increases capacity by 0 or 100000 with equal probability. The legal seller maximises their expected discounted profit at all times.

We solve for the optimal pricing scheme as in section 3 by assuming that the shock takes the low value and calculating the profits emerging from every possible pricing strategy. We retain the pricing sequences for years four and five that give the highest discounted (deterministic) value for each possible sequence in years one to three. The solution and retention are then repeated assuming the shock takes the high value. For each sequence for years one to three, we average the discounted deterministic values for the optimal price

Figure 11: Price variation as a function of piracy in the presence of a capacity shock. The black dashed lines show pricing before and after a high shock, while the red dotted lines show pricing before and after a low shock. The piracy rates are 0 (top left), 0.000025 (top right), 0.0001 (bottom left), and 0.0003 (bottom right).

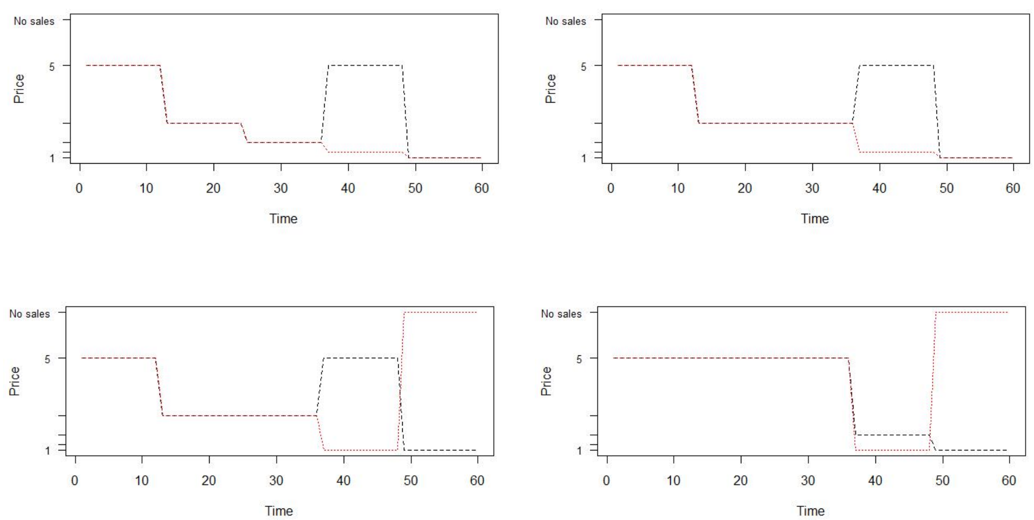

sequences from the low and high sequences, and keep the two sequences (one for the low and one for the high shock) sharing the first three years' prices and with the highest mean value.

Figure 11 shows prices before and after low and high shocks on the same graphs. The top left graph has no piracy. Perfect intertemporal price discrimination is practiced until the shock. After a low shock, the skimming continues as before. After a high shock, prices rise to sell to the newly entered highest valuing consumers, before falling to the lowest rate. The top right graph has a piracy rate of 0.000025. Before the shock, prices start off at the highest valuation before falling to the next highest valuation and remaining there until the shock. Thus, there is a departure from skimming. After the shock, pricing is the same as in the absence of piracy. The bottom left graph shows pricing when the piracy rate is 0.0001. The same pricing before the shock is observed as for the previous graph. After a high shock the same pricing is again observed, but after a low shock prices fall immediately to the lowest valuation followed by market exit in the last period. In the bottom right graph, the piracy rate is 0.0003. Before the shock, the price is held at the highest valuation. After it, there are steep price declines, with some limited skimming being practiced after a high shock alone.

The deterministic analysis of pricing after the shock is familiar from the earlier analysis. Intertemporal price discrimination is optimal if there is no piracy, or if there are large valuation differences among consumers. Piracy compresses price changes.

The analysis of expected value pricing prior to the shock has to consider the effect of piracy after the shock. Without piracy, a pricing strategy to extract value from capacity available in the period prior to the shock does not reduce the value of a strategy to extract value from capacity emerging from the shock. However, with piracy a pricing strategy in the first three years can reduce the value of a strategy applicable to capacity emerging from the shock, as it may alter the number of pirates present after the shock. The reduction would be more severe if the piracy rate is higher. It is therefore beneficial to wait to see whether a high or low shock occurs, and adjust strategy accordingly. Thus, when piracy is higher there are fewer price reductions prior to the shock.

7 Conclusion

In this paper, we have specified a model of pricing in the presence of piracy and with heterogeneous consumers. Piracy is found to lower the profitability of a skimming strategy in favour of a compressed price reduction scheme. In a fixed capacity market, piracy increases welfare, but in a growing market it reduces welfare. The optimal piracy rate choices of consumers, pirate providers, and legal sellers do not generally coincide.

Piracy is found to trade off two effects on sales time and welfare in the presence of market capacity growth, by both delaying product launch and accelerating subsequent sales.

Further work could clarify the mechanisms algebraically and classify them among other possible ones with related impact. Among the possible modifications could be allowance of a role for information spreading as in Givon ***et al.*** (1995), consideration of the impact of partial transfer of product value forward after the end of its lifetime, and inclusion of incentives to innovate if they enter in an analytically substantial way.

We examined uncertainty briefly. More work could examine alternative distributional forms and uncertainty on variables other than capacity. Diffusion timing and welfare consequences could be examined as well. Our model differs from earlier work in assuming persistence in consumer heterogeneity and an increasing marginal propensity to pirate. Both could be tested. For the second of the two assumptions, our recent work (Waters, 2013) provides some evidence that the propensity does increase with the number of users.

Our model is highly stylised, and the assumptions of no rival entry and non-competitive pricing for pirate copies could be relaxed for a more realistic model. The model could then be tested quantitatively or qualitatively. For all the model's stylisation, it seems reasonable that some of its principal predictions would hold even in its current form. For example, a company would feasibly delay a product launch in a developing country prone to piracy until the market was large enough to gain short term profits from entry.

References

Bae, S. and Choi, J. (2006) A model of piracy, *Information Economics and Policy*, **18**, 303–20.

Belleflamme, P. (2002) Pricing Information Goods in the Presence of Copying, Queen Mary, University of London, Department of Economics Working Paper No. 463.

Besen, S. (1986) Private copying, reproduction costs, and the supply of intellectual property, *Information Economics and Policy*, **2**, 5–22.

Besen, S. and Kirby, S. (1989) Private Copying, Appropriability, and Optimal Copying Royalties, *Journal of Law and Economics*, **32**, 255–80.

Cremer, H. and Pestieau, P. (2009) Piracy prevention and the pricing of information goods, *Information Economics and Policy*, **21**, 34–42.

Geroski, P. (2000) Models of technology diffusion, *Research Policy*, **29**, 603–25.

Givon, M., Mahajan, V. and Muller, E. (1995) Software Piracy: Estimation of Lost Sales and the Impact on Software Diffusion, *Journal of Marketing*, **59**, 29–37.

Givon, M., Mahajan, V. and Muller, E. (1997) Assessing the Relationship between the User-Based Market Share and Unit Sales-Based Market Share for Pirated Software

Brands in Competitive Markets, *Technological Forecasting and Social Change*, **55**, 131–44.

Haruvy, E., Mahajan, V. and Prasad, A. (2004) The Effect of Piracy on the Market Penetration of Subscription Software, *Journal of Business*, **77**, S81–107.

Herings, P.-J., Peeters, R. and Yang, M. (2009) Piracy on the internet: Accommodate it or fight it? A dynamic approach, Maastricht University Faculty of Economics and Business Administration Mimeo RM/09/034.

Johnson, W. (1985) The economics of copying, *Journal of Political Economy*, **93**, 158–74.

Khouja, M., Hadzikadic, M., Rajagopalan, H. and Tsay, L.-S. (2008) Application of complex adaptive systems to pricing of reproducible information goods, *Decision Support Systems*, **44**, 725–39.

Khouja, M. and Smith, M. (2007) Optimal pricing for information goods with piracy and saturation effect, *European Journal of Operational Research*, **176**, 482–97.

Liu, H. (2010) Dynamics of Pricing in the Video Game Console Market: Skimming or Penetration?, *Journal of Marketing Research*, **47**, 428–43.

Liu, Y., Cheng, H., Tang, Q. and Eryarsoy, E. (2011) Optimal software pricing in the presence of piracy and word-of-mouth effect, *Decision Support Systems*, **51**, 99–107.

Nair, H. (2007) Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video-games, *Quantitative Market Economics*, **5**, 239–92.

Novos, I. and Waldman, M. (1984) The Effects of Increased Copyright Protection: An Analytic Approach, *Journal of Political Economy*, **92**, 236–46.

Prasad, A. and Mahajan, V. (2003) How many pirates should a software firm tolerate? An analysis of piracy protection on the diffusion of software, *International Journal of Research in Marketing*, **20**, 337–53.

R Development Core Team (2009) *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, http://www.R-project.org.

Takeyama, L. (1994) The Welfare Implications of Unauthorized Reproduction of Intellectual Property in the Presence of Demand Network Externalities, *Journal of Industrial Economics*, **42**, 155–66.

Waters, J. (2013) Variable marginal propensities to pirate and the diffusion of computer software, Mimeo: Nottingham University Business School.

Yoon, K. (2002) The optimal level of copyright protection, ***Information Economics and Policy***, **14**, 327–48.