Network Externalities and Copyright Enforcement

David Blackburn*

May 2006

Abstract

A monopolist selling a good subject to consumer copying and network externalities triggers two effects by not fully enforcing copyright protection. The copying effect reduces demand due the introduction of copies as an outside alternative. However, a "network augmenting" effect increases demand through the increased network size due to copies. If the marginal network externality is large enough, the monopolist increases profits by allowing some level of copying to occur. In a dynamic setting, I show that the typical pattern of copyright enforcement involves little enforcement at product introduction, followed by increased enforcement as the market and network grow.

*National Economic Research Associates, david.blackburn@nera.com. A previous version of this paper, titled "Complementarities and Network Externalities in Casually Copied Goods" was presented at the Conference on Economic Development, Technology, and Human Resources in Tafi de Valle, Argentina in December 2001 and published in the conference summary in *Estudios de Economia*, Vol. 29 (1). p 71-88. June 2002. This version of the paper is a revision of the first chapter of my 2005 Harvard University PhD dissertation. The author would like to thank Julie Mortimer, Markus Mobius, and Walter Cont as well as seminar participants at Harvard University, Universidad de Chile, and Universidad Nacional de Tucuman (Argentina) for helpful comments and suggestions. All remaing errors are my own.

1 Introduction

In recent years, there has been much focus in the popular media on the detrimental effects that unauthorized duplication of original work, such as software piracy and copying of music and film, has had on the affected industries. For example, the Business Software Alliance (BSA) estimates that over 20% of software installed in the United States in 2004 was unlicensed, a loss of US$6.6 billion due to unauthorized reproduction of software. Furthermore, the Recording Industry Association of America (RIAA), which has engaged in multiple lawsuits regarding unauthorized downloads of copyrighted recordings, has repeatedly blamed billions of dollars of lost sales on piracy. The estimates from the BSA and RIAA tend to be based on the assumption that all copied good would otherwise be purchased through authorized channels, an assumption that economic logic can not support. In fact, it is possible that very few of those with illegal copies would have purchased these items if it were impossible to copy originals.[1]

In this paper, I examine the impacts of unauthorized copying and copyright enforcement in the market for goods that have strong network externalities and/or are strong complements with goods in another market. In particular, I assume that copyright protection is available and then examine when and to what extent a monopolistic producer of a network good would choose to enforce that protection to prevent unauthorized copying.[2] I find that in circumstances where the marginal network externality to consumers is sufficiently large, the monopolist will actually increase her profits by choosing a less than full level of copyright enforcement.

Although this may seem paradoxical, the intuition is quite simple. By choosing not to enforce copyright protections fully, the monopolist triggers two effects. First, the po-

[1]See Blackburn (2004), Oberholzer-Gee and Strumpf (2004), and Zentner (2003) for various empirical analysis of the effects of copying on the sales of recorded music.

[2]Even if the firm choses not to enforce copyright protection and thus implicitly allow copying, we still will refer to copying as unauthorized.

tential availability of copies to consumers increases the outside alternative to purchasing a new good, thereby reducing the amount of surplus the monopolist could potentially extract, which would necessarily push towards a reduction in the monopolist's profits. However, a second, "network augmenting" effect pushes the monopolist's profits in the other direction. Copying increases the size of the network beyond what it would be if it were restricted to sales. The increase in the size of the network increases each consumer's willingness to pay for the good, through the network externality (or through the complementary market), thereby increasing the amount of potentially extractable surplus. If this "network augmenting" effect is strong enough, it will outweigh the traditional copying effect and the monopolist will achieve higher profits through reduced copyright enforcement.

I then consider a dynamic game and the implications that this has on the monopolist's evolving choice of copyright enforcement over time. Under simple assumptions about the form of the externalities, the dynamic analysis demonstrates that as time passes and the market for the good becomes more mature, firms will increasingly seek to enforce tighter and tighter copyright protection. The intuition here is simple as well. As the market becomes more mature, the good has a smaller marginal network externality from an additional good, so the "network augmenting" effect that suggests that a monopolist may prefer some positive level of copying is getting weaker and weaker. Thus, the monopolist will attempt to make it harder for copies to spread.

Previous work by Takeyama (1994) has studied, in a discrete consumer model, the impacts of copying on static monopolist profits in the face of demand network externalities, with very similar results.[3] Takeyama (1997) also studies copying in a dynamic setting,

[3]Takeyama analyzes a static monopolist facing a fixed level of copying. She finds that, given a fixed level of degredation between copies and originals, a firm which faces a large enough network externality prefers not to enforce a copyright. The model developed in this paper extends this idea to allow firms the ability to enforce a copyright with variable degrees of vigor and to change this behavior over time. Thus, a firm facing the same network externality in both models might prefer no copying in the Takeyama world while choosing to allow a low level of copying in this model.

absent network externalities, and finds that copying can allow the monopolist to overcome problems associated with the Coase conjecture.[4] Work on shared goods such as libraries, families, and rental markets has found that the impact that sharing has on producer profits depends critically on the assumption made about the formation of groups and the ability of the monopolist to appropriate the additional surplus that the secondary, shared uses add to the 'single-use' good.[5] In particular, Varian (2000) and Bakos, Brynjolfsson, and Lichtman (1999) find that when groups are sufficiently heterogenous, the monopolist is able to achieve larger profits than if it was necessary to sell a single good to each user. In an empirical examination of the academic journal industry, Leibowitz (1985) finds that the introduction of photocopying technology, rather than decreasing profits for journal publishers, increased demand for journals at collegiate libraries.

In a similar vein, there exists work that suggests that it may be best for a monopolist to license the production of her good to other firms and compete in an oligopoly. Shepard's (1987) model finds that second-sourcing commits the monopolist to higher quality levels, which can increase industry demand enough to offset the loss in market power. However, Economides (1996), in a network externality framework like minepower and withhold supply. Of course, licensing differs sufficiently from using "unauthorized" copying to expand the network size. While the technology is licensed to create a large market size, the previous monopolist must now compete with other firms to serve high value consumers. However, when unauthorized copying is used, the firm retains monopoly power and faces no competitors for high value consumers.[6]

The rest of the paper proceeds as follows: First, in Section 2, I describe the set-up of the

[4]The Coase conjecture, recall, states that the monopolist producer of a fully durable good will necessarily price at marginal cost (net of any discouting premium), as she essentially competes against future incarnations of herself.

[5]See, for example, Varian (2000), Besen and Kirby (1989), and Bakos, Brynjolfsson, and Lichtman (1999).

[6]Although, the introduction of copies into the market implies that the monopolist can not extract the same amount of surplus as if copying was not possible.

environment that the monopolist faces. I describe the market for goods complementary to the monopolist's in Section 3. Section 4 describes the monopolist's good and performs the static analysis of the monopolist's problem under an exogenous level of copying (where the firm's choice is between not allowing copying at all or allowing copying at the exogenous level). Section 5 extends this to analyzes the monopolist's optimal choice over the level of copying, and section 6 introduces a dynamic game to examine how the choice of copy protection by the monopolist changes over time as the market grows. Finally, Section 7 concludes.

2 The Model

The model considers the monopolist producer of a hardware good (such as an operating system). The hardware good that the monopolist produces is subject to unauthorized copying by those consumers who do not purchase an original. Furthermore, in the spirit of Katz and Shapiro (1985), the hardware good exhibits network externalities in that the value of the good to consumers increases in the number of goods that is consumed by other consumers (the network size). For simplicity I assume that the monopolist faces constant marginal cost, and normalize that value to zero.

In addition to the direct network externality that the good exhibits, the hardware good is used in conjunction with complementary software goods. The monopolist is not a participant in the software market, but the availability of third-party software impacts the utility of consumers of the hardware good. That is, consumers have zero demand for goods in the software market unless they consume the hardware good as well. The software market is assumed not to be subject to concerns about copying.

While it may seem strange to think of hardware as being subject to copying, rather than software, it is not hard to think of examples which fit into the framework. For example, Microsoft Windows could be the hardware, and the software market is then the market for third-party applications. Alternatively, a web browser or application (Excel) could be the relevant hardware, and software would represent add-on modules or plug-ins that are used in conjunction with the application. And a parallel could be drawn to the relationship between recorded music sales and complimentary products like concert ticket and merchandise sales, though it is not a perfect parallel because tickets and merchandise can be sold to non-consumers of recorded music. The exact assumed relationship between hardware and software is explained below, where it is shown that (as in Economides (1996)), the existence of the complementary third-party software market reduces the studied problem to one of a traditional network externality for consumption of the hardware good.

3 Software Market

The software market is assumed to be a standard differentiated product model (a circular city a la Salop) with circumference equal to 1 and uniform density equal to the size of the hardware market, S. Essentially, then, every consumer of a hardware good exists on a point on the software circular city. That is, as the hardware market grows, so does the potential size of the software market.

Consumers in the circular city have gross surplus from a piece of software equal to K (large enough to ensure that all consumers will purchase software) and linear transportation cost equal to d.[7] There is a fixed cost of entry into the software market of f, and I assume that there is zero marginal cost for production of a good. This yields a price for firm i of $p_i = d/n$ and a quantity for each firm of $x_i = S/n$. The profit of a firm in the industry is then $\pi = \frac{S}{n}\frac{S*d}{n} - f$, and thus the zero profit condition implies that the number of software

[7] This is a simplifing assumption. Relaxing it does not change the results qualitatively.

firms is:

$$n = S\sqrt{\frac{d}{f}}$$

Thus, a larger installed hardware base implies a larger number of firms (more variety) in the software industry.

The model makes two important assumptions about the software market. First, the hardware producer does not enter the software market and second, there is no copying in the software market. Intuitively, allowing the monopolist to produce in the software market would increase the monopolist's willingness to allow copying. The monopolist faces the same effects as discussed below in the hardware market, but also in the software market. Of course, as specified in the model, free entry leads to a zero-profit condition in the software market and thus allowing the monopolist into this market would have no effect. If, however, there was some market power in the software market, then the monolopolist's existence in that market should increase the willingness to allow copying of hardware.

On the other hand, the effects of allowing copying in the software market are not as simple. In the simple software market model above, copying in the software market would simply reduce the profits made by software firms, thus reducing the number of firms in the market, resulting in less variety and reduced gains to consumers in the hardware market. Thus the hardware firm may wish to enforce copyright protection in the software market, in order to protect that variety, or even to enter the software market to help provide sufficient variety to entice consumers to adopt the hardware.[8]

For simplicity and in order to isolate the hardware firm's incentives to enforce copyright protection in the hardware market, these extentions are left to future work and I proceed with the model outlined above.

[8]Of course, one could imagine a more complicated model of the software market in which firms have market power and consumers buy mulitple pieces of software in which copying could have many different effects on consumers and firms in both markets.

4 Hardware Market

I first consider the case of the monopolist's decision whether or not to allow copying in a static framework, where the degree of copying is exogenous, and represented by the parameter α, described below.[9] The timing is as follows: first consumers form expectations about the size of the hardware network and the number of firms in the software market, then the monopolist prices, taking into account the expectations of consumer, and finally consumers make purchase and copying decisions, and the size of the network and the software market are realized.

4.1 Consumers

The model of consumer behavior follows the lines of Katz and Shapiro (1985). Consumers in the hardware market form expectations about the size of the hardware market as well as the implied number of firms in the software market. Given an expected size of the hardware market of S^E and an expected number of firms in the software market of n^E, each consumer has expected surplus associated with consumption of the hardware good equal to $r + v(S^E) + w(n^E)$. I assume that $v', w' > 0$, $v'', w'' \leq 0$, and also that $v(0) = w(0) = 0$; that is, there are no network effects resulting from a network of size 0.[10] Thus, r can be interpreted as the consumer's valuation of the hardware independent of any other consumer's choice or the availability of any software products.

The function $v(S)$ is a standard direct network externality function, while $w(n)$ is an indirect network externality which represents the increase in each consumer's value of the good as the amount of variety in the software market increases. Each consumer has the

[9]This section follows Takeyama (1994) closely.

[10]The assumption of (weakly) concave network externality functions is standard in the literature. It is sufficient to assume that there is a network size S^* and a market size n^*, such that for all $S > S^*$ and $n > n^*$, $v'(S) + w'(n) < 1$, which rules out the possibility of infinite network sizes being realized.

same v and w, so the expected sizes of the network and the complementary market have the same effect on all consumers. However, consumers differ in the level of the independent valuation r. I assume a continuum of consumers, with values of r that are uniformly distributed from $-\infty$ to θ.[11] Thus, the value of the hardware to a consumer who pays a price p for the hardware is:

$$U^{\text{buying}} = r + v(S^E) + w(n^E) - p$$

Although consumers are forming expectations about network size, utility need not integrate over potential network sizes, because I will focus on rational expectations equilibrium, where expected network size is actualized. Thus, forward looking consumers know that the expected network size will be realized, and there is no uncertainty.

The expected size of the hardware network, S^E, implies that the expected number of firms in the software market is $n^E = S^E\sqrt{\frac{d}{f}}$. Defining $\widetilde{w}(x) = w(x\sqrt{\frac{d}{f}})$, an individual consumer's net surplus of the hardware good as:

$$U^{\text{buying}} = r + v(S^E) + \widetilde{w}(S^E) - p$$

Thus the software/hardware paradigm is observationally equivalent to, and can be modeled as, an additional traditional network externality.

Consumers who decide not to purchase the hardware from the monopolist but rather decide to obtain an unauthorized copy of the good receive surplus that equals an exogenous

[11]This assumption will create a linear demand curve in the market. Furthermore, the normalization of marginal cost to zero somewhat alters the interpretation of r, as it can now be interpreted as the consumer's value of the good above marginal cost, absent any network effect. Thus, negative values of r have a natural interpretation. However, there is an implicit assumption in this set-up that the marginal cost of a copy is equal to the marginal cost of an orginal. For goods distributed on CDs, this assumption is approximately correct.

fraction, $\alpha \in [0,1]$, of the gross surplus from buying the good, so that:

$$U^{\text{copy}} = \alpha(r + v(S^E) + \widetilde{w}(S^E))$$

The fraction α can be interpreted as the probability that the consumers receives a (perfect) copy of the original, and so α represents the degree to which unauthorized copying exists in the market.[12] A high α corresponds to a large degree of copying (little copyright enforcement on the part of the monopolist), while a low α corresponds to little copying (or strong enforcement).[13] It is assumed that this is the only price associated with a copy (the marginal cost of a copy is assumed to be zero) and $\alpha = 0$ represents a no-copying regime with absolute copyright enforcement.[14]

Given the monopolist's price p, a consumer will prefer to purchase rather than copy if

$$(1-\alpha)(r + v(S^E) + \widetilde{w}(S^E)) > p$$

or

$$r > \frac{p}{1-\alpha} - v(S^E) - \widetilde{w}(S^E)$$

[12] It is likely possible to obtain an illegal copy of almost any piece of software. However, in practice, differing levels of copy protection mean that it will take longer to obtain copies of some software. Search and time, however, do not exist in this model, but can be thought to exist through the parameter α. Increased copyright enforcement makes it harder (in terms of realized search time) for consumers to obtain a copy. If consumers are only willing to search for finite time, then they become more likely not to be successful in their search as copy protection is increased. Thus more copy protection results in a low value of α.

[13] It may seem appropriate to have the probability of obtaining a copy be proportional to the amount of originals that are sold. However, given that a monopolist facing a straight-line demand curve maximizes profit by selling to half the market, it will always be that one half of consumers who have a positive gross valuation of the good purchase originals. Therefore, nothing is lost (within this model) by assuming that α is independent of sales.

[14] Under an alternative interpretation, the consumer is always able to obtain a copy of an original, but the copy is imperfect, and α represents the degree of imperfection or degradation of a copy. This degradation could be the result of a lack of technical support, locked features that require registration, or simple imperfections from copying, such as a reduction in image or audio quality. The two interpretations of α yield slightly different results, since they imply different sizes of the final network. I focus on the probability of acquisition alternative because the degredation interpretation yields a degenerate optimal strategy- set α as close to zero as possible and enjoy the full benefit of the network effect with virtually no lost sales.

Thus the mass of consumers who will purchase an original, given consumer expectations equal to S^E and level of copying equal to α, is

$$\{r : \frac{p}{1-\alpha} - v(S^E) - \widetilde{w}(S^E) < r < \theta\}$$

and the demand curve faced by the monopolist is

$$D = \theta - \frac{p}{1-\alpha} + v(S^E) + \widetilde{w}(S^E)$$

due to the uniformity assumption.

4.2 Monopolist

The monopolist's profit function is then:

$$\pi(p, S^E) = p(\theta - \frac{p}{1-\alpha} + v(S^E) + \widetilde{w}(S^E))$$

Solving the FOC implies that:

$$p(S^E) = \frac{1-\alpha}{2}(\theta + v(S^E) + \widetilde{w}(S^E))$$
$$q(S^E) = \frac{1}{2}(\theta + v(S^E) + \widetilde{w}(S^E))$$

where $q(S^E)$ is the quantity of originals copies supplied to the market by the monopolist when consumer expectations of market size are S^E. It follows that profit for the monopolist is

$$\pi(S^E) = \frac{1-\alpha}{4}(\theta + v(S^E) + \widetilde{w}(S^E))^2$$

As Katz and Shapiro noted, there are many market quantities corresponding to different expectations of network size, and thus, as mentioned above, it will be worthwhile to focus on rational expectation equilibrium where the consumers expectations about the size of the hardware network (and thus the number of software firms) are realized, $S^E = S$.

4.3 Equilibrium

The actual network size is equal to the sum of the number of consumers who purchase the hardware and the number of consumers who are able to obtain a copy. Thus,

$$S = q(S^E) + \text{quantity copied}$$

The total quantity copied will be $\alpha(\frac{1}{2}(\theta + v(S^E) + \widetilde{w}(S^E)))$, and we have that $S^C = \frac{1+\alpha}{2}(\theta + v(S^E) + \widetilde{w}(S^E))$.

Imposing rational expectations equilibrium requires that $S = S^E$, and so we have that S^C (the network size with copying) and S^{NC} (the network size with no copying) solve:

$$S^C = \frac{1+\alpha}{2}(\theta + v(S^C) + \widetilde{w}(S^C))$$
$$S^{NC} = \frac{1}{2}(\theta + v(S^{NC}) + \widetilde{w}(S^{NC}))$$

Given the concavity assumptions on v and w, we know that $v + w$ is also concave and thus S^{NC} and S^Chave unique solutions, as can be seen in Figure 1.

It is immediate that $S^C > S^{NC\ (\alpha=0)}$; that is, the size of the network is increasing in the amount of copying, α.

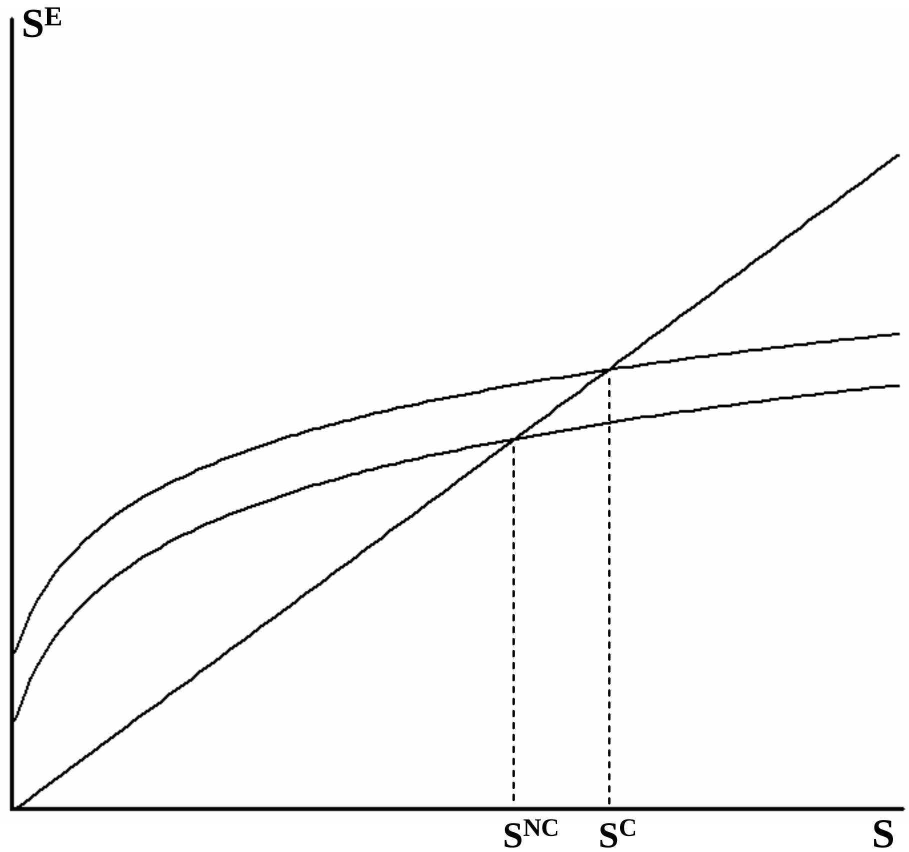

Figure 1: Rational Expectations Equilibrium

4.4 Copyright Enforcement

I begin by examining the simple binary choice of whether the monopolist should allow copying or not, taking α as given.[15] Not exercising full copyright enforcement and thus allowing unauthorized copying generates two effects which are illustrated in Figure 2. First, copying $(\alpha > 0)$ implies a larger outside option for consumers in the hardware market. This copying effect results in the demand curve rotating inward. By itself this implies that the firm's profits will be lower, as sales do not decrease, but the price is reduced. However, the second effect is that the increased installed hardware base that copying creates

[15]This is the result found in Takeyama (1994).

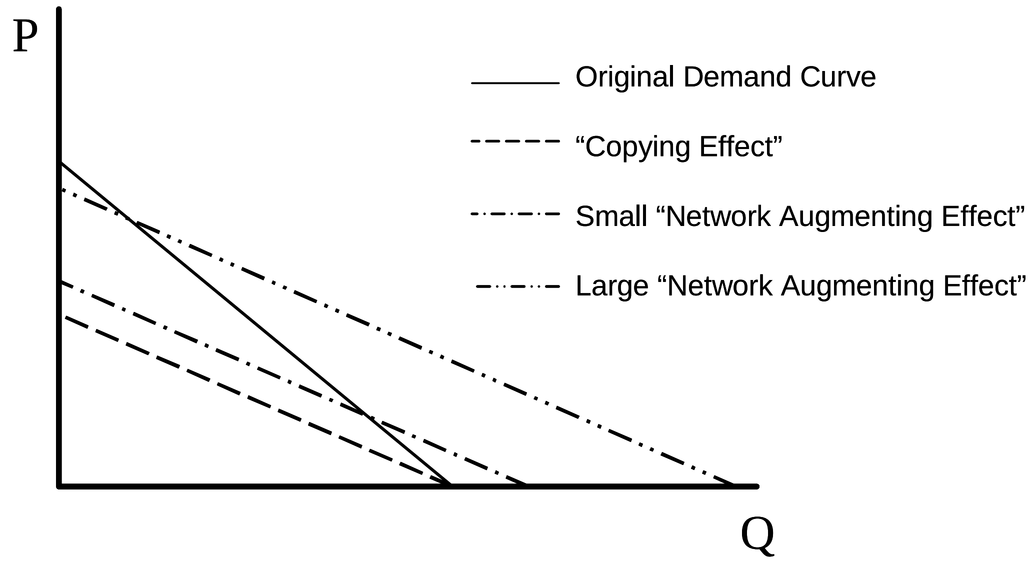

Figure 2: “Copying Effect” and “Network Augmenting Effect” of Copying on Demand

implies a larger willingness to pay by consumers, which shifts the demand curve outward. This “network-augmenting” effect can take place through two potential channels. First, there could be a pure network effect due to the fact that there are many more hardware consumers (the $v(\cdot)$ function). Secondly, the larger base of hardware components implies larger variety in the complementary software network (the $w(\cdot)$ function). These two effects imply that the demand curve is shifted outward, which results in a larger price and quantity for the monopolist.

The net impact of these two effects depends on the size of α (which is the rotation of the demand curve in Figure 2) and the marginal size of the network effects $v(\cdot)$ and $w(\cdot)$, which is the shift in the demand curve in Figure 2. A larger α implies a larger inward rotation of the demand curve (as well as a larger jump in S from copying). And larger network effects imply a larger outward shift in the demand function. If the second effect dominates, so that there is a large marginal “network-augmenting” effect relative to the copying effect, firms will find it profitable to allow copying at a level α. Thus, we might expect that in an industry with large network effects, monopolists would rather turn a blind eye to copying. However, if the marginal network effect is small, then the inward rotation of demand will dominate and the firm will prefer to take means to prevent copying and restore $\alpha = 0$.

Mathematically, the difference in profits to the monopolist (which depends on how α is interpreted) is:

$$\begin{aligned}\pi_C - \pi_{NC} &= \frac{1-\alpha}{4}(\theta + v(S^C) + \widetilde{w}(S^C))^2 - \frac{1}{4}(\theta + v(S^{NC}) + \widetilde{w}(S^{NC}))^2 \\ &= \frac{-\alpha}{4}\theta^2 + \frac{1-\alpha}{2}\theta(v(S^C) + \widetilde{w}(S^C)) - \frac{1}{2}\theta(v(S^{NC}) + \widetilde{w}(S^{NC})) \\ &+ \frac{1-\alpha}{4}(v(S^C) + \widetilde{w}(S^C))^2 - \frac{1}{4}(v(S^{NC}) + \widetilde{w}(S^{NC}))^2\end{aligned}$$

The first term ($\frac{-\alpha}{4}\theta^2$) is the copying effect which represents the unambiguous loss of extractable surplus from the consumer. The second term $\frac{1-\alpha}{2}\theta(v(S^C)+\widetilde{w}(S^C))-\frac{1}{2}\theta(v(S^{NC})+\widetilde{w}(S^{NC})) + \frac{1-\alpha}{4}(v(S^C) + \widetilde{w}(S^C))^2 - \frac{1}{4}(v(S^{NC}) + \widetilde{w}(S^{NC}))^2$ is the "network-augmenting" effect on demand which represents the larger surplus that each consumer receives because of the larger network size when copying is allowed. The difference in profits is positive (and the monopolist prefers that copying exists) when the second effect is larger. Notice that the monopolist is more likely to allow copying when the second interpretation of α is taken, as this implies that the network size will be as large as possible (subject to the rational expectations equilibrium restriction), and thus the network externalities will be maximized.

As noted in Takeyama (1994), the potential advantage here is one of price discrimination. The monopolist is able essentially to price discriminate between high and low value consumers, setting a positive price at which high value consumers will purchase and a zero price for low value consumers. Thus, the monopolist is able to induce a large network size without setting a uniformly low price to induce low value consumers to purchase, and instead the monopolist uses copying to achieve this effect and extract more rents from high valuation consumers.

4.5 Sales Displacement

The fact that the monopolist can use copying to price discriminate depends crucially on the assumption above that there is no correlation between r, a consumer's baseline valuation, and α, the extent of copying. As modeled above, every consumer faces the same α; no consumers have a high or lower proclivity for (or cost of) copying, or ability to copy. Thus, consumers sort only on their valuation of the good (and network), and the monopolist is potentially able to exploit that to their advantage, and copying either does not displace sales at all, or sales are displaced only for the low end of buyers.

However, it is possible that different types of consumers may have differing copying abilities. If there are groups of consumers who would face a high level of α, so that they can either obtain a copy with higher probability or obtain a more functional copy (depending on the interpretation of α), it may well be that the existence of copying is crippling to the monopolist. Of course, if it is the case that high value consumers are the ones with the lowest proclivities for copying (so that there is a negative correlation between r and α), copying is more appealing to the monopolist. In this case, the consumers who copy will naturally be the ones to whom the hardware monopolist will not price and results similar, but slightly stronger in favor of copying, to the ones above would be obtained. Again, displacement occurs only for the lowest value consumers, and no high-valuation sales are displaced by copying.

The trouble for the hardware producer is if there was a positive correlation between r and α. In this case, the availability of copies can have a debilitating effect on the hardware producers. If the consumers with the high valuations of the good and network are those for whom it is easiest to obtain a functional copy, then copying can cannibalize the hardware market. The monopolist can no longer extract sufficient surplus from those with the highest valuations in order to profit from copying's network-enhancing effects, as those consumers can easily obtain a copy rather than purchase. In the extreme case where $a = 1$ for the highest valuation consumers, it is impossible to obtain any surplus from these high valuation consumers as copying yields the full value of the good and network at no cost. If copying existed in a form like this, all the monopolist could do would be to serve the middle of the market by setting an intermediate price at which consumers with a high enough valuation, but low enough copying skills, would purchase from them, while low valuation consumers and the highest valuation consumers do not purchase and instead copy.

I now turn to a numerical example to highlight the effect of a positive correlation between consumer's valuations of the goods and their proclivity for copying. I must use a numerical example because analytical solutions to even simple forms of correlations between r and α are unavailable. The primitives of the model are as such:

$$v(S^E) + \widetilde{w}(S^E) = \frac{3}{4}S^E$$

$$\theta = 100$$

With these primatives, it is easy to solve the model if there is no copying whatsoever. In that case, the monopolist sets a price of $p^{NC} = 80$ and sells a quantity of $q^{NC} = 80$, resulting in a network size of $S^{NC} = 80$ and profits of $\pi^{NC} = 6400$. It is also simple to solve the model in the face of simple copying, when $\alpha = 1/3$ for all consumers. In this case, the monopolist now sets a lower price of $p' = 66.67$ but, due to the increased network size from copying, is able to sell a total of $q' = 100$, and a total of 33.33 are obtained by non-purchasing consumers. This results in a total profit to the hardware monopolist of

$\pi' = 6667$, a 4% increase in profits over the no copying case.[16]

To see how asymmetric copying can cannibalize the profits of the hardware producer, I numerically solve the hardware producer's problem when copying takes the form:

$$\alpha_i = \max\{0, \frac{1}{3} + \frac{r_i}{3\theta}\}$$

Thus, the average amount of copying (among those consumers for whom $r_i > -\theta$) is equal to $\frac{1}{3}$, as above. However, now copying proclivity is positively correlated with consumer's valuations, and thus we will see that this makes taking advantage of copying impossible for the hardware producer. Now, the best the hardware producer can do is to sell at a price of $p'' = 70.4$, thus selling a total quantity of 72.9, for a total profit of only $\pi'' = 5132.2$. Consumers copy a total of $41.\,2$ copies. What causes this? The main problem is that copies are sold not to the consumers with the highest valuations, but rather to consumers on $r \in (-31.6, 41.3)$, as consumers with an independent valuation greater than 41.3 are able to obtain a copy with high enough probability (or high enough quality) that the hardware monopolist is unable to profitably price to them.

Thus, we see the problem that a positive correlation between copying and valuations can have for the monopolist, even when the average amount of copying in the market essentially the same as the ideal amount of copying, when copying is constant across consumers. Throughout the rest of the paper, I will return to the case of symmetric copying abilities, though the lesson explored here is germane for the rest of the applications in the paper. That is, a positive correlation between r and α is as destructive for the hardware producer in the other settings explored as well. In order for copying to provide and advantage to the monopolist, there can not be "too much" correlation between consumer's valuations and copying.

[16]As will be shown in the next section, $\alpha = 1/3$ is the profit-maximizing level of copying if the combined network externality function is $\frac{3}{4}S^E$.

5 Choosing a Level of Copyright Enforcement

I now return to the case where there is no correlation between r and α, and extend the model to allow the monopolist to choose a level of copyright enforcement.[17] That is, the firm can choose a value for α. We can imagine this happening through the firm withholding support for non-registered users, locking features within the good, or employing sophisticated duplication prevention methods in production of the good, or more diligent attempts at enforcing copyrights either privately or through law enforcement. Whatever form copy protection takes, the level of copy protection associated with a good is certainly one choice variable for a firm in product design. In this framework, the monopolist is simply able to choose α.[18]

From above, note that the firm's profits are purely a function of the choice of α. Given an α, the realized network size and profit level are determined from the optimal pricing behavior discussed above. Therefore, the firm can directly consider the effect that a change in α has on profits.

$$\frac{d\pi}{d\alpha} = \frac{\partial \pi}{\partial \alpha} + \frac{\partial \pi}{\partial S}\frac{dS}{d\alpha}$$

That is, the change in the monopolist's profits from a differential decrease in the level of copy protection (as α increases, the level of copy protection decreases) is composed of the two effects mentioned above. Again, there is the "copying effect" which is the direct negative impact on the profit function ($\frac{\partial \pi}{\partial \alpha} < 0$). As it becomes easier for consumers to get copies, the amount of surplus that the monopolist can extract is reduced because

[17] The effects of correlation between valuations and copying abilities are similar throughout all the extentions considered.

[18] The reason the alternative interpretation of α as a copy-degredation parameter is not interesting can be seen here. Because the size of the network does not depend on α, the firm's optimal choice would be to set α as close to zero as possible, which would take full advantage of the network-augmenting effect with (essentially) zero loss associated with the copying effect.

the consumer now has a better outside alternative to purchasing the good. But it is the second effect that suggests that a monopolist might want to allow copying. The second term ($\frac{\partial \pi}{\partial S}\frac{dS}{d\alpha} > 0$) is the "network-augmenting" effect that copying has on profits through the increased network size that copying creates. As the level of copy protection decreases, realized network size increases, which increase the extractable surplus of consumers. If the "network-augmenting" effect outweighs the copying effect, then the monopolist will find it optimal to increase α and allow an increase in level of copying.

From above,

$$\frac{dS}{d\alpha} = \frac{\frac{1}{2}(\theta + v(S) + \widetilde{w}(S))}{1 - \frac{1+\alpha}{2}(v'(S) + \widetilde{w}'(S))} > 0$$

Note that $\frac{dS}{d\alpha} > 0$ because in equilibrium the concavity assumptions on v and w imply that $v'(S) + \widetilde{w}'(S) < 1$. So, increasing α has a bigger effect on the realized hardware base when the marginal network externality is larger.

Differentiating the profit function completely yields

$$\begin{aligned}\frac{d\pi}{d\alpha} &= -\frac{1}{4}(\theta + v(S) + \widetilde{w}(S))^2 + \frac{1-\alpha}{2}(\theta + v(S) + \widetilde{w}(S))(v'(S) + \widetilde{w}'(S))\frac{dS}{d\alpha} \\ &= (\theta + v(S) + \widetilde{w}(S))^2 \left[\frac{-1}{4} + \frac{1-\alpha}{2}\frac{(v'(S) + \widetilde{w}'(S))}{2 - (1+\alpha)(v'(S) + \widetilde{w}'(S))}\right]\end{aligned}$$

So, the copying effect is proportional to $\frac{1}{4}$ while the network effect is proportional to $\frac{1-\alpha}{2} \times \frac{(v'(S)+\widetilde{w}'(S))}{2-(1+\alpha)(v'(S)+\widetilde{w}'(S))}$. For simplicity letting $f(S) = v(S) + \widetilde{w}(S)$, we have that the network effect is proportional to $\frac{1-\alpha}{2}\frac{f'(S)}{2-(1+\alpha)f'(S)}$. We know immediately that the firm will not choose $\alpha = 1$, because at $\alpha = 1$, the firm makes zero profit, while any other level of α yields positive profits. Noting that $\frac{d\pi}{d\alpha}$ is monotonically decreasing in α, to determine if the firm prefers some level of copying to no copying at all, it is sufficient to check the sign of

$\frac{d\pi}{d\alpha}$ at $\alpha = 0$

$$\frac{d\pi}{d\alpha}(\alpha = 0) > 0 \Leftrightarrow \frac{-1}{4} + \frac{1-\alpha}{2}\frac{f'(S)}{2-(1+\alpha)f'(S)} > 0$$
$$f'(S) > \frac{2}{3}$$

Therefore, if the composite marginal network externality is strong enough ($f'(S) = v'(S) + \widetilde{w}'(S) > \frac{2}{3}$), then the monopolist increases her profits by relaxing copyright protection somewhat. This is intuitive, because the firm will benefit from some level of copying if the boost in network size it gets through the increased network externality as copies proliferate is large enough to overcome the negative copying effect.

Allowing the firm to optimize over the choice of α leads to the condition $\frac{d\pi}{d\alpha} = 0$, subject to the constraint that $\alpha \in [0, 1]$, which yields

$$\alpha^* = \left\{ \begin{array}{cc} \frac{3f'(S)-2}{f'(S)} & f'(S) > \frac{2}{3} \\ 0 & f'(S) \leq \frac{2}{3} \end{array} \right\}$$

Where S is determined through the rational expectations equilibrium given α.[19] So, as the marginal network externality decreases, the monopolist is willing to allow less and less copying. The intuition for this is straightforward: as the demand network externality becomes less and less important, the monopolist gains less by allowing non-purchasers the ability to get an 'unauthorized' copy of the good, and thus desires a higher level of copy protection. Thus, for goods that have a strong marginal network externality, the monopolists should be willing to allow a (relatively) large degree of copying, while a good that has little (or no) marginal network externality would attempt to enforce full copyright protec-

[19] If the combined network externality function is linear, $f(S) = kS$, so that $f'(S) = k$, (with $k < 1$) then we have that $\alpha^* = \frac{3k-2}{k}$ for $k > \frac{2}{3}$ and $\alpha^* = 0$ for $k \leq \frac{2}{3}$. So the hardware monopolist prefers some level of copying if the marginal network effect is bigger than $\frac{2}{3}$. This corresponds to the linear externality case in Economides (1996), who finds that a monopolist will prefer to license to at least one other firm if $k > \frac{2}{3}$.

tion.

Although I have assumed throughout that there is no cost involved in enforcing copyright protection, introducing a cost of enforcing a copyright would have little qualitative impact, and would just lower the bound on the marginal network externality above which the firm would allow copying. That is, were copyright enforcement to be costly, the monopolist would be more willing to allow copying and thus, in this static framework, little is gained by doing so. After discussing the dynamic game in the next section, I will discuss again how costly enforcement might affect the analysis and implications.

6 A Dynamic Game

Finally, I introduce a dynamic game in which the monopolist sells hardware over multiple periods to examine what path copy protection takes over the life-cycle of a product. There are T periods $t = 1, 2, ..., T$, and the monopolist faces a distinct, but identical, set of consumers in each period. That is, consumers live for only one period, while the monopolist lives forever. This allows us to avoid the typical problems related to monopolist pricing in dynamic settings, such as the Coase conjecture, as well as to abstract away from the possibility of intertemporal substitution by consumers that live throughout multiple periods. Instead I focus on the copyright enforcement choice of the monopolist.[20]

I assume that some of the network externality that exists in period t continues on to period $t + 1$. In particular, I will assume that the variety of products in the software market in period t is still available in period $t + 1$, along with any new products that arise in the software market in period $t+1$ as a result of the hardware network in period $t+1$. Thus, let

[20] This problem is also discussed in Dinlersoz and Pereira (2004). It is noted within their discussion that "introducing dynamic decision for consumers is not difficult. It is well-known that (see, e.g., Sargent and Ljungqvist (2000)), in such a dynamic environment, a consumer faces a buy-or-wait problem and holds a reservation price in deciding whether to buy." Unfortunately, the tractability of the algebra in this model makes additional complications not worthwhile.

n_t be the number of new software products developed in period t (so that $n_t = \sqrt{\frac{d}{f} S_t}$).[21] Then the total number of software products available in period t is $N_t = \sum_{s=0}^{t} n_t$.[22] A possible interpretation for this carry over is that the hardware product in each new period is an upgraded version of the previous period's hardware product which has backwards compatibility with the software produced for the old hardware. An example would be the Windows operating system, which features backward compatibility with all previous versions of the operating system, while each new version offers many new features. However, one could imagine alternative stories where a new generation of consumers would yield benefits arising from large networks of users existing before them.

The major implication of such a dynamic game is the following result concerning the time path of optimal copyright enforcement.

Proposition 1 *In the dynamic game explained above, the optimal time path of copyright enforcement is one of (weakly) increasing enforcement every period. That is,* $\alpha_t \leq \alpha_{t-1}$ *for all* t.

Proof. See Appendix. While the algebra of the proof is somewhat complicated, the intuition of the proof is very simple. In an earlier period, the monopolist faces a smaller installed base of software products, and thus a higher marginal network effect from allowing copying. Furthermore, as reduced copyright enforcement fuels the generation of more software goods, and these software goods persist over time, it is beneficial to the monopolist to build up the software base earlier. Thus, there is greater incentive to allow more copying (and thus impose weaker copyright enforcement) in earlier periods. ■

[21] Technically, I am now assuming that all new consumers each period buys one unit of new software, but still obtains utility from the existence of software from previous periods. This assumption is unnecessary for the propotion below, but greatly simplifies the analysis.

[22] Limited backwards compatability would imply, for example, $N_t = \sum_{s=0}^{t} \delta^{t-s} n_t$, with $\delta < 1$, but would not change the result concerning the optimal time path of copyright enforcement.

The proposition above demonstrates that as a good (or product-line) becomes more mature, the level of copyright enforcement chosen by the monopolist will increase. The logic here is clear: once goods have formed established networks, there is little gain to consumers from adding other consumers to the network. Another Microsoft analogy illustrates the thought process of a monopolist in such a dynamic game. When Windows 3.11 was released (and MS-DOS before it), Microsoft Windows and Macintosh OS were both relatively widely used. The addition of another consumer to the Windows 3.11 hardware network then had a relatively strong network externality on all other users. As the network grew, third-party software manufactures found it to be more worthwhile to focus on the Windows market, rather than the Macintosh market. However, by the time that Microsoft released Windows XP, almost all third-party software applications that are released are already focused on the Windows market, and so there is less gain in the form of additional software variety from adding another consumer to the Windows market, which is arguably as mature as possible.[23]

Thus one would expect that producers of new goods would desire low levels of copyright enforcement, while producers of goods that already have established networks would (ceteris paribus) desire high levels (or a full level) of copyright protection. This proposition is anecdotally verifiable. Again, returning to the Microsoft example, it is the case that over time, Microsoft has increased the level of copy protection on its operating systems. DOS and Windows 3.x, came on floppy disks and included a back-up program, Windows 95 jumped from floppy disks to CDs, which at the time were expensive to duplicate, Windows 98 required the entry of 12 digit CD-key to unlock the software, and Windows XP features a tight copy protection scheme that sends information about the computer on which it is installed to Microsoft, and then will not work for more than 30 days if it installed on a

[23] A report from OneStat.com (2005), a Dutch company that tracks operating system use through website access, finds that in 2002, over 97% of website visitors were using Microsoft Windows, compared to 1.5% for Apple Macintosh, and less than 1% for Linux.

computer that does not match the one in the Microsoft database. And plans for the next version of the Windows operating system include even stronger copy protection procedures. The proposition above explains the rationality of such a time path given that the marginal network externality for Microsoft OS's is (arguably) far lower now than it ever has been.

Here, though, is where the assumption of costless enforcement may begin to affect the analysis. A model without the externality-based benefits of copying, but with costly enforcement, could generate a similar dynamic implication if the cost of enforcement was changing over time. That is, it is possible that Microsoft's increased enforcement over the years is due in part (or completely) to a reduction in the costs of enforcing copyright protection, rather than a decrease in the gains from allowing copying. In fact, many of the anti-copying techniques employed by Microsoft today (such as the internet-based activation system in Windows XP) would have been impossible in the past, as web access was not widespread enough to make it feasible. Nevertheless, additional anti-copying security features were available at relatively little cost to Microsoft in the past, and yet they chose not to employ them, suggesting that their it is their dominance in the market that is playing the stronger role in their increase in enforcement, rather than cost-based changes, as it was common even in the mid-1980s for software to be encrypted to make copying the floppy disks difficult or even impossible. In fact, PRO-DOS, the version of DOS used by the Apple II, was copy-protected.

Interestingly, Apple which has lagged far behind Microsoft in the market for operating systems (which is admittedly strongly tied to the computer hardware market) no longer copy-protects its operating system software. In fact, in a New York Times review of the most recent version of the Macintosh OS named Tiger, which was released in April 2005, the lack of copy protection is discussed in a way that suggests that this is a "feature" of which users should take advantage (Pogue 2005): "you could make the case that Tiger is overpriced at $95... although it's worth noting that Mac OS X is not copy-protected and requires no Windows-style activation." Were the course of history different, and instead Apple had come to dominate the operating system market over the past 20 years, it is hard to imagine that Apple's operating system would be unprotected; it seems sure that Apple would have had progressively more and more copyright enforcement built into the software, as Microsoft has.

7 Conclusion

This paper shows that firms which produce goods which feature strong network externalities (either directly or through complimentarities with goods in another market) may benefit from relaxing copyright enforcement on unauthorized copying. This stems from the fact that the additional network size that is generated by the copying may increase consumers' willingnesses to pay enough as to counter-balance the loss in extractable surplus that arises from the availability of another option to consumers. By allowing copying to spread the good to low value consumers that the monopolist does not want to price to, the monopolist is able to benefit from a large network size without having to resort to a single low price.

Extending the static market into a simplified dynamic setting demonstrates that as the market for the product matures, the monopolist hardware producer will want to increase the level of copyright enforcement associated with her good. This is due to the fact that as the network gets larger over time, the marginal network externality decreases for these future generations of consumers. Because it is the benefits of a substantial network externality that cause the monopolist to be willing to allow copying, a smaller externality on the margin leads the monopolist to attempt to reduce the amount of copying among (potential) consumers of the good.

Future work could lead towards the analysis of a richer dynamic game with multi-generational consumers to see how the ability of copying to overcome the Coase conjecture would affect a firm's desire not to allow copying in mature networks. Similarly, a natural extension would be to examine the impacts of copying on the choice of consumers between competing technologies, such as Windows or Macintosh, or Windows and Linux. Moreover, an examination of the impacts of the monopolist's involvement in the third-party software market as well as introducing copying into the complementary market may well yield different implications about the role of copying. In the case where the monopolist plays a role in the complementary market, it may be that the monopolist prefers complete copying in the hardware market, if it results in large enough increases in demand in the complementary market. Such a finding would explain why Adobe, for example, distributes Acrobat Reader for free (this is equivalent to $\alpha = 0$ in the model presented here) and sells Acrobat Writer at a high price.

References

[1] Besen, S. M. and Kirby, S. (1989), “Private Copying, Appropriability, and Optimal Copyright Royalties” *Journal of Law and Economics*, vol. 32, pp. 255-280.

[2] Blackburn, D. (2004), “On-line Piracy and Recorded Music Sales” mimeo.

[3] Business Software Alliance (2005), “Second Annual BSA and IDC Global Software Piracy Study” website, ww.bsa.org/globalstudy/upload/2005-Global-Study-English.pdf. Downloaded September 25, 2005.

[4] Bykos, Y., Brynjolfsson, E. and Lichtman, D. (1999), “Shared Information Goods” *Journal of Law and Economics*, vol. 42, no. 1, pp. 117-155.

[5] Dinlersoz, E., and Pereira, P. (2004), “On the Diffusion of Electronic Commerce” mimeo.

[6] Economides, N. (1997), ”Network Externalities, Complementarities, and Invitations to Enter” *European Journal of Political Economy*, vol. 12, no.2, pp. 211-33.

[7] Katz, M. L. and Shapiro, C.(1985), “Network Externalities, Competition, and Compatibility” *American Economic Review*, vol. 75, no.3, pp. 424-440.

[8] Leibowitz, S. J. (1985), “Copying and Indirect Appropriability: Photocopying of Journals” *Journal of Political Economy*, vol. 93, no. 5, pp. 945-957.

[9] Oberholzer-Gee, F and K. Strumpf (2004), "The Effect of File Sharing on Record Sales: An Empirical Analysis." HBS mimeo.

[10] OneStat.com (2005), “Microsoft’s Windows OS global market share is more than 97% according to OneStat.com” website,

http://www.onestat.com/dutch/dutch_aboutus_pressbox10.html, downloaded April 20, 2005.

[11] Pogue, D. (2005), “Apple’s Tiger May Even Have PC Owners Longing for a Mac to Put It In” *The New York Times*, April 28, 2005, Section C, p. 1.

[12] Sargent, T. and Ljungqvist, L. (2000) *Recursive Macroeconomic Theory*, MIT Press.

[13] Shepard, A. (1987), ”Licensing to Enhance Demand for New Technologies” *Rand Journal of Economics*, vol. 18, no.3, pp. 360-368.

[14] Takeyama, L. (1994) “The Welfare Implications of Unauthorized Reproduction of Intellectual Property in the Presence of Demand Network Externalities” *Journal of Industrial Economics*, vol. 42, no. 2, pp. 155-166.

[15] Takeyama, L. (1997) “The Intertemporal Consequences of Unauthorized Reproduction of Intellectual Property” *Journal of Law and Economics*, vol. 40, pp. 511-522.

[16] Varian, H. (2000), ”Buying, Sharing, and Renting Information Goods” *Journal of Industrial Economics*, vol. 48, no. 4, pp. 473-88.

[17] Zentner, A. (2003), "Measuring the Effect of Online Music Piracy on Music Sales," University of Chicago Working Paper.

Appendix A Proof of Proposition 1

Proof. To show that $\alpha_t \leq \alpha_{t-1}$, for all t, look at the generic FOC that defines α_t. First, denote by π_t the static profits in period t and by Π_t the discounted sum of all remaining future profits. Thus:

$$\pi_t = \frac{1-a^t}{4}(\theta + v(S_t) + w(N_t))^2$$
$$\Pi_t = \sum_{s=0}^{T} \beta^s \pi_{t+s}$$

where β is a traditional discount factor. The FOC that defines α_t then is

$$\frac{\partial \Pi}{\partial \alpha_t} = \frac{\partial \pi_t}{\partial \alpha_t} + \sum_{s=1}^{T} \beta^s \frac{\partial \pi_{t+s}}{\partial N_{t+s}} \frac{\partial N_{t+s}}{\partial \alpha_t} = 0$$

because α_t only enters future profits through the increase in complementary products that it creates. The FOC is then:

$$\frac{\partial \Pi_t}{\partial \alpha_t} = (\theta + v(S_t) + w(N_t))^2 (\frac{-1}{4} + \frac{1-\alpha_t}{2} \frac{v'(S_t) + w'(N_t)}{2 - (1+\alpha_t)(v'(S_t) + w'(N_t))})$$
$$+ \sum_{s=0}^{T} \beta^s \sqrt{\frac{d}{f}} w'(N_{t+s}) \frac{(\theta + v(S_{t+s}) + w(N_{t+s}))^2}{1 - \frac{(1+\alpha_{t+s})}{2}(v'(S_{t+s}) + w'(N_{t+s}))} \frac{1 - a_{t+s}}{4} = 0$$

In order to see that $\alpha_t \leq \alpha_{t-1}$, compare the FOC for α_t with that for α_{t-1}, noting the changes and seeing how the changes affect the optimal choice of α. The FOC for α_{t-1} is, obviously, similar:

$$\frac{\partial \Pi_{t-1}}{\partial \alpha_{t-1}} = (\theta + v(S_{t-1}) + w(N_{t-1}))^2 (\frac{-1}{4} + \frac{1-\alpha_{t-1}}{2} \frac{v'(S_{t-1}) + w'(N_{t-1})}{2 - (1+\alpha_t)(v'(S_{t-1}) + w'(N_{t-1}))})$$
$$+ \sum_{s=0}^{T} \beta^s \sqrt{\frac{d}{f}} w'(N_{t-1+s}) \frac{(\theta + v(S_{t-1+s}) + w(N_{t-1+s}))^2}{2 - (1+\alpha_{t-1+s})(v'(S_{t-1+s}) + w'(N_{t-1+s}))} \frac{1 - a_{t-1+s}}{2} = 0$$

There are two sets of differences between these two FOC. First, note that $N_{k-1} = \sum_{j=0}^{k-1} n_j < N_k = \sum_{j=0}^{k} n_j$ for all k. Every occurrence of N_k in the FOC for α_t is replaced by the smaller N_{k-1} in the FOC for α_{t-1}. (The replacement of S_k with S_{k-1} changes nothing because S_k is determined solely by a static choice.) Second, note that the FOC for α_{t-1} has one more term in it, $\beta^{T-t+1}\sqrt{\frac{d}{f}}w'(N_{T-t+1})\frac{(\theta+v(S_{T-t+1})+w(N_{T-t+1}))^2}{2-(1+\alpha_{T-t+1})(v'(S_{T-t+1})+w'(N_{T-t+1}))}\frac{1-a_{T-t+1}}{2}$. I will show that both of these changes lead to a higher choice of α, and thus $\alpha_t \leq \alpha_{t-1}$.

I start by showing that $\frac{d\alpha_t}{dN_s} < 0$ for all s. I first construct the partial $\frac{\partial}{\partial N_t}(\frac{\partial \Pi_t}{\partial \alpha_t})$:

$$\frac{\partial}{\partial N_t}(\frac{\partial \Pi_t}{\partial \alpha_t}) = \frac{\partial(\frac{\partial \pi_t}{\partial \alpha_t})}{\partial N_t} + \sum_{s=1}^{T} \beta^s \frac{\partial(\frac{\partial \pi_{t+s}}{\partial N_{t+s}}\frac{\partial N_{t+s}}{\partial \alpha_t})}{\partial N_t}$$

The first term above is:

$$\begin{aligned}\frac{\partial(\frac{\partial \pi_t}{\partial \alpha_t})}{\partial N_t} &= \frac{1-\alpha_t}{2}\frac{(2-(1+\alpha_t)(v'(S_t)+w'(N_t))w''(N_t)+(1+\alpha_t)(v'(S_t)+w'(N_t))w''(N_t)}{(2-(1+\alpha_t)(v'(S_t)+w'(N_t))^2} \\ &= \frac{1-\alpha_t}{2}\frac{2w''(N_t)}{(2-(1+\alpha_t)(v'(S_t)+w'(N_t))^2} < 0\end{aligned}$$

because $w'' < 0$ and $v' + w' < 1$ by assumption and $0 \leq \alpha_t \leq 1$. The second term is:

$$\sum_{s=1}^{T} \beta^s \frac{\partial(\frac{\partial \pi_{t+s}}{\partial N_{t+s}}\frac{\partial N_{t+s}}{\partial \alpha_t})}{\partial N_t} = \sum_{s=1}^{T} \beta^s \sqrt{\frac{d}{f}}\frac{1-a_{t+s}}{2}\frac{\theta+v(S_{t+s})+w(N_{t+s})}{2-(1+\alpha_{t-1+s})(v'(S_{t-1+s})+w'(N_{t-1+s}))}\times$$
$$\left[w''(N_{t+s})(\theta+v(S_{t+s})+w(N_{t+s}))\left\{1+\frac{1+\alpha_{t+s}}{2-(1+\alpha_{t+s})(v'(S_{t+s})+w'(N_{t+s}))}\right\}+2w'(N_{t+s})^2\right]$$

The term outside the brackets is positive, from the assumption on v', w', and $v' + w'$. The term inside the brackets is negative. This is because the term is a quadratic function of $w'(N_{t+s})$. Evaluating the term in brackets at $w'(N_{t+s}) = 0$, it is negative and the entire term has no real valued roots. Thus, if it has no roots and is negative at one point, it must be everywhere negative. Thus both terms in $\frac{\partial}{\partial N_t}(\frac{\partial \Pi_t}{\partial \alpha_t})$ are negative, and therefore $\frac{\partial}{\partial N_t}(\frac{\partial \Pi_t}{\partial \alpha_t}) < 0$. It is then immediate that $\frac{\partial}{\partial N_s}(\frac{\partial \Pi_t}{\partial \alpha_t}) < 0$ as well, since $\frac{\partial}{\partial N_s}(\frac{\partial \Pi_t}{\partial \alpha_t})$ includes only the second term in $\frac{\partial}{\partial N_t}(\frac{\partial \Pi_t}{\partial \alpha_t})$.

Next construct the partial $\frac{\partial}{\partial \alpha_t}(\frac{\partial \Pi_t}{\partial \alpha_t})$:

$$\begin{aligned}
\frac{\partial}{\partial \alpha_t}(\frac{\partial \Pi_t}{\partial \alpha_t}) &= \frac{\partial}{\partial \alpha_t}(\frac{\partial \pi_t}{\partial \alpha_t}) + \sum_{s=1}^{T} \beta^s \frac{\partial}{\partial \alpha_t}(\frac{\partial \pi_{t+s}}{\partial N_{t+s}} \frac{\partial N_{t+s}}{\partial \alpha_t}) \\
&= \frac{-1}{2} \frac{v'(S_t) + w'(N_t)}{2 - (1+\alpha_t)(v'(S_t) + w'(N_t))} + \frac{1-\alpha_t}{2} \frac{(v'(S_t) + w'(N_t))^2}{(2 - (1+\alpha_t)(v'(S_t) + w'(N_t)))^2} \\
&+ \sum_{s=1}^{T} \beta^s \frac{\partial}{\partial N_{t+s}}(\frac{\partial \pi_{t+s}}{\partial N_{t+s}} \frac{\partial N_{t+s}}{\partial \alpha_t}) \frac{\partial N_{t+s}}{\partial \alpha_t} \\
&= \frac{1}{2} \frac{v'(S_t) + w'(N_t)}{2 - (1+\alpha_t)(v'(S_t) + w'(N_t))} \left[-1 + \frac{1-\alpha_t}{\frac{2}{v'(S_t)+w'(N_t)} - (1+\alpha_t)} \right] \\
&+ \sum_{s=1}^{T} \beta^s \frac{\partial}{\partial N_{t+s}}(\frac{\partial \pi_{t+s}}{\partial N_{t+s}} \frac{\partial N_{t+s}}{\partial \alpha_t}) \frac{\partial N_{t+s}}{\partial \alpha_t}
\end{aligned}$$

Notice that the term outside of the brackets is positive, as $v' > 0$, $w' > 0$, $v' + w' < 1$, and $0 \leq \alpha_t \leq 1$. Thus, the sign of the first term in $\frac{\partial}{\partial \alpha_t}(\frac{\partial \Pi_t}{\partial \alpha_t})$ depends on the sign of the term in brackets. However, $\frac{1-\alpha_t}{\frac{2}{v'(S_t)+w'(N_t)} - (1+\alpha_t)}$ is always less than 1, since $v' + w' < 1$ and $0 \leq \alpha_t \leq 1$. Therefore, the term in brackets is negative, and thus the first part of the sum is negative. The second term in the sum is:

$$\sum_{s=1}^{T} \beta^s \frac{\partial}{\partial N_{t+s}}(\frac{\partial \pi_{t+s}}{\partial N_{t+s}} \frac{\partial N_{t+s}}{\partial \alpha_t}) \frac{\partial N_{t+s}}{\partial \alpha_t} = \sum_{s=1}^{T} \beta^s \frac{\partial}{\partial N_t}(\frac{\partial \pi_{t+s}}{\partial N_{t+s}} \frac{\partial N_{t+s}}{\partial \alpha_t}) \frac{\partial N_t}{\partial \alpha_t}$$

because $N_{t+s} = N_t + \sum_{j=1}^{s} n_{t+j}$. I have already shown that $\frac{\partial}{\partial N_t}(\frac{\partial \pi_{t+s}}{\partial N_{t+s}} \frac{\partial N_{t+s}}{\partial \alpha_t}) < 0$, so I need only to focus on $\frac{\partial N_t}{\partial \alpha_t}$. However, this term again has already been shown to be positive in the main text. Thus the second term, $\sum_{s=1}^{T} \beta^s \frac{\partial}{\partial N_{t+s}}(\frac{\partial \pi_{t+s}}{\partial N_{t+s}} \frac{\partial N_{t+s}}{\partial \alpha_t}) \frac{\partial N_{t+s}}{\partial \alpha_t}$, is negative. Therefore $\frac{\partial}{\partial \alpha_t}(\frac{\partial \Pi_t}{\partial \alpha_t}) < 0$ as it is the sum of two negative terms.

Thus, $\frac{d\alpha_t}{dN_s}$ is negative:

$$\frac{d\alpha_t}{dN_s} = -\frac{\frac{\partial}{\partial N_s}(\frac{\partial \Pi_t}{\partial \alpha_t})}{\frac{\partial}{\partial \alpha_t}(\frac{\partial \Pi_t}{\partial \alpha_t})} < 0$$

because both the term in the numerator and the term in the denominator are negative. The replacement of every occurrence of N_k in the FOC for α_t by the smaller N_{k-1} in the FOC for α_{t-1} would imply that $\alpha_t \leq \alpha_{t-1}$. (The equality is due to the fact that if $\alpha_{t-1} = 0$, α_t can not be reduced.) However, it still remains to check that the addition of the extra term in the FOC would also serve to decrease the optimal choice of α.

This is comparatively simple. Imagine that the FOC for α_t includes the extra term

$$\phi \left(\beta^{T-t+1} \sqrt{\frac{d}{f}} w'(N_{T-t+1}) \frac{(\theta + v(S_{T-t+1}) + w(N_{T-t+1}))^2}{2 - (1 + \alpha_{T-t+1})(v'(S_{T-t+1}) + w'(N_{T-t+1}))} \frac{1 - a_{T-t+1}}{2} \right)$$

but $\phi = 0$ in the FOC for α_t and $\phi = 1$ in the FOC for α_{t-1}. Then it is only necessary to check the sign of $\frac{d\alpha_t}{d\phi}$. I have already determined that $\frac{\partial}{\partial \alpha_t}(\frac{\partial \Pi_t}{\partial \alpha_t}) < 0$ and $\frac{\partial}{\partial \phi}(\frac{\partial \Pi_t}{\partial \alpha_t}) > 0$ is immediate. Thus,

$$\frac{d\alpha_t}{d\phi} = -\frac{\frac{\partial}{\partial \phi}(\frac{\partial \Pi_t}{\partial \alpha_t})}{\frac{\partial}{\partial \alpha_t}(\frac{\partial \Pi_t}{\partial \alpha_t})} > 0$$

The changes in the FOC for α_{t-1} relative to the FOC for α_t can be summarized as simply a decrease in N_s for all s as well as a change from $\phi = 0$ to $\phi = 1$. Since the decreases in N_s serve to increase the optimal α, and the increase in ϕ serves to increase the optimal α as well, it follows that $\alpha_t \leq \alpha_{t-1}$. QED. ■