

PERFECTLY COMPETITIVE INNOVATION

MICHELE BOLDRIN AND DAVID K. LEVINE

ABSTRACT. Based partially on the belief that innovation is not possible under perfect competition, many thousands papers have been written about the nature of innovation under monopoly or oligopoly. In fact, competitive rents can and do sustain innovation in the complete absence of monopoly power. However, little is known about the source and significance of these rents. We begin to remedy this imbalance by examining the way in which competitive innovators earn rents both in theory and in practice.

Keywords: Economic Theory, Innovation, Perfect Competition, Intellectual Property

JEL Classification: A1; A2

Date: First Version: October 3, 1997, This Version: 28th August 2005.

Many ideas presented here were contained in a previous, unpublished, paper titled “Growth Under Perfect Competition.” This paper is based upon work supported by the National Science Foundation under Grants SES 01-14147, and 03-14713. Boldrin also acknowledges research support from the Spanish BEC2002-04294-C02-01 and the University of Minnesota Grants in Aid Program for financial support. We benefited from comments from seminar participants at Toulouse, the London School of Economics, Humboldt University, UC Berkeley, Cornell, Chicago, Wisconsin-Madison, Iowa State, New York University, Stanford, Univ. of Pennsylvania, Columbia, Oxford, CEMFI, Carlos III, Rochester, DELTA-ENS Paris, Venice, Padova, IGIER-Bocconi, and the University of Minnesota. Andrew Postlewaite and Jim Schmitz also made a number of valuable suggestions.

Corresponding Author: Michele Boldrin, Department of Economics, University of Minnesota, Minneapolis, MN 55455, USA. Phone: 612-624 4551. Email: mboldrin@econ.umn.edu.

1. INTRODUCTION

It would be efficient ex post to make the existing discoveries freely available to all producers, but this practice fails to provide the ex ante incentives for further inventions. A tradeoff arises between restrictions on the use of existing ideas and the rewards to inventive activity. Robert Barro and Xavier Sala-i-Martin (1999)

Evidently, we must depart from the common practice in neoclassical growth theory of assuming that all firms act as price takers in an environment of perfect competition. Firms must be able to sell their products at prices in excess of unit production costs if they are to recover their up-front outlays on research and development. In other words, some imperfect competition is necessary to support private investments in new technologies. Gene Grossman and Elhanan Helpman (1994)

The idea that monopoly is the necessary handmaiden of innovation is traditional and widespread in economics. It was the dominant theme of Schumpeter [1942] famous book. We find it formalized in the work of Arrow [1962] and subsequently developed by Shell [1966, 1967]. The theme re-emerges as central to the “new growth” theory of Romer [1990], and plays a leading role in the work of authors such as Grossman and Helpman [1991] and Aghion and Howitt [1992]. Indeed, this idea is such a strong part of conventional wisdom that most textbooks speak in terms similar to the passage from the Barro and Sala-i-Martin textbook quoted above.

That an idea is widely believed and quoted does not, however, make it true. The claim that monopoly is necessary for innovation is not correct either as a matter of theory, or as a matter of fact. Indeed, economists such as Stigler [1956] seem to explicitly reject the idea, and recent authors such as Irmen and Hellwig [2001] or Boldrin and Levine [2002] question it as well. However, while many thousands of papers have been written exploring innovation under monopoly, very few have examined the incentives to innovate “assuming that all firms act as price takers in an environment of perfect competition.” This paper is an effort to fill that gap. Understanding innovation under competition is significant for two reasons. First, innovation under competition is more pervasive than innovation under monopoly and cannot be understood in the context of a model that rules it out as a matter of principle: what your model cannot see, you cannot measure. Second, the socio-economic relevance of intellectual property legislation is becoming every day more apparent, and an intelligent and informed debate about

policy is impossible if the consequences of eliminating the current monopolistic system of intellectual property cannot be conceived.

The theory of innovation is a theory of the costly creation of new factors of production, and of their adoption. That competitive equilibrium results in prices equal to marginal costs is as widely believed - even by economists - as it is untrue. It is true only in the absence of capacity constraints. Since, in reality, only a limited number of units of a new factor of production are available, as is painfully evident to students of general equilibrium analysis, fixed factors earn competitive rents. They can and do compensate innovators for the cost of their innovations.

Some authors appear to believe, and some currently argue, that capacity constraints are irrelevant to the theory of innovation because ideas flow costlessly and instantaneously once they are created. The irony of professors, whose substantial salary derives entirely from the difficulty and time consuming nature of communicating old ideas, arguing that ideas are freely and costlessly transmitted, should not escape us. But however incredible this might be as a general principle, certainly it is true that the development of information technology has greatly accelerated the transmission of copies of ideas in electronic form. If the price at which the creator can sell under competition is p , rapid reproduction results in k available copies, and the cost of creating the innovation is C , then the return on the innovation is $R = pk/C$. If k is very large, then p must be small, and so the conventional wisdom holds, certainly $R < 1$, meaning it is unprofitable to innovate. This argument is problematic: it may well be true that p is close to zero, but it is only close to zero because k is large, and there is no theorem that a small number times a large one is small. Indeed pk is simply revenue, and the behavior of revenue for large k depends on the elasticity of demand. Moreover, an often overlooked fact is that the same technological change that results in the increase in k may result also in a massive decrease in C . Revenue in the music industry has fallen less than 20% due to piracy, including rapid electronic reproduction.¹ On the other hand, computer technology has reduced C by over an order of magnitude.² When there are many small and large numbers around, we should not be too quick to round off selected ones to zero.

¹Hui and P'ng [2003] have the most detailed analysis of the impact of piracy on revenues, although it lacks estimates of the decreased revenue due to the legitimate seller setting a lower price. Leibowitz [2002] indicates also a relatively small effect.

²>From 1992-2002, estimates in Karr [2002] are that the cost for a sound studio has fallen from \$50,000 to \$1,000. Since the 1970s the fall in cost has been far more dramatic.

We also consider the impact that unpriced spillover externalities have on the incentive to innovate under competition. The key point is that competitive rents are reduced but positive. This highlights the issue of appropriability. Competitive equilibrium enables innovators to appropriate some of the social surplus from their innovation, but not all of it.³ Ideas that generate a great deal of social surplus will be invented regardless of whether or not there is monopoly. On the other hand, some marginally socially valuable ideas will not be produced under perfect competition. It is worth noticing, though, that competition over marginal ideas is not likely to be fierce and, to the extent that loss of appropriability is due to the imitative effort of competitors, appropriability for marginal ideas may be quite high. The welfare tradeoff is complex, since policies that increase appropriability increase other distortions as well. We do not explicitly consider welfare tradeoffs here, but when the point of departure is a model in which perfect competition allows no appropriation, a correct analysis of welfare becomes impossible. We propose a model in which welfare considerations become meaningful.

Our goal in this paper is to examine the main determinants of the competitive rents that sustain innovation in the absence of monopoly power. We do this both through theory, and through case studies. As a matter of practice, it is not easy to find markets in which monopoly power and other frictions are completely absent. One market that does closely fit the theory of perfectly competitive equilibrium is the market for open source software. Through the voluntary renunciation of both copyright and trade-secrecy, software products released under open source licenses may be freely copied and resold under conditions as close to perfect competition as we are likely to find in any market. The very existence of such a market is a challenge to the conventional idea that there can be no innovation without monopoly.

So let us grab our mallet, enter the glass house of conventional wisdom, and shatter a few illusions.

2. INNOVATION UNDER COMPETITION

We consider the situation of an innovator or creator who is contemplating a fixed cost of C to create her new idea. Once the fixed cost is incurred, the innovator will have produced a single initial copy of a new good - the first copy of the new novel, the first prototype of the new machine, or more broadly the template by which copies may be created. Since our interest is in perfect competition, we are going to assume that from the moment in which the innovator sells any fraction of this first copy, she is in immediate

³Monopoly generally allows a greater amount of appropriability, but still less than full appropriability.

and perfect competition with the purchaser(s) of this idea. To focus things still more, we will analyze the case in which every copy of the idea contains a template for the idea that is a perfect substitute for the original idea. That is, any purchaser of the idea can make copies of the idea using exactly the same technology for producing copies that is available to the original innovator. Indeed, to get the ball rolling, we will make the extreme assumption that ideas are so freely reproducible that there is no need to invest resources in either reverse engineering or distributing them - simply, in the process of using ideas, more copies are created that may be freely sold.

Specifically we suppose that for the initial cost of C a single initial copy is made available in period $t = 0$. In all periods t , if k_t copies are available then those copies are consumed, yielding a utility of $u(k_t)$ and at the same time reproduce themselves and additional copies, so that βk_t copies will be available at time $t + 1$ where $\beta > 1$... Quah [2002] refers to this as the 24/7 case, meaning that copies are being reproduced 24/7 regardless of what other uses they are put to. As for ownership and possession, we assume that each owner of a copy is endowed the following period with β copies. This model can be thought of as a model of music in the form of MP3 electronic files distributed on a peer-to-peer network. One initial copy is produced by actually writing and recording the music at a cost of C . This is used to “seed” the peer-to-peer network by making the initial copy available. Subsequently, users download a copy of the MP3 file, and - while listening to it - make their copy available to other users of the network for further downloading. As time goes on, the number of individuals who have copies of the file grows exponentially. We assume also a common discount factor of $0 \leq \delta < 1$.

This model is straightforward to analyze. In the first period and subsequently, the price of copies is proportional to marginal utility. Since the initial template is the only fixed factor, there is constant returns to scale, and there is competition to produce subsequent copies, all competitive rents accrue to the innovator. The number of copies in circulation naturally grows exponentially: $k_t = \beta^t$. The innovator’s rents are easily computed to be the discounted present value of the revenue stream from renting copies

$$q_0 = \sum_{t=0}^{\infty} (\beta\delta)^t u'(\beta^t).$$

Unless the single initial copy satiates the market so that $u'(1) = 0$ this competitive rent is always positive, and the innovator will be willing to innovate provided that $R = q_0/C \geq 1$.

Our goal is to understand the determinants of q_0 both as a matter of theory and practice.

2.1. Improved Reproduction Technology. Let us examine first what happens as β , the rate at which copies can be made, increases. This seems to approach the conventional case in which prices fall quickly to zero, and competition is alleged to fail to produce innovations. This conventional wisdom fails for two reasons: first, it ignores the initial period. During this initial period, no matter how good the reproduction technology, only the initial copy is available. In other words, $q_0 \geq u'(1)$ regardless of β . The amount that will be paid for the initial copy will never fall to zero, no matter how many copies will be available in the immediate future. If, however, the same technological improvement that leads to increased β also reduces the time it takes to make additional copies - that is, periods get shorter increasing the discount factor δ - then $u'(1)$ does effectively fall to zero.

Conventional wisdom also fails for a less apparent, reason: increasing β may increase, rather than decrease competitive rents. We compute

$$\frac{dq_0}{d\beta} = (1/\beta) \sum_{t=0}^{\infty} t(\beta\delta)^t u'(\beta^t) \left[1 + \frac{\beta^t u''(\beta^t)}{u'(\beta^t)} \right].$$

This of course can be either positive or negative, depending on whether demand elasticity $-u''(k)k/u'(k)$ is greater or smaller than one. If demand is elastic, then naturally as the reproduction rate increases, competitive rents become infinite. Notice that this is true regardless of the lead time for making copies - that is, with elastic demand revenue goes to infinity as $\beta \rightarrow \infty$ even if $\delta \rightarrow 1$. If, on the other hand, demand is inelastic, prices may fall to zero sufficiently quickly that revenue falls to zero. This latter case supports the conventional conclusion - but notice that it relies on the assumption that demand elasticity for new goods is less than one.⁴ If marginal cost is zero, there is no capacity constraint and demand is elastic throughout its range in a static model, then the monopolist would choose to produce arbitrarily large amounts, resulting in an arbitrarily small price. Since we do not observe monopolists doing this, we might be tempted to infer, as some have done, that demand cannot be elastic throughout its range. However, since marginal cost is not zero, there is a capacity constraint, and the world is not static, this inference is not justified.

2.2. General Reproduction Technology. The 24/7 production technology is obviously very special - no other inputs are used other than copies, and there is no tradeoff between consuming and making copies. Introducing a more general technology for reproduction, however, does not change the basic picture. There will still be a positive competitive rent q_0 accruing to the scarce initial fixed factor, and the innovation will still be introduced if

⁴Note that if price actually falls to zero for some finite k then demand is necessarily inelastic near that k .

$R = q_0/C \geq 1$. Although a more general technology makes computation of q_0 more complex, it provides us with interesting insights into the factors affecting competitive rents.

Suppose that the representative consumer's utility function is $\sum_{t=0}^{\infty} \delta^t [u(c_t) - wL_t]$, where c_t is the flow of consumption services from copies, and L_t is labor, supplied at the constant wage w . We assume that copies depreciate at a fixed rate so that, without additional reproduction, ζk_t units are available tomorrow; we allow $\zeta > 1$ to include the 24/7 case.

Copies and labor may be used either in the consumption or in the copying sector. Let k_t^c, ℓ_t^c be the inputs employed in the first and k_t^k, ℓ_t^k those employed in the second sector; both have neoclassical production functions. Consumption services are $c_t = F(k_t^c, \ell_t^c)$ and additional copies are $x_t = G(k_t^k, \ell_t^k)$. The stock of copies evolves according to $k_{t+1} = \zeta k_t + x_t$, and as before we assume perfect competition from the instant of first sale of the initial copy $k_0 = 1$.

In each period $t = 0, 1, \dots$, equilibrium solves two maximization problems. First, given the stock k_t , the labor supply L_t , and the planned reproduction x_t , inputs are allocated to sectors in order to

$$\max_{0 \leq k_t^c \leq k_t, 0 \leq \ell_t^c \leq L_t} c_t = F(k_t^c, \ell_t^c)$$

subject to

$$0 \leq x_t \leq G(k_t - k_t^c, L_t - \ell_t^c),$$

The solution to this yields a production possibility frontier

$$c_t = T(k_t, x_t, L_t).$$

Under standard regularity assumptions on F and G the production possibility frontier T is increasing in k_t and L_t and decreasing in x_t and is concave. Define also $\bar{x}(k_t) = \lim_{L \rightarrow \infty} G(k_t, L)$ - this represent the greatest amount of new capital that can be produced from a given starting capital stock.⁵

Second, the period labor supply L_t solves

$$\max_{L_t} u[T(k_t, x_t, L_t)] - wL_t,$$

which has a unique solution $L_t = L(k_t, x_t)$, for given w . Notice that, from the first order condition

$$u'[T(k_t, x_t, L_t)] \frac{\partial T}{\partial L}(k_t, x_t, L_t) = w,$$

the relation between L_t and either k_t or x_t is ambiguous, as different factor intensity rankings, and the possibility of factor intensity reversal, may force $L(k_t, x_t)$ to depend on k_t and x_t in non-monotone form. We rule out these

⁵We allow the possibility that $\bar{x}(k_t) = \infty$; if G is strictly increasing then this must be the case.

complicated, and, for our purposes, altogether irrelevant cases, and assume that L is sufficiently well behaved that $T(k_t, x_t, L(k_t, x_t))$ is increasing in k_t , decreasing in x_t , and has a non-negative cross-partial second derivative. Regardless of factor intensities it is always concave. Define the period return function $V(k_t, x_t) = u[T(k_t, x_t, L(k_t, x_t))]$; this then is also increasing in k_t , decreasing in x_t , and strictly concave in both arguments.

Given $k_0 = 1$, the intertemporal competitive equilibrium of this economy is summarized by sequences $\{k_t\}_{t=0}^{\infty}$ solving

$$v(k_0) = \max_{\{k_t\}_{t=1}^{\infty}} \sum_{t=0}^{\infty} \delta^t V(k_t, k_{t+1} - \zeta k_t)$$

subject to

$$\zeta k_t + \bar{x}(k_t) \geq k_{t+1} \geq \zeta k_t.$$

Notice immediately that $q_0 = v'(k_0) > 0$ under exactly the same conditions as in the simple model, that is: as long as consumers are not satiated by the amount of the durable good made available by the innovator in the first period.

Consider first the special case in which labor is not used as an input, so that $c_t = k_t^c$ and $x_t = \beta k_t^k$. The Bellman equation for this optimization problem is

$$v(k) = \max_{0 \leq c \leq k} \{u(c) + \delta v(\beta k - (\beta - \zeta)c)\}.$$

When this problem is decentralized as a competitive equilibrium, the price of consumption services in period t is given by $p_t = u'(c_t)$ From the resource constraint

$$c_t = \frac{\beta k_t - k_{t+1}}{\beta - \zeta}.$$

If ζ is large enough relative to β it may be optimal not to invest at all and to reproduce solely by consuming, as for example in the 24/7 case. When the optimum involves a strictly positive savings rate, by standard dynamic programming arguments, the price q_t of copies k_t can be computed as

$$q_t = v'(k_t) = p_t \frac{\beta}{\beta - \zeta}.$$

As $p_t > 0$, $q_t > 0$ for all t . The zero profit condition here implies that q_t decreases at a rate of $1/\beta$ per period of time.

Next, let us assume a 24/7 production technology, but one in which labor is required along with copies to produce consumption. That is, $c_t = F(k_t, L_t)$ and $k_t = \beta^l$. Consider first the case in which copies and labor must be used in fixed proportion - that is $c_t = \min\{k_t, L_t\}$ - similar to, say, the Netflix technology. DVDs are kept at a central location, every period they are delivered to users who rent them for one period then return them to the central

location. Here, labor is needed for delivering and recovering the DVDs. In this case, the optimum is obviously to allow $c_t = k_t = L_t$ until some threshold c^* is passed, then keep c_t constant as k_t continues to (costlessly and irrelevantly) grow. Utility is $u(c_t) - wc_t$, so the threshold is simply where $u'(c^*) = w$. Since the competitive rent per unit of consumption services in period t is the marginal social value $u'(c_t) - w$, it decreases as productive capacity accumulates period after period and, at the threshold, it falls to zero. After that period no further rents are earned by the innovator. This situation in which unit competitive rents fall to zero in a finite and possibly short period of time is the closest one gets to conventional wisdom.

By way of contrast, sticking with the 24/7 technology, suppose there is a Cobb-Douglas production function for “delivering” consumption $c_t = (k_t)^\alpha (L_t)^{1-\alpha}$. Here labor and copies are substitutes. For example, one might have many copies of a recording, one in your bedroom, one in your living room, one in each car and so forth. Alternatively, you could get by with a single copy, but at the expense of substituting the labor of moving the copy from one location to another.⁶ Suppose also that the utility function has the complementary CES/constant relative risk aversion form $u(c_t) = c_t^\gamma$ for $1 > \gamma > 0$. Then the period utility is $(k_t)^{\gamma\alpha} (L_t)^{(1-\alpha)\gamma} - wL_t$. The first order condition for the optimum labor supply is

$$(1 - \alpha)\gamma(k_t)^{\gamma\alpha} (L_t)^{(1-\alpha)\gamma-1} = w.$$

This can be substituted back into the utility function to find utility as a function of capacity (stock of copies) and the wage. We may then easily differentiate this to find the rental price of capital $p_t = du/dk_t$ and the revenue earned in each period by the innovator

$$p_t k_t = \frac{\alpha\gamma}{1 - (1 - \alpha)\gamma} \left[\frac{(1 - \alpha)\gamma}{w} \right]^{\frac{(1-\alpha)\gamma}{1-(1-\alpha)\gamma}} k_t^{\frac{\alpha\gamma}{1-(1-\alpha)\gamma}}.$$

This has the property that as $k_t \rightarrow \infty$ the per period revenue becomes infinite. In particular in this case, as the reproduction technology improves and $\beta \rightarrow \infty$ the present value competitive rent accruing to the innovator becomes infinite - the opposite of the conventional case.

2.3. Spillover Externalities. In the literature on innovation, the idea that there are spillover externalities is widespread. That cheap imitation is possible is undeniable - and indeed our entire model so far has been based upon cheap imitation. The best example we know of is the invention, by Travelpro, of the modern wheeled roll-on suitcase with a retractable handle.⁷

⁶The labor in this example is less likely to be the effort of physically moving the disk than the need to remember to do so.

⁷We are grateful to Ivan P'ng for bringing this to our attention.

Obviously such an idea cannot be both useful and secret - and once you see a wheeled roll-on suitcase rolling across the airport terminal it is not difficult to figure out how to make one of your own.

As an empirical matter small unpriced externalities are certainly widespread. The empirical issue of how strong and quantitatively relevant they are in markets for innovations has scarcely been addressed. Our best guess as to what the literature on spillover externalities has in mind is either the reverse engineering of existing products or the spread of knowledge embodied in employees who move from one firm to another. Neither of these, however, are unpriced externalities. In the case of reverse engineering, it is generally necessary to buy something to reverse engineer it. As to the movement of employees, we agree with Becker [1971]

Firms introducing innovations are alleged to be forced to share their knowledge with competitors through the bidding away of employees who are privy to their secrets. This may well be a common practice, but if employees benefit from access to salable information about secrets, they would be willing to work more cheaply than otherwise.

Regardless of their empirical significance, what is the consequence of unpriced spillover externalities on competitive rents? In the 24/7 model, each copy produces β copies in the following period. We have assumed that these copies belong to the owner of the copy from which the duplicates are made. With an unpriced spillover externality, some lucky individuals will instead get copies for free. In other words, we may take $\beta = \beta_0 + \beta_S$ where β_0 are the number of copies that wind up in the hands of the original owner, and β_S are the number of copies that wind up in the hands of fortunate passers-by. When $\beta_S = 0$ all copies are the fruit of the original copy, and as is standard in competitive equilibrium, all competitive rents go to the original innovator. With the spillover externality, the value of the fruits that are given away involuntarily accrue to the lucky beneficiary, and none of the fruits that stem from the free copy accrue to the original innovator. In other words, the competitive rent accruing to the innovator is only

$$q_0 = \sum_{t=0}^{\infty} (\beta_0 \delta)^t u'(\beta^t).$$

Price is driven by the total number of copies β^t , but the innovator collects rents only on those copies β_0^t that have not escaped his control. Unless $\beta_0 = 0$, so that all copies have to be given away to passers-by, this competitive rent is still positive. This result is not terribly surprising: theft after all is a real phenomenon, and yet markets continue to function even though a fraction of goods and services are not sold but are stolen instead.

We should note that when the 24/7 model is extended to allow for costly reproduction of copies, the presence of spillover externalities does have an important consequence. Without them, conditional on innovation taking place, there is no distortion and the welfare theorems apply. In the presence of spillover externalities - just as with the presence of theft - fewer copies will be reproduced than in the first-best.

2.4. Complementary Sales. Our model is one of perfect competition *post* innovation. Prior to innovating, no innovator can be a price taker, as her innovation will certainly have a non-trivial impact on relevant prices. For example, an innovation that lowers the cost of making cars will certainly have an impact on price of cars; writing a new novel will certainly have an effect on the price at which that novel can be sold. Creation of a new idea or fixed factor will generally affect prices in other markets as well. The standard case of perfect competition assumes perfect divisibility - that the initial unit may be produced in arbitrarily small quantities. In the theory of innovation, we have dropped that assumption - recognizing that two first halves of a book are a poor substitute for the whole. With perfect divisibility and perfect competition, we can ignore the effect that production has on the prices of substitutes and complements, and who owns these other factors does not matter. In the case of innovation, where a new good is produced in a discrete amount, we can no longer safely do so. For example, writing a new novel will have a significant effect on the demand for the author's services on the lecture circuit; writing a new song will have a significant effect on the demand for the singer's live performances; creating a new software package will have a significant effect on the demand for the author as a consultant, and so forth and so on.

For simplicity, we assume there is a single other commodity, the fixed quantity of which we denote by a . We now write the utility of the representative consumer by $u(c_t, a)$. Conditional on the innovation taking place, our analysis of the time path of c_t is not affected by the presence of a . What we are interested in is the net change in the price of a due to innovation, assuming perfect competition post innovation. The answer is simply

$$Q = \sum_{t=0}^{\infty} \delta^t [u_a(c_t, a) - u_a(0, a)] = \sum_{t=0}^{\infty} \delta^t \int_0^{c_t} u_{ac}(c, a) dc.$$

>From this we see that if a is a substitute for c_t , so that the cross partial is negative, then innovation causes the price of a to fall, while conversely if a is a complement of c_t , so that the cross partial is positive, then the innovation causes the price of a to rise.

Consider first the case of exogenous ownership. If the innovator does not own (any) a then Q does not matter to the process of innovation. If

the innovator owns (some portion of) a then the incentive to innovate is decreased/increased as the complementary good is a substitute/complement. For example, the owner of a car factory will have a reduced incentive to create a new type of car, because this will lower the value of his existing cars, while the writer of a song will have an increased incentive to create a new song, because this will raise the value of his live performances. The case of recorded versus theatrical performances is an interesting case. Overall, the behavior of the existing industry - carefully avoiding DVD releases until after theatrical performances are concluded - seems to suggest that theatrical performances are a substitute for the DVD. Examination of practices in the music industry suggest the opposite is true for music. Notice however, that even under perfect competition, there is no obligation to release a recordable version of a product, so even without government intervention the movie industry would still be free to release DVDs only after the first theatrical run was complete.⁸

The case of complements is particularly important, since in practice it provides a significant source of competitive rents. We refer to the sale of a complementary product by the innovator as a *complementary sale*. Notice that in the case of complementary sales, raising the efficiency of reproduction (increasing β) always increases Q . In particular, with complementary sales, it is perfectly possible to have innovation under perfect competition even when the reproduced good is given away for free. For most part of the history of the radio and television industries this was exactly the case: the product was given away for free, and profit - substantial profit - came from the complementary sale of advertising.

Ownership, of course, is not exogenous. The owner of a car factory can sell his factory; even such things as the rights to revenue from the live performance of music can be traded in asset markets. This leads to a key point, first pointed out by Hirshleifer [1971]: an innovator has a substantial first mover advantage from his ability to trade in markets on the basis of his inside information about his innovation. The owner of a car factory who invents a better car can sell the factory prior to announcing his invention, for example. Moreover, in asset markets, short-sales may be possible as well - and in the extreme case considered by Hirshleifer in which the innovator is a price-taker in asset markets, she can generate essentially infinite profit through inside knowledge that prices are about to change by a tiny amount. In the Hirshleifer account, the private value of innovation under perfect competition is generally much greater than the social value. In other words if innovators and creators are small enough players in financial

⁸The theft of master copies to make reproductions (or for any other reason) is covered under ordinary laws concerning theft, independent of any laws on intellectual property.

markets and liquidity constraints are not too important, then under perfect competition there will be too much rather than too little innovation.

3. CASE STUDIES

In practice there are few markets that operate without some degree of monopoly power. We examine four case studies in which monopoly power is weak. Obviously this is not an exhaustive list, nor do these markets correspond in all respects to the idealized perfect competition of our model. They do signal that many examples exist of industries in which innovation thrives where there is fierce competition between innovators and imitators, and that competitive rents play an important role in sustaining the incentive to innovate.

3.1. Open Source Software. As we indicated in the introduction, the market for open source software is the most striking example of competitive innovation. Open source software is characterized by the voluntary renunciation of copyright and patent - buyers are entitled to make their own copies, modified or not, and sell them. In addition there is a voluntary renunciation of trade-secrecy as the original creator publishes the source code - the “blueprint” for writing the software - along with the software itself, and buyers are also entitled to make copies modified or not and sell them.⁹

Since the conditions of the market are essentially those of perfect competition, two questions arise. First, what exactly are the competitive rents in this market? Second, is the market itself significant - or, as it is sometimes alleged, does the market simply free-ride off the proprietary market making cheap imitations of software that never would have been produced in the first place absent monopoly power?

First, the source of competitive rents appears to be largely through the complementary sale of expertise. That is, the actual duplication of copies is sufficiently quick and conditions of demand are such that only small rents seem to be obtainable through the actual sale of copies. It is true that, historically, physical copies have been sold for greater than marginal cost. Red Hat is a company that, at one time, sold a distribution of Linux - a modified and customized Linux system with many features, which can be easily installed. Although the underlying Linux system is obtained by Red Hat for free, the customization and testing conducted by Red Hat is costly. Using prices quoted on the Internet on July 10, 2002, Red Hat charged \$59.95 for a package containing its system. Because it is based on the underlying GNU/Linux system, competitors can legally duplicate and sell the exact

⁹Some open source software also requires that as a condition of use, buyers make their modifications available under the same terms.

same “Red Hat” system. In fact, at least two two companies, Hcidesign and Linuxemporium, did exactly this. On July 10, 2002, Hcidesign offered for sale Red Hat Linux 7.2 for a price of \$16.00, about 1/3rd of the price charged by Red Hat. Linuxemporium.co.uk offered a similar deal. The striking fact being that Red Hat sold many more \$59.95 packages than Hcidesign and Linuxemporium did with \$16.00 packages and never represented a dangerous market threat to Red Hat.

However, despite their ability to earn revenue from this sale of expertise, Red Hat eventually concluded that they were not selling enough \$59.95 copies and switched to a different revenue model. What had previously been sold is now given away for free as “Fedora Core” and is used as a platform to get feedback on features that are incorporated into the commercial system called “Red Hat Enterprise Linux” which is available only by annual subscription at a price that - depending on features - ranged on August 24, 2005 from \$349 to \$2499. The following blurb from Red Hat promotional material makes it clear what it is that is being paid for

Unlimited access to service and support: Subscriptions include ongoing service and support to guarantee your systems remain secure, reliable, and up-to-date. When you have a technical question, you’ll speak to Red Hat Certified Software Engineers. Or you can access a self-serve knowledge base of technical information and updates.

What does this offer that imitators cannot? If you have a problem with software, would you prefer to consult with the people who wrote it or the people who copied it?

Similar observations are made by Lerner and Tirole [2004] about the financial benefit to individual developers of contributing to open source projects. For example, Apache is the leading webserver on the internet, holding a greater than 65% market share.¹⁰ The team of programmers that develop Apache are ranked according to the significance of their contributions, and hold other jobs. Work by Hann et al [2004] shows that the salaries they receive in these other jobs is heavily influenced by their rank within the Apache organization. In other words, the “expertise” model is much like that in academia - the software writers write software in order to receive recognition and financial payment for the expertise they demonstrate through their published product. While there is no doubt that some contributions to open source are due to altruism, it is equally certain that this massive industry is in fact largely financed through competitive rents.

¹⁰See http://news.netcraft.com/archives/web_server_survey.html for current statistics on webserver market share.

>From the perspective of the theory, the relevant competitive rents are generated by the complementary sales of expertise rather than the direct sales of boxed or downloaded software. Notice that the stock of expertise unfolds over time much as in the two-sector model - originally the writer/creator of the software is the only one with expertise. But as time goes on, others, such as Red Hat Certified Software Engineers are taught the knowledge, and the stock of expertise expands, and the price at which it can be sold drops. Of course, in the meanwhile new innovations are created, and new expertise is generated.

Finally, we turn to the question of just how massive and successful the open source industry is. It could be that it exists only because it is able to free-ride off of the innovations created in the proprietary part of the industry in which the monopoly power of copyright plays a key role. Certainly it is true that Linux is a knock-off of Unix and that Openoffice Writer is a knock-off of Microsoft Word. But this means little, because practically all software, proprietary or not, is an imitation of some other software. Microsoft Windows is an imitation of the Macintosh, which is an imitation of Smalltalk. Microsoft Word is an imitation of Wordperfect, which is an imitation of Wordstar. Microsoft Excel is an imitation of Lotus 1-2-3 which is an imitation of Visicalc. And so forth and so on.

A good example is the webserver.¹¹ The first webserver was written by Tim Berners-Lee at CERN in 1991 and was followed shortly by the NCSA webserver written by Robert McCool. Neither of these ever saw much commercial use, both were public domain, and both were effectively publicly funded. This initial pattern is similar to the way that basic research, for example, in pharmaceuticals is generally publicly funded. Following this, Netscape corporation introduced a proprietary webserver and at about the same time Apache took over the code from the NCSA webserver. Both of these servers survive today, with the Netscape server having mutated into the the Sun One webserver, and Apache having become the dominant force in the webserver industry. Many new features have been added to these servers since their inception, as well as to the competing Microsoft product - the evidence suggests that Apache has been at least as innovative as the others in introducing new features. Certainly there is no evidence here that open source model was less able to turn a basic experimental idea into a commercially viable product than the proprietary model, or that it free-rode off of ideas developed in a proprietary product.

¹¹Information about the history of the webserver is from the Wikipedia.

Another interesting case is that of word processing. Many open source alternatives to Microsoft Word exist, including Kword, AbiWord and OpenOffice Writer, the latter being the most widely used. How did the cost of developing this software - financed as it was by an open source model - compare to the cost of developing Microsoft Word? The fact is that most of the cost of writing software is not in the observation that it might be nice to have a button to justify text, or in the algorithms for spacing lines - which were after all developed by Gutenberg back in 1450 - but rather in the detailed implementation and debugging of the computer code. As far as we know, none of these open source projects benefited at all from the work done by Microsoft in developing their detailed computer code. Indeed, it appears as if the development of these open source projects was probably more expensive than the development of Microsoft Word - the single most difficult and expensive programming task faced by the developers of these projects appears to be the need to reverse engineer Microsoft Word documents and provide compatible formatting capability so that documents in Microsoft Word are usable and documents can be exchanged with Microsoft Word. Had these projects gone first, this substantial cost would have been avoided. It is also worth noting that the competitive rents generated by these projects is significantly smaller than they would have been had they hit the market before Microsoft Word did. So it seems reasonable to conclude that perfect competition would have delivered both these programs, as it did, and Microsoft Word as well.

Probably the most innovative program in the last few years is BitTorrent, a program that decentralizes and vastly increases the speed at which very large files can be downloaded off the internet. It is commercially successful in the sense that 50,000 copies a day are downloaded.¹² It is also sufficiently innovative that it is now being imitated- by Microsoft.¹³ BitTorrent, however, is open source and, according to their website Bram Cohen, the author, maintains the program for a living.

3.2. Books. Most innovation tends to bring with it some degree of monopoly power due to the first mover advantage. In our remaining case studies, we will examine markets where there is in fact some degree of monopoly power despite the absence of barriers to entry. It is important to emphasize the difference between monopoly profits, which are due to artificial scarcity, and competitive rents, which are due to natural scarcity. The former has distortionary welfare effects, the latter does not. Never-the-less, when potential competition is fierce, as it is in our remaining case studies,

¹²>From the statistics for the project provided by the SourceForge website.

¹³The Microsoft knockoff is called "Avalanche." Nicolai [2005] has details.

the effects of the initial monopoly are minimal, and the market approaches perfect competition.

Consider the production of books and literature. Today of course copyright is nearly ubiquitous, with works in the United States copyrighted as a matter of course, unless the author explicitly rejects it. This was not always the case. For example, in the United States, in the 19th Century, foreign works - including books written in England - were not entitled to copyright in the United States at all.

yet American publishers found it profitable to make arrangements with English authors. Evidence before the 1876-8 Commission shows that English authors sometimes received more from the sale of their books by American publishers, where they had no copyright, than from their royalties in [England] Arnold Plant [1934]

We should note that in 1850 U.S. population was 23.2 million; in 1851 U.K. population was 27.5 million. Per capita GDP in those same years, in 1996 U.S. dollars, was roughly \$1930 in the U.S. and \$2838 in the U.K. The literacy rates in both countries were roughly 85%. In other words, the U.S. market was smaller than the U.K. market, but of similar size.

American publishers found it profitable to pay the English authors so that they could get the initial copy ahead of their rivals, earning a substantial profit before copiers would have time to enter the market. That the monopoly distortion from this first mover advantage was small is indicated by the enormous price differential between the sale price of books without copyright in the U.S. and with copyright in the U.K. For example, Dickens' *A Christmas Carol* sold for six cents in the US, while it was priced at roughly two dollars and fifty cents in England.

However relevant this may be for our understanding of how innovations have been historically created, one may wonder whether data from an age of clipper ships and hand presses is relevant to an age of cheap electronic reproduction. The ubiquity of modern copyright makes this question difficult - but not impossible - to answer. Documents produced by the U.S. government are not subject to copyright, and a few have been commercial best sellers. The most significant government best seller of recent years has the rather off-putting title of *The Final Report of the National Commission on Terrorist Attacks Upon the United States*, but it is better known as the *9/11 Commission Report*. The report was released to the public at noon on Thursday July 22, 2004. At that time, it was freely available for downloading from a government website. A printed version of the report published by W.W. Norton simultaneously went on sale in bookstores. Norton had signed an interesting agreement with the government.

The 81-year-old publisher struck an unusual publishing deal with the 9/11 commission back in May: Norton agreed to issue the paperback version of the report on the day of its public release. (An indexed hardcover edition will follow.) Norton did not pay for the publishing rights, but had to foot the bill for a rush printing and shipping job; the commission did not hand over the manuscript until the last possible moment, in order to prevent leaks. The company will not reveal how much this cost, or when precisely it obtained the report. But expedited printings always cost extra, making it that much more difficult for Norton to realize a profit. In addition, the commission and Norton agreed in May on the 568-page tome's rather low cover price of \$10, making it that much harder for the publisher to recoup its costs. (Amazon.com is currently selling copies for \$8 plus shipping, while visitors to the Government Printing Office bookstore in Washington, D.C. can purchase its version of the report for \$8.50.) There is also competition from the commission's Web site, which is offering a downloadable copy of the report for free. And Norton also agreed to provide one free copy to the family of every 9/11 victim. Brendan Koerner [2004]

To be clear: what Norton received from the government was the right to publish first, and the right to use the word “authorized” in the title. What they did not get was the usual copyright - the right to exclusively publish the book. Because it is a government document, the moment it was released, other individuals, and more important, publishing houses, had the right to buy or download copies and to make and resell additional copies - electronically or in print - at a price of their choosing - in direct competition with Norton. In other words - after the release of the book on July 22, the market became a conventional competitive market.

The right to compete with Norton was not a purely hypothetical one. Another publisher, St. Martin's, in collaboration with the New York Times, released their own version of the report in early August - about two weeks after Norton¹⁴ - and this version contained not only the entire government report - but additional articles and analysis by New York Times reporters. Like the Norton version, this version was also a best seller. In addition, it

¹⁴This date and description of the content is from Wyatt [2004].

is estimated that 6.9 million copies of the report were (legally) downloaded over the Internet.¹⁵ Competition, in short, was pretty fierce.

Despite this fierce competition, the evidence suggests that Norton was able to turn a profit. We do not know, unfortunately, how much they would have paid up front to the “author” had the rights to go first been put out to bid. But we do have some idea of how much they made after the fact. First, we know that they sold about 1.1 million copies, and that they charged between a dollar and a dollar fifty more than St. Martin’s did. Other publishers also estimated Norton made on the order of a dollar of profit on each copy. Assuming that St. Martin’s has some idea of how to price a book to avoid losing money, this suggests Norton made at the very least on the order of a million dollars. We also know that their contract with the government called upon them to donate their “profits” to charity - and they did in fact “donate \$600,000 to support the study of emergency preparedness and terrorism prevention.”¹⁶ Since the entire Hollywood movie industry has managed by creative accounting to avoid earning a profit during its entire history, we can be forgiven if we suspect that Norton earned a bit more than the \$600,000 they admitted to.

To put these numbers in perspective, the 9/11 commission report was in paperback and, including free downloads, there seem to be about about 8 million copies in circulation. In contrast the initial print run for *Harry Potter and the Half-Blood Prince* was reported to be 10.8 million hardcover copies.¹⁷ So we can realistically conclude that if J. K. Rowling were forced to publish her book without the benefit of the monopoly conferred by copyright, she might reasonably expect to sell the first copy of the book to a publishing house for several million dollars - or more. This is certainly quite a bit less money than she earns with her current legal monopoly. But it seems likely, given her previous occupation as a French teacher,¹⁸ that it would still give her adequate incentive to produce her great works of literature.

Returning to the 9/11 commission report, the bulk of Norton’s profits were from the short-term monopoly of copies during the two weeks prior to entry.¹⁹ Since at most 300,000 people preferred to wait two weeks to

¹⁵May [2005] reports sales and download estimates. The Norton version sold 1.1 million copies and the other publisher St. Martin’s sold about 900,000 copies.

¹⁶This was reported in the Associated Press [2005].

¹⁷This figure was widely reported. See for example www.veritaserum.com.

¹⁸Reported in an on-line biography at gaga.essortment.com.

¹⁹Norton initially printed 600,000 copies, followed according to Koerner [2004] quickly by a print run of an additional 200,000 copies. It appears that these copies were sold prior to the entry by St. Martin’s. According to Wyatt [2004] the St. Martin’s version was available roughly two weeks after the Norton version. Subsequently, according to May [2005] an additional 300,000 copies were sold by Norton. Since according to best-seller

purchase a copy for a dollar less than Norton, the deadweight loss from the monopoly is at most \$300,000, and a more plausible estimate is to assume a uniform distribution of values, which would make the loss \$150,000. If we assume that the 800,000 people who purchased the book for \$8.00 during the first two weeks were indifferent to purchasing it two weeks later for \$7.00 and discount the value two weeks later by a factor²⁰ of 50%, then they value purchase of the book at \$9.00. So the 300,000 people that waited had values between \$7.00 and \$9.00. A reasonable approximation is to assume that the 800,000 people who did not wait had values ranging uniformly over \$9.00 to \$15.00. This implies an average per consumer surplus of \$4.00 for each 800,000 consumers who bought the \$8.00 copy during the first two weeks. In addition we estimate an average per consumer surplus of \$1.00 for the remaining 1.2 million copies sold after the first two weeks. So the total consumer surplus is about \$4.3 million, about thirty times the deadweight loss from monopoly. By way of contrast with unrestrained monopoly, linear demand, and constant marginal cost, the consumer surplus is equal to the deadweight loss from monopoly - and this is a welfare triangle that does not seem to trouble economists a great deal.

In the end it should be no great surprise that ideas of great social value are going to be produced under competition. The great blockbuster novels; the life-saving drugs - all generate such great surplus relative to the cost of creation that relatively little of that surplus need be captured by the innovator to make it worth her while. And indeed, the great works of Shakespeare and Mozart were created under conditions of perfect competition. What about more socially marginal creations? Naturally, the creator will have to capture a greater share of social surplus if these are to be created. On the other hand, competition is likely to be less fierce. While *Harry Potter and the Half Blood Prince* was scanned and illegally released onto the Internet within hours of appearing in print, we have been unable to find any trace of pirate versions of Sara Rath's opus *Star Lake Saloon and Housekeeping Cottages: A Novel* published six days earlier.

What is the effect of decreased reproduction costs on marginal creations? A widely held belief seems to be that Internet piracy will eliminate these creations unless the government intervenes, for example, by increasing the penalties for "piracy." Interestingly, the Internet seems to have increased rather than decreased the production of marginal ideas. A case in point is

reports in the Washington Post the Norton version outsold the St. Martin's version during the weeks immediately following the release of the St. Martin's version, we assume that the number of people who preferred to wait two weeks was no more than 300,000.

²⁰This is the "impulsive" discount factor used in hyperbolic discounting theory, so it is at the high range of how much we might imagine the near future reasonably being discounted. See Laibson [2005].

the creation of comic strips with very small audiences. With fixed costs of print runs and distribution - to say nothing of the cost of finding a few interested readers in a population of billions - such comics were never produced prior to the advent of the Internet. Now they are. Realistically, such small scale productions are never going to benefit from copyright or government intervention - it would never pay to sue someone over copying a comic strip that few people read, and would be equally hard to get the FBI interested in pursuing copyright violators. So the profits of these small productions come from competitive rents, and a large share of these competitive rents are due to complementary sales, as theorized above. In an earlier version of this paper we joked that novels would still be written under perfect competition as long as authors were able to sell signed copies of t-shirts. It turns out that this is not a joke for small audience comic strips.

Rosenberg raves that he has been able to make five times as much off his merchandising as off his subscriptions and that advertising doesn't come close to generating the revenue he gets off t-shirts, noting a profit margin of up to 50%, which would be as much as \$9 per item in some cases. Stevens quotes \$4-\$5 as his margin. Rosenberg further claims to have tripled his 2003 income by switching to t-shirt sales in the last three months of 2003. Todd Allen [2005]

3.3. Financial Securities. Prior to 1998, investment bankers and other firms selling financial securities operated without the “benefit” of IP protection. The rapid pace of innovation in financial securities until the late 1980s is well documented, for example by Tufano [1989]. Tufano estimates that roughly 20% of new security issues involve an “innovative structure.” He reports developing a list of some 1836 new securities over a 20 year period and remarks that this

severely underestimate[s] the amount of financial innovation as it includes only corporate securities. It excludes the tremendous innovation in exchange traded derivatives, over-the-counter derivative stocks (such as the credit derivatives, equity swaps, weather derivatives, and exotic over-the-counter options), new insurance contracts (such as alternative risk transfer contracts or contingent equity contracts), and new investment management products (such as folioFN or exchange traded funds.)

Herrera and Schroth [2004] examine the forces underlying this innovation. They show how a market for expertise in the presence of learning by doing

leads to substantial competitive rents, combined with a small amount of monopoly power that unravels over time.²¹

Like the case with books, this highlights the fact that there is generally some degree of “natural” monopoly power generating distortionary profits along with competitive rents. This poses an interesting issue for policy: the optimal term of patent or copyright may be negative. That is, rather than granting and enforcing private monopolies, the optimal policy might be for the government to reduce the monopoly power from being the first mover. The most obvious source of monopoly power that could be eliminated by public policy is that from secrecy. For example, a requirement for selling a machine would be the publication of its blueprints; a requirement for selling software would be the publication of the source code; and rather than enshrining “Digital Rights Management” into the Digital Millennium Copyright Act, it could be made illegal.

3.4. Pharmaceuticals. In any discussion of innovation under perfect competition, the pharmaceutical industry quickly comes up. The cost of bringing a new drug to market is large - on the order of \$200 million 1989 dollars.²² Moreover, companies are required to disclose the chemical formula for their products as part of the FDA approval process, and to make available to other manufacturers the results of their clinical trials. That is, without patents, this industry would operate under “negative” patent in which the government forces disclosure as a condition of doing business. It is widely perceived that with the elimination of patents in this industry, generics would enter the market at roughly the same time as the original and there would be no profit or rent with which to cover the high cost of creating new drugs.

It turns out that this is far from obvious. From Lanjouw [1999] we can examine the behavior of the Indian pharmaceutical industry. Since 1972 product patents on pharmaceuticals have not been recognized in India. Nevertheless, it takes about 5 years for a new drug to enter the Indian market as a generic following its introduction elsewhere. There are two reasons for this. First, the generic manufacturers generally wait a year to see how the new drug does on the market before making the decision to enter.²³ This

²¹For more details on the financial securities and the pharmaceutical industry, considered next, see Boldrin and Levine [2005].

²²Hansen et al [1991].

²³Entry of generics following the expiration of a patent in the U.S. is much quicker - according to the CBO [1998] only about a month. However, when the patent expires, the generic manufacturer has had nearly 10 years to observe the drug in use and to make plans for entry. For this reason, the Indian market where entry is possible at any time gives a better indication of the effect of abolishing patent protection.

highlights a problem with the common view that imitators have an advantage because they only need to imitate successful products. While it is true that few would spend a lot of money imitating a product that sold very little at a low price, by the time the imitator has learned that the original is selling a lot of units at a high price, the innovator has already pocketed quite a bit of money.

Second, the process of actual imitation and clearing regulatory hurdles takes 3-4 years. Lanjouw conjectures that the amount of time to imitate is short and that the primary delay is the regulatory one, but there not yet much data that would shed light on this. Regardless, from the creators point of view, in the absence of patent protection - which due to the fact that the patent must be filed long before a drug is approved and marketed lasts only about 10 years - it appears that the innovator of a new product will enjoy a 5 year rather than 10 year monopoly. The evidence suggests that when generics enter the price of the original does not change much, and the original retains about 50% market share.²⁴ Suppose a 6% real interest rate and normalize the flow of revenues from monopoly to be 1. Then with 10 years of protection present value is roughly $.6 \times 1 + .4 \times .5 = .85$. With 5 years of protection this becomes $.3 \times 1 + .7 \times .5 = .65$. In other words, even if we continue to enforce the revelation of trade secrets for free in this industry and eliminate patents entirely, the loss of present value would be only about 25% of current rents. As a public policy issue, this must be set against the fact that patents serve as a discouragement to future innovation as well as a spur to current innovation, so that the steady state effect of patents on innovation can easily be negative.

4. CONCLUSION

There is one question that we have carefully avoided asking - given that innovation can and does thrive in the absence of monopoly power, and certainly in the absence of the artificial monopoly power imposed by government copyright and patent regulation - what should government policy be? From a social point of view are copyright and patents a good idea? Or perhaps government policy should go the opposite direction - should trade secrecy be abolished? We have no intention of trying to answer those questions here. The trade-off is complicated - certainly competitive rents can and do sustain innovation, and substantial amounts of innovation. But competitive rents are not equal to social value, so in the presence of fixed costs, there will be socially valuable innovations that will not occur. So we cannot say that government sustained monopoly over ideas is necessarily a bad idea. Regardless of what ones priors are on this however, we do not see

²⁴See, for example, CBO [1998].

how it is possible to have a sensible, let alone correct, discussion of policy without understanding first how and why the absence of monopoly by no means implies the absence of innovation.

REFERENCES

- [1] Aghion, P. and P. Howitt (1992), "A Model of Growth through Creative Destruction," *Econometrica* **60**, 323-351.
- [2] Allen, T. [2005], *Online Comics Vs. Printed Comics: A Study in E-Commerce and the Comparative Economics of Content*, IndignantOnline.
- [3] Arrow, K.J. (1962), "Economic Welfare and the Allocation of Resources for Invention," in Richard Nelson (ed.), *The Rate and Direction of Inventive Activity*, Princeton, NJ: Princeton University Press.
- [4] Associated Press (2005), "Publisher of 9/11 Report to Make Donation," *Associated Press*, July 21.
- [5] Barro, R.J. and X. Sala-i-Martin (1999), *Economic Growth*, Cambridge, Mass.: MIT Press.
- [6] Becker, G., [1971], *Economic Theory*, Knopf Publishing Co.
- [7] Boldrin, M. and D. Levine (1997), "Growth under Perfect Competition," UCLA and Universidad Carlos III de Madrid, October.
- [8] Boldrin, M. and D.K. Levine (2002), "The Case Against Intellectual Property," *The American Economic Review (Papers and Proceedings)* **92**, 209-212.
- [9] Boldrin, M. and D.K. Levine (2005), *Against Intellectual Monopoly*, forthcoming.
- [10] CBO (1998), "How Increased Competition from Generic Drugs has Affected Prices and Returns in the Pharmaceuticals Industry," Congressional Budget Office.
- [11] Grossman, G.M. and E. Helpman (1991), "Quality Ladders in the Theory of Growth," *Review of Economic Studies* **58**, 43-61.
- [12] Grossman, G. and E. Helpman (1994), "Endogenous Innovation in the Theory of Growth," *The Journal of Economic Perspectives*, **8**, 23-44.
- [13] Hann, I., J. Roberts, S. Slaughter, and R. Fielding (2004), "An Empirical Analysis of Economic Returns to Open Source Participation," Carnegie-Mellon University.
- [14] Hansen, R. W., H. G. Grabowski and L. Lasagna (1991), "Cost of Innovation in the Pharmaceutical Industry," *Journal of Health Economics* **10**, 107-142.
- [15] Hellwig, M. and A. Irmen (2001), "Endogenous Technological Change in a Competitive Economy," *Journal of Economic Theory*, **101**, 1-39.
- [16] Herrera, H. and E. Schroth (2004), "Developer's Expertise and the Dynamics of Financial Innovation: Theory and Evidence," FAME-Univ. of Lausanne, Research Paper No. 124, October; and mimeo, ITAM, Mexico D.F. .
- [17] Hui, K.L. and I. P'ng (2003), "Piracy and the Legitimate Demand for Recorded Music," *Contributions to Economics Analysis and Policy* **2**, 11.
- [18] Karr, R. (2002), "TechnoPop: The Secret History of Technology and Pop Music," NPR, <http://www.npr.org/programs/morning/features/2002/technopop/>.
- [19] Koerner, B. I. (2004), "Who Gets the 9/11 Report Profits? The publisher's printing at least 800,000 copies. Where will all the money go?" *State*, July 27, 2004.
- [20] Laibson, D. (2005), "Intertemporal Decision Making," forthcoming in *Encyclopedia of Cognitive Science*.
- [21] Lanjouw, J. O. (1999), "The Introduction of Pharmaceutical Product Patents in India: Heartless Exploitation of the Poor and Suffering?," NBER WP# 6366

- [22] Leibowitz, S. (2002), "Record Sales, MP3 Downloads and the Annihilation Hypothesis," mimeo, University of Texas, Dallas.
- [23] Lerner, J. and J. Tirole (2004), "The Economics of Technology Sharing: Open Source and Beyond," NBER Working Paper 10956.
- [24] May, E. R. (2005), "When Government Writes History: The 9-11 Commission Report," *History News Network*, <http://hnn.us/articles/11972.html>.
- [25] Niccolai, J. (2005), "Microsoft readies BitTorrent alternative Avalanche technology could make it easier to distribute big files over the Internet," *IDG News Service*, June 16.
- [26] Plant, A. (1934), "The Economic Aspect of Copyright in Books," *Economica*, 167-195.
- [27] Quah, D. (2002), "24/7 Competitive Innovation," mimeo, London School of Economics.
- [28] Romer, P.M. (1990), "Endogenous Technological Change," *Journal of Political Economy* **98**, S71-S102.
- [29] Schumpeter, J. (1942), *Capitalism, Socialism, and Democracy*, Harper Perennial (1962).
- [30] Shell, K. (1966), "Toward a Theory of Inventive Activity and Capital Accumulation," *The American Economic Review (Papers and Proceedings)*, **56**, 62-68.
- [31] Shell, K. (1967), "A Model of Inventive Activity and Capital Accumulation," in *Essays on the Theory of Optimal Economic Growth* (K. Shell, ed.), Cambridge, Massachusetts: MIT Press, 67-85.
- [32] Stigler, G. J. (1956), "Industrial Organization and Economic Practice," in *The State of the Social Sciences* edited by Leonard D. White, University of Chicago Press.
- [33] Tofuno, P. (1989), "First Mover Advantages in Financial Innovation," *Journal of Financial Economics*, **3**, 350-370.
- [34] Wyatt, E. (2004), "For Publisher of 9/11 Report, a Royalty-Free Windfall," *New York Times*, July 28.

BOLDWIN: UNIVERSITY OF MINNESOTA, LEVINE: UCLA, BOTH: FEDERAL RESERVE BANK OF MINNEAPOLIS