

Employment Mobility Laws and Competition

April M. Franco and Matthew F. Mitchell*

July 22, 2005

Contents

1	Introduction	1
2	Background	4
3	Do Employment Mobility Restrictions Matter?	5
4	Overview of the Analysis	6
5	The Formal Model	8
5.1	Environment	8
5.2	Retain or Not? No CNC Protection	8
5.3	Retain or Not? CNC Protection	9
5.4	Value of Entry	11
5.5	Number of Firms	13
6	Conclusion	15

1 Introduction

Hovenkamp discusses the relationship between antitrust laws and intellectual property (IP) laws. He points out that the conflict between the two

*Department of Economics, University of Iowa, Iowa City, IA 52245. This paper was prepared for the Antitrust Enterprise Conference. We thank Herb Hovenkamp for helpful suggestions and participants at the conference for comments.

“is explained by deep uncertainty about the optimal amount and scope of IP protection. As long as that uncertainty remains there will always be tension between IP and antitrust.”¹

In this paper we study a particular sort of intellectual property protection, laws that help firms to impose barriers on employee mobility, such as covenant not to compete (CNC). Our goal is to analyze, in a simple economic model, the role that the scope of these IP laws have on the degree of competition. We see this as a necessary step toward understanding the complicated interactions underlying the optimal scope of IP protection.

In practice, the degree of protection afforded employers suggests that the issue is an important one. States differ significantly in their policies toward the intellectual property of employers, and have seen different histories as a result. In Massachusetts, where the legal system is based on the English system, employment contracts with CNC are enforceable by the courts, while in California, where the law, written in the mid 1800’s, is codified and does not specifically allow for the enforcement of such contracts.² There are some states, such as Texas, where CNCs are typically not enforced.³ The differential enforcement of these clauses has been constant for more than a century in Massachusetts and California, and the results are strikingly different.⁴

There is evidence of hyper-mobility between high tech firms in California, which is not evident in Massachusetts.⁵ In the hard drive industry, more than a quarter of the start-ups in California were spin-outs, firms started by former employees of incumbent firms, while almost all of the start-ups in Massachusetts were de novo firms.⁶ In addition, while the number of firms in

¹H.J. Hovenkamp, “United States Antitrust Policy in an Age of IP Expansion,” U Iowa Legal Studies Research Paper No. 04-03 (2004).

²Ronald J. Gilson, “The Legal Infrastructure of High Technology Industrial Districts: Silicon Valley, Route 128, and Covenants Not To Compete,” 74 *New York University Law Review*, 579 (1999).

³See id. at 627.

⁴Annalee Saxenian, *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*, Cambridge: Harvard University Press, (1994).

⁵Bruce Fallick, Charles A. Fleishman, and James B. Rebitzer, “Job-Hopping in Silicon Valley: Some Evidence Concerning the Micro-Foundations of a High Technology Cluster,” Finance and Economics Discussion Series, Federal Reserve Board, (2005).

⁶April Franco and Darren Filson, “Spin-outs: Knowledge Diffusion through Employee Mobility,” Federal Reserve Bank of Minneapolis Staff Report 272, (2000).

Massachusetts in the electronics industry and more specifically, in the hard drive industry, was higher initially, but by the mid 1980's, the number of firms in California was significantly higher and continuing to grow, while the number in Massachusetts was in significant decline.⁷

We are motivated to study this form of IP law because of the great uncertainty about the optimal scope of such regulations. As a first step toward understanding optimal policy, we seek to understand these laws' implications for competition. History suggests that these policies clearly affect the number of firms in an industry, and therefore are relevant to an antitrust authority concerned with the level of competition. There are at least two ways that these rules might affect the number of firms. On the one hand, to the extent that these policies reduce start-ups by employees, they thwart competition. On the other hand, to the extent that they protect IP, they encourage firms to enter with new ideas when the protection is better.

The purpose of this paper is to put forward a simple model of the two effects, in order to determine under what conditions one or the other might dominate. The model is intended as a starting point for discussing the role of these policies on the level of competition in industries where employee turnover is an important source of new firms. As such, our intent is to provide a guide by including a few key economic ingredients necessary in the construction of such a model.

Our results suggest that, in this case, IP protection and competition need not be in conflict. In fact, in some cases, more competition results from greater IP protection, since it encourages entry by entrepreneurs with new ideas. We find this to be the case when either additional competition reduces profits greatly, or when additional competition has nearly no effect on firm's profits. The conflict between IP protection and antitrust is most likely to arise in the intermediate cases. There, strong IP protection for employers will tend to foster fewer firms in the long run. This result can help to focus the empirical debate over the net effects of IP protection on competition.

Moreover, our results are suggestive of the inherent complexity in the link between IP and competition, as Hovenkamp suggests. We find that the key variable in determining the effect of IP on competition is the degree to which profits are affected by the inclusion of one additional firm. This variable,

⁷See Christensen, Clayton M., "The Rigid Disk Drive Industry: A History of Commercial and Technological Turbulence," *Business History Review*, 67 (1993); Saxenian *supra* note 4.

however, is typically related to competition itself: the greater is competition, the less is the impact of one additional firm. So our results suggest that not only does the scope of IP affect competition, but competition affects the impact of IP protection.

In the following section, we give a brief overview on the enforcement of CNCs. The third section is a discussion on the relevance of CNCs to turnover and conditions under which CNCs effect turnover and competition. A brief overview of the model is described in the fourth section. The formal model is presented in the fifth section, along with the conditions under which enforcing CNCs will increase or decrease competition. The final section concludes.

2 Background

The effects of IP protection on competition have been considered indirectly in both the legal and economic literature. In the legal literature, there is general consensus that IP protection in the form of CNCs has important implications for the creation of new firms. Gilson argues that the number of firms in California is higher than that in Massachusetts because of a “serendipitous result of historical coincidence;” specifically the creation of California’s codified legal system in the 1870’s. Because California’s legal code disallowed the enforcement of CNCs, employees were able to create new firms which often competed with their former employers. Another legal scholar, Hyde , argues that the reason for this difference in firm numbers is directly related to the differential enforcement of trade secrets laws and inevitable disclosure rules, other forms of IP protection.⁸ He argues that because trade secret laws provide a secondary net to prevent the spread of tacit knowledge, the number of firms in Massachusetts is lower since the California legal system does not enforce them as strictly.

Recent work in the economics literature on sequential innovation, including O’Donoghue et al. and Hopenhayn et al. has the feature that IP protection, in particular patents, granted to current innovators might be detrimental to future innovators.⁹ In a sense, then, these papers are in the spirit

⁸Alan Hyde, *Working in Silicon Valley: Economic and Legal Analysis of a High-Velocity Labor Market*, M.E. Sharpe, (2003).

⁹T. O’Donoghue, Scotchmer, S., and Thisse, J-F.. “Patent Breadth, Patent Life and the Pace of Technological Improvement.” 7 *Journal of Economics and Management Strategy*, (1998); H. Hopenhayn, Llobet, G. and Mitchell, M. “Rewarding Sequential Innovations:

of the papers on CNCs, but are aimed at a different sort of IP protection. This paper extends those ideas to a model of an employer-employee relationship, rather than to a framework of sequential innovations which build on one another.

In contrast, the classic economics literature, dating back to at least Arrow, stresses that in order for entrepreneurs to be interested in pursuing new ideas, which may be costly and take time to recoup development costs, they must be afforded some level of IP protection.¹⁰ If the new ideas are particularly expensive in terms of start-up costs and the per period profits are below these initial costs, these ideas will not be developed without some form of monopoly protection. Hence only those ideas with initial profits that cover the development will result in firms without any IP protection. So, by providing some IP protection, the government increases the number of firms and the level of competition.¹¹

3 Do Employment Mobility Restrictions Matter?

Pakes and Nitzan point out that there may be ways for employers to protect themselves against competition from employees even when no statutory protection is available.¹² One way to do this would be for the employee to post a bond at the beginning of his employment. If he stays, the bond is returned. If the employee leaves, the employer receives the bond payment. Equivalently, the contract could stipulate a wage for both the first period of employment and the second period of employment. This alternative to bond posting requires that the wage is artificially low in the first period and

Patents, Prizes, and Buyouts,” Federal Reserve Bank of Minneapolis, Staff Report 273, (2000).

¹⁰K. Arrow, “Economic Welfare and the Allocation of Resources for Inventions”, in *The Rate and Direction of Inventive Activity: Economic and Social Factors*, Princeton, N.J., National Bureau of Economic Research, (1962).

¹¹A. Franco and M. Mitchell, “Covenants Not to Compete, Labor Mobility, and Industry Dynamics,” manuscript, University of Iowa, (2004) studies competition between two regions with different CNC enforcement policy and finds that the salient features of the industry lifecycle is very similar to that seen in California and Massachusetts.

¹²Ariel Pakes and Saul Nitzan, “Optimal Contracts for Research Personnel, Research Employment and Establishment of ‘Rival’ Enterprises,” 1 *Journal of Labor Economics*, (October, 1984).

then second period wage is significantly higher. If the worker stays, the high second period wages compensate for the low first period wages; if the worker leaves, the low first period wages, without further compensation, work like the bond's payoff, protecting the employer from competition with a former employee.

Such an arrangement provides contractual protection without a true CNC, but might be problematic. If the courts are unwilling to enforce a CNC, would they enforce the bond repayment in the second period, which is performing the same function? The alternative, a low first period wage, might run afoul of either minimum wage laws, or be impossible for the employee to accept, if credit market imperfections keep him from being able to finance a period of arbitrarily low wages. Nonetheless, it is important to realize that it is possible that, regardless of whether a firm is in a state that enforces CNCs or not, actual protection could be the same, due to other forms of contractual protection.

Even if no form of contractual relief is available for the employer, the Coase conjecture suggests that, absent additional frictions, allocations are independent of the allocation of property rights, in this case the right to determine the employee's second period activity. However, this would require efficient *ex ante* contracting; in particular, it would require that all of the firm's decisions (including the entrepreneur's decision to form the firm in the first place) be made after the contracts with employees are written. Clearly such early contracting is typically impossible, so the allocation of property rights matter. We will proceed by assuming that no dynamic contracts can be enforced in the absence of CNC protection. In practice, firms have some protection even without CNCs via other IP protection; however, in order to highlight the effects of different levels of employee mobility protection, we abstract from other forms of IP protection to make the differences as stark as possible.

4 Overview of the Analysis

Imagine an entrepreneur has an idea for a new firm. In order to implement her idea, she needs to hire an outsider, who we will call the employee. In the course of employment, the employee may learn the idea; only the employee knows if he has learned. If the employee does learn the idea, he can start up a firm of his own for free.

After introducing the formal structure, we start the analysis in the situation where the employee may or may not have learned, and the employer can decide whether or not to keep the employee. The employer must weigh two factors. First, she must consider her enterprise's expected value without competition from the employee, versus the expected value with such competition. Second, she must consider how much it will cost to keep competition from occurring. The second part is where the legal system will matter: with CNCs, the employer need only pay the employee the wage, while without CNCs, the employer must compensate the employee so that the employee is indifferent between leaving or staying.

The goal of this exercise is to determine under what conditions the enforcement of CNCs will increase the number of firms and when it will not. What we find is that, since enforcing CNCs has two effects, it depends on how profits change with competition. There are three cases. In the two extreme cases where either profits are largely unaffected by the entrance of a spin-out, or profits are severely affected by the entrance of a spin-out, CNC enforcement leads to more firms. In the case where profits after the entrance of a spin-out are in the moderate range, CNC enforcement leads to less firms.

The critical insight is that more CNC enforcement may not lead to more firm formation by employees. This is precisely what happens in the cases where CNC enforcement leads to more firms: despite retention being more costly without CNCs, in some cases the benefits of retention are either so large or so small that the decision to retain is the same. In particular, when the profits are largely unaffected by the spin-out, the benefits of retention are small, and the worker is not retained even when a CNC allows retention at low (but not zero) cost. On the other hand, when the profits are hugely affected by the spin-out, the benefits of retention are so great that the employee is retained even if the lack of CNC protection makes the costs high. In these cases, the only effect of CNCs is to encourage entry by protecting entrepreneurs.

The intermediate situation, where retention depends on whether or not a CNC is available, is precisely the case where CNCs can decrease the number of firms. It is worth noting that this case is considered the empirically plausible range for the computer industry, given the literature that suggests that CNC protection led to more turnover when CNCs were not allowed.

5 The Formal Model

5.1 Environment

The entrepreneur (or, later, the employer) has an idea for a new firm. Implementing the idea will cost c . The value of c is drawn from the uniform distribution, so that the probability of a cost below any given value of c , denoted by $F(c)$, is proportional to c . In addition to paying c , the entrepreneur must hire an outsider, who we call the employee, and show that outsider the idea. Upon seeing the idea, the employee learns the idea and how to implement the idea with probability λ . Only the employee knows if he has learned, and the employee can only successfully form a competing firm if he has learned.

The entrepreneur's profits depend on whether or not she faces competition from the outsider. If the employee does not form a firm, the employer can earn π_h . If employee does form a firm, the employer earns $\pi_l \leq \pi_h$, and the employee earns $\sigma\pi_l$. The variable $\sigma \leq 1$ allows for the possibility that the employee does not learn how to operationalize the idea as well as its original inventor. These profits include market wages to all employees; any above-market wages are deducted from these profits. We assume both π_l and π_h are strictly positive.

We assume that adverse selection concerns keep the employer from charging the outsider to see the idea, along the lines of Pakes and Nitzan.¹³ Instead, the employer has to decide whether or not to pay the employee to keep him from forming a firm (in the case where no CNC protection is available), or whether or not to allow the worker to leave, in exchange for a payment (if CNC protection is available).

5.2 Retain or Not? No CNC Protection

Suppose that employee is free to leave; there is no protection of any kind afforded the employer. Then the employer can only retain by offering at least $\sigma\pi_l$ for the employee to stay; this can be interpreted as an above market wage. As a result, the employer makes $\pi_h - \sigma\pi_l$ if she retains the employee, regardless of whether the employee has actually learned. If, on the other hand, she chooses not to pay the employee, she faces possible competition. With probability λ the employer earns π_l ; however, it is possible (with complementary

¹³id.

probability $1 - \lambda$) that the employee did not learn, and therefore the employer will not face competition. Therefore the employer makes $\lambda\pi_l + (1 - \lambda)\pi_h$ if she offers no inducement to keep the employee. The employer will choose not to retain the employee if this expected profit is greater than the profit from retaining the employee; in other words, the employer allows the employee to leave if

$$\lambda\pi_l + (1 - \lambda)\pi_h > \pi_h - \sigma\pi_l$$

or

$$\pi_l/\pi_h > \frac{\lambda}{\lambda + \sigma}$$

Note that the profits ratio can range between 0 and 1. The first boundary is defined by the case when the profits are driven to zero when the employee enters the market. In this case, the employer was making positive economic profits, but with the entry of a new firm, due to Bertrand competition, the employer's and employee's profits are driven to zero. The second boundary is defined by the case when the profits are the same regardless of whether the employee becomes a competitor or not. A classic economic example of this would be a perfectly competitive market; individual profits are unchanged with a single additional firm, since a single firm is tiny relative to the industry and therefore has no effect on prices.

5.3 Retain or Not? CNC Protection

We model CNC protection as the ability to keep the employee from forming a firm. Moreover, we give the employer the ability to waive that right, and allow the employee to start a firm, in exchange for a fee. The employer can ask for as much as $\sigma\pi_l$ as such a payment, or not allow a buyout and make profits of π_h by retaining the employee for sure. If the employer allows the buyout, then she makes $\pi_l + \sigma\pi_l$ with probability λ (if the employee learns), and π_h otherwise. In other words, the employer will allow the buyout if

$$\lambda(1 + \sigma)\pi_l + (1 - \lambda)\pi_h > \pi_h$$

or

$$\pi_l/\pi_h > \frac{1}{1 + \sigma}$$

We are left with the following depiction of employee mobility, as a function of the legal environment and π_l/π_h :




Figure 1: Worker Retention

If the profits ratio is below $\frac{\lambda}{\lambda+\sigma}$, then there is no employee turnover, regardless of the legal structure. Firms with the benefit of CNC enforcement find it more profitable to prevent employees from leaving because the resulting profits for both firms is too low. Firms without the benefit of CNC enforcement find it more profitable to keep employees by paying them the above market wage of at least $\sigma\pi_l$. If the profits ratio is between $\frac{\lambda}{\lambda+\sigma}$ and $\frac{1}{1+\sigma}$, then employees leave only if the employer does not have CNC protection. As a result, employees who have learned will leave and become competitors. Employees who have not learned will remain with the firm. In the case where the profits ratio is greater than $\frac{1}{1+\sigma}$, the turnover will again be the same regardless of the IP protection: either way, employees will be allowed to leave if they learn. The difference is that, with a CNC, the employer receives the buyout payment. The employees must pay $\sigma\pi_l$ upon leaving.

The simple intuition is the lack of CNC protection makes keeping employees more costly, and therefore makes it occur for a smaller set of profit ratios π_l/π_h . Moreover, the decision is monotone, in the sense that, the greater is π_l/π_h , the greater is the incentive to let the employee leave under either IP regime.

5.4 Value of Entry

We now can compute the value of entry under each legal environment. The value of entry without CNC protection, denoted V_{np} , is

$$V_{np} = \begin{cases} \pi_h - \sigma\pi_l, & \text{if } \pi_l/\pi_h \leq \frac{\lambda}{\lambda+\sigma} \\ \lambda\pi_l + (1-\lambda)\pi_h, & \text{if } \pi_l/\pi_h > \frac{\lambda}{\lambda+\sigma} \end{cases}$$

The first branch, when the worker is retained, shows that the value of entry is equal to profits minus above-market wages. The second branch, when the worker is allowed to leave, is the expected value of the firm given that the worker may or may not leave.

The value of entry with CNC protection, denoted V_p , is

$$V_p = \begin{cases} \pi_h, & \text{if } \pi_l/\pi_h \leq \frac{1}{1+\sigma} \\ \lambda(1+\sigma)\pi_l + (1-\lambda)\pi_h, & \text{if } \pi_l/\pi_h > \frac{1}{1+\sigma} \end{cases}$$

Again, the first branch is the case where the worker is not allowed to compete; the second branch is where the worker is allowed to compete. In that case the worker must pay the buyout $\sigma\pi_l$. The ratio V_{np}/V_p tells us the relative value that results from different legal environments. Dividing the equations for the two values gives

$$V_{np}/V_p = \left\{ \begin{array}{l} 1 - \sigma(\pi_l/\pi_h), \text{ if } \pi_l/\pi_h \leq \frac{\lambda}{\lambda+\sigma} \\ (1-\lambda) + \lambda(\pi_l/\pi_h), \text{ if } \pi_l/\pi_h \in \left(\frac{\lambda}{\lambda+\sigma}, \frac{1}{1+\sigma}\right] \\ (\lambda\pi_l + (1-\lambda)\pi_h) / (\lambda(1+\sigma)\pi_l + (1-\lambda)\pi_h), \text{ if } \pi_l/\pi_h > \frac{1}{1+\sigma} \end{array} \right\}$$

By graphing the function, it is easy to see its non-monotonicity in π_l/π_h :




Figure 2: Relative Value of Entry

Figure 2 is in contrast with the decision to retain the worker (Figure 1), which is monotone in π_l/π_h : the greater is π_l/π_h , the less the employer has to lose in letting the employee compete, and therefore the greater is the incentive to let the worker leave. The effect on value is much more complicated. For the values of π_l/π_h less than $\frac{\lambda}{\lambda+\sigma}$, where the employee is not allowed to leave under either legal environment, the difference between the two IP regimes is that, without protection, the employer must pay $\sigma\pi_l$; in either case, the firm makes π_h . As π_l/π_h increases, this payment *to* the worker is becoming a greater fraction of the firm's profits, and therefore the relative value of the firm without CNC protection is falling.

At the other end of the spectrum, when π_l/π_h are greater than $\frac{1}{1+\sigma}$, the employer is allowing the employee to leave if he learns regardless of the IP regime. The difference in value comes from the fact that, with a CNC, the employer is entitled to the buyout payment $\sigma\pi_l$. Again, as π_l/π_h increases, this payment *by* the worker is becoming relatively more important, and so the relative value of being without protection is falling.

When the value of π_l/π_h is between $\frac{\lambda}{\lambda+\sigma}$ and $\frac{1}{1+\sigma}$, the employee leaves only when there is no CNC enforcement. With CNC enforcement, the firm's value is π_h ; without it, the firm sometimes makes π_h (if the employee fails to learn) but sometimes makes π_l (if the employee does learn, and therefore leaves to start his own firm). As π_l/π_h rises, the latter outcome is becoming closer to the former, and therefore the employer is losing less when the employee leaves. This means that a rise in π_l/π_h makes the relative value without protection rise, since the risk of the employee starting a competing firm is become less and less of a loss in profits.

As one would expect, value is always lower without the option of a CNC; this is the usual sort of IP result that greater protection yields higher return for the owner of the IP, in this case the employer. However, note that there is an interesting role of π_l/π_h in determining just how much benefit the IP protection provides. It is natural to think of V_{np}/V_p as being inversely related to the benefit of IP protection for the employer; it is interesting to note that one cannot determine the benefit of IP protection independently of π_l/π_h . In fact, the benefit of IP protection is not even monotone in the profit ratio.

While our focus is the role of IP protection in determining the amount of firms, there is also a sense in which our results suggest a feedback effect, that the number of firms may affect the way in which IP protection works. In general, economic models typically have the feature that the greater is competition, the less is the impact of one additional firm. In other words, the greater is competition, the greater is π_l/π_h . Our results imply that the benefit of CNC protection for individual firms at any given time depends on the level of competition in a nontrivial way.

5.5 Number of Firms

Now that we have calculated the value of entry, we can compute the expected number of firms per initial idea of an entrepreneur. The initial idea is implemented if cost c is less than or equal to value of entry. Since we are using the uniform distribution, the probability of entry, per innovator, is, for any value of entering V , equal to $F(V)$, which is proportional to V . The expected number of firms, per entrant, is $1 + \lambda$ if the employer allows the employee to compete, and 1 otherwise.

Given these calculations, the expected number of firms per idea, N_{np} ,

without protection, is proportional to value:

$$N_{np} \simeq \begin{cases} V_{np}, & \text{if } \pi_l/\pi_h \leq \frac{\lambda}{\lambda+\sigma} \\ V_{np}(1+\lambda), & \text{if } \pi_l/\pi_h > \frac{\lambda}{\lambda+\sigma} \end{cases}$$

Likewise, the expected number of firms per idea, N_p , with protection, is

$$N_p \simeq \begin{cases} V_p, & \text{if } \pi_l/\pi_h \leq \frac{1}{1+\sigma} \\ V_p(1+\lambda), & \text{if } \pi_l/\pi_h > \frac{1}{1+\sigma} \end{cases}$$

The constant of proportionality in both expressions is identical.

To understand the impact of the legal environment in determining the number of firms, we focus on the ratio N_{np}/N_p . It tells us the relative amount of competition (in terms of the number of firms) that results from different legal environments. We see that the relative number of firms is

$$N_{np}/N_p = \begin{cases} V_{np}/V_p, & \text{if } \pi_l/\pi_h < \frac{\lambda}{\lambda+\sigma} \\ (1+\lambda)V_{np}/V_p, & \text{if } \pi_l/\pi_h \in \left(\frac{\lambda}{\lambda+\sigma}, \frac{1}{1+\sigma}\right] \\ V_{np}/V_p, & \text{if } \pi_l/\pi_h > \frac{1}{1+\sigma} \end{cases}$$

In the first and last cases the ratio is surely less than one; there are more firms with the more protective legal environment. This is because turnover is identical across legal environments for these cases, so the only difference is that stronger protection encourages greater initial entry. However, in the middle case, there is more turnover when less protection is afforded, but fewer entrepreneurs choose to develop ideas. In this case, the first effect dominates and the ratio is greater than one, as seen in the graph below:




Figure 3: Relative Firm Numbers

The model rationalizes the idea that limited IP protection for firms might be good for competition, because it promotes spin-outs, even though it does discourage development by initial entrepreneurs.

6 Conclusion

While this is an extremely stylized model of labor mobility, it is a first step at understanding the effects of IP legislation on competition. This model helps to highlight the two opposing effects on competition from restrictions on employee mobility. While enforcement of CNCs helps to lower the value appropriation by employees, it may also help to increase initial investment and entry by firms, thus increasing competition. Only in the case where the profits ratio is in the moderate range does enforcing CNCs lower competition. These results point to the possibility that instead of using a one size fits all type of CNC legislation, it might be better to consider the industry and the current level of competition first.

The model also demonstrates the delicate balance between IP and competition. We have highlighted the fact that IP protection affects competition, since it influences both the incentives for entrepreneurs to pursue new ideas, and the ability of workers to use those ideas to form firms of their own. But there is a further feedback effect: to the extent that π_l/π_h depends on the degree of competition, we see that competition is affecting the implications of IP protection. As Hovenkamp suggests, this is a delicate balance; understanding this delicate balance in a more fully developed model is a clear direction for future research.