Antitrust in Innovative Industries

By Ilya Segal and Michael D. Whinston*

We study the effects of antitrust policy in industries with continual innovation. Antitrust policies that restrict incumbent behavior toward new entrants may have conflicting effects on innovation incentives, raising the profits of new entrants, but lowering those of continuing incumbents. We show that the direction of the net effect can be determined by analyzing shifts in innovation benefit and supply, holding the innovation rate fixed. We apply this framework to analyze several specific antitrust policies. We also show that, in some cases, the tension does not arise, and policies that protect entrants necessarily raise the rate of innovation. (*JEL* K21, L13, L14, L40, O30)

This paper is concerned with the effects of antitrust policy in markets in which innovation is a critical determinant of competitive outcomes and welfare. Over the last two decades, a large share of the economy—the so-called "new economy"—has emerged that shares this feature (see, for example, David S. Evans and Richard Schmalensee 2002). Traditionally, however, antitrust analysis has tended to ignore issues of innovation, focusing instead on the price/output effects of contested practices.[1] Sparked by this disparity, and the recent Microsoft case (Civil Action No. 98-1232), a number of commentators have questioned whether traditional antitrust analysis is poorly suited to maximizing welfare in such industries. Evans and Schmalensee (2002) state the concern succinctly:

> [In these industries] firms engage in dynamic competition *for the market*—usually through research and development (R&D) to develop the "killer" product, service, or feature that will confer market leadership and thus diminish or eliminate actual or potential rivals. Static price/output competition on the margin *in the market* is less important.

Unfortunately, the effects of antitrust policy on innovation are poorly understood. In the Microsoft case, for example, Microsoft argued that while as a technological leader it may possess a good deal of static market power, this is merely the fuel for stimulating dynamic R&D competition, a process that it argued works well in the software industry. Antitrust intervention would run the risk of reducing the rate of innovation and welfare. The government argued, instead, that Microsoft's practices prevented entry of new firms and products, and therefore both raised prices and retarded innovation.[2] How to reconcile these two views, however, was never clear in discussions of the case.

On closer inspection, these two conflicting views reveal a fundamental tension in the effects of antitrust policy on innovation. Policies that protect new entrants from incumbents raise a successful innovator's initial profits and may thereby encourage innovation, as the government

* Segal: Stanford University, Department of Economics, 579 Serra Mall, Stanford, CA 94305-6072 (e-mail: ilya.segal@stanford.edu); Whinston: Northwestern University, Department of Economics, 2003 Sheridan Road, Evanston, IL, 60208-2600 (e-mail: mwhinston@northwestern.edu). We thank Shane Greenstein, Ken Hendricks, Volker Nocke, Richard Schmalensee, Scott Stern, and participants in seminars at the Federal Trade Commission, Harvard University, Northwestern University, University of Pennsylvania, UT-Austin, the 2003 NBER Summer Institute meeting on Innovation Policy, the 2003 International Industrial Organization Conference, and the 2004 Duke-Northwestern-Texas IO Theory Conference, for helpful discussions and comments. We also thank the National Science Foundation (SES-0318438) and the Searle Foundation for financial support.

[1] Issues of innovation have been considered when discussing "innovation markets" in some horizontal merger cases in which there was a concern that a merger might reduce R&D competition. See, e.g., Richard J. Gilbert and Steven C. Sunshine (1995).

[2] For further discussion, see Franklin M. Fisher and Daniel L. Rubinfeld (2000), Richard J. Gilbert and Michael L. Katz (2001), Whinston (2001), and Evans and Schmalensee (2002).

argued. But new entrants hope to become the next Microsoft, and will want to engage in the same sorts of entry-disadvantaging behaviors should they succeed. Thus, by lowering the profits of incumbency, protective policies may actually retard innovation, as Microsoft alleged. Disentangling these two effects is the central focus of this paper.

We study the effects of antitrust policy in innovative industries using models in which innovation is a continual process, with new innovators replacing current incumbents, and holding dominant market positions until they are themselves replaced. Although a great deal of formal modeling of R&D races has occurred in the industrial organization literature (beginning with the work of Glenn Loury (1979), and Tom Lee and Louis L. Wilde (1980); see Jennifer Reinganum (1989) for a survey), this work has typically analyzed a single, or at most a finite, sequence of innovative races.[3] Instead, our models are closer to those in the recent literature on growth (e.g., Gene Grossman and Elhanan Helpman 1991; Philippe Aghion and Peter Howitt 1992; and Aghion et al. 2001). The primary distinction between our analysis and this growth literature lies in our explicit focus on how antitrust policies affect equilibrium in such industries.[4]

The paper is organized as follows. In Sections I and II, we introduce and analyze a simple stylized model of antitrust in an innovative industry. Our aim is to develop a model that yields some general insights into the effect of antitrust policies on the rate of innovation, and that we can apply to a number of different antitrust policies. In Section I, we study the simplest version of this model, in which in each period a single potential entrant conducts R&D. The model captures antitrust policy in a reduced form way, by assuming that it alters the profit flows that an incumbent and a new entrant can earn in competition with each other, as well as the profits of an uncontested incumbent. In the model, a more protective antitrust policy—one that protects entrants at the expense of incumbents—increases a new entrant's profits, but also affects the profitability of continuing incumbents. Since successful entrants become continuing incumbents, both of these effects matter for the incentive to innovate.

Using this simple stylized model, we characterize equilibrium in terms of "innovation benefit" and "innovation supply" functions, which provides a very simple approach to comparative statics, and we develop some general insights into the effect of antitrust policies on the rate of innovation. We show that a policy change will increase innovation when a certain weighted sum of the profit changes for a new entrant and a continuing incumbent increase, holding fixed the level of innovation. Those weights imply that a more protective antitrust policy tends to increase the level of innovative activity by "front-loading" an innovative new entrant's profit stream. We also use this condition to provide some general comments on the effects of limits on R&D-deterring activities and voluntary deals, and on how the degree of market growth can alter the effects of antitrust policy. Most significantly, we argue that, in a range of settings, limits on R&D-deterring activities necessarily increase innovation because, holding the rate of innovation fixed (as our comparative statics results instruct us to do), they increase both new entrant and continuing incumbent profits, so the tension identified above is absent.

[3] Three exceptions, however, are Ted O'Donoghue, Suzanne Scotchmer, and Jacques-Francois Thisse (1998) and Robert M. Hunt (2004), who use continuing innovation models to examine optimal patent policies, and Drew Fudenberg and Jean Tirole (2000), who study dynamic limit pricing in markets with network externalities using a model of continuing innovation. Quirmbach (1993) studies the effect of collusion—a change in product market competition that can be addressed in our model (see Section I)—using a model with a one-stage patent race. As a result of this difference, the tension that we focus on differs from the innovation-versus-deadweight loss trade-off that arises in his model.

[4] The growth literature often considers how changes in various parameters will affect the rate of innovation, sometimes even calling such parameters measures of the degree of "antitrust policy" (e.g., Aghion et al. (2001) refer to the elasticity of substitution as such a measure). Here we are much more explicit about what antitrust policies toward specific practices do. This is not a minor difference, as our results differ substantially from those that might be inferred from the parameter changes considered in the growth literature. As an example, one would get exactly the wrong conclusion if one interpreted results showing that more inelastic demand functions lead to more innovation (e.g., Aghion and Howitt 1992) to mean that allowing an incumbent to enhance its market power through long-term contracts leads to more innovation.

In Section II, we extend our comparative statics approach to substantially more general innovation benefit and supply settings. The extension to more general innovation supply, for example, allows us to consider supply settings with N potential entrants, with free entry, or with "free entry with a limited idea" as in work by O'Donoghue, Scotchmer, and Thisse (1998) and Fudenberg and Tirole (2000). We show, using our innovation benefit and supply approach, that, in each case, or in any other setting satisfying several basic properties, the condition characterizing comparative statics is the same.

With the comparative statics results of Section II in hand, in Section III we provide two illustrations of the use of our framework. First, we study a model of long-term (exclusive) contracts and show that a more protective antitrust policy necessarily stimulates innovation, and also raises both aggregate and consumer welfare. Then, we consider a model of compatibility choice in an industry characterized by network externalities. Here, we identify cases in which innovation necessarily increases when incumbents are forced to make their products more compatible with those of future entrants, as well as cases in which innovation may decline.

The analysis of Sections I–III makes the strong assumption that only potential entrants conduct R&D. While useful for gaining understanding, this assumption is rarely descriptive of reality. In Section IV, we turn our attention to models in which both incumbents and potential entrants conduct R&D. Introducing incumbent investment has the potential to complicate our analysis substantially by making equilibrium behavior depend on the level of the incumbent's lead over other firms. We study two models in which we can avoid this state dependence. Interestingly, in both models there is a wide range of circumstances in which a more protective policy can increase the innovation incentives of *both* the incumbent and potential entrants.

Sections I–IV focus on antitrust policies that alter entrant and incumbent profit flows, shifting only innovation benefit. In Section V, we consider antitrust policies toward two other types of behaviors, predatory activities and actions that shift innovation supply.

Section VI concludes and discusses the relation of our analysis to issues in intellectual property protection, where some similar issues arise.

Proofs of all lemmas and propositions appear in the Appendix.

I. A Stylized Model of Antitrust in Innovative Industries

We begin by developing a stylized model of continuing innovation. Our aim is to develop a model that yields some general insights into the effect of antitrust policies on the rate of innovation, and that we can apply to a number of different antitrust policies in the remainder of the paper. In this section, we develop the simplest possible version of this model, which we substantially generalize in the next section.

The model has discrete time and an infinite horizon. There are two firms that discount future profits at rate $\delta \in (0,1)$. In each period, one of the firms is the "incumbent" I and the other is the "potential entrant" E. In the beginning of each period, the potential entrant chooses its R&D rate, $\phi \in [0,1]$; the cost of R&D is given by a twice differentiable function $c(\phi)$, with $c'(\cdot) > 0, c''(\cdot) > 0$.[5] The R&D of the potential entrant yields an innovation—which we interpret to be an improvement in the quality of the product—with probability ϕ. If the potential entrant innovates, it receives a patent, enters, competes with the incumbent in the present period, and then becomes the incumbent in the next period, while the previous incumbent then becomes the potential entrant. In this sense, this is a model of "winner-take-all" competition. While the patent provides perfect protection (forever) to the innovation itself, the other firm may overtake the patent holder by developing subsequent innovations.

Antitrust policies can have an impact on incentives for innovation in various ways. Throughout most of the paper, we will be interested in the effects of an antitrust policy α that affects the incumbent's competition with an entrant who has just received a patent. Many antitrust policies are of this type, and we will analyze two of these in detail in Section III. (In Section V we discuss policies that have other effects.) To this end, we denote the incumbent's profit in competition

[5] Note that $c(\cdot)$ must be convex if the entrant can randomize over its R&D strategy. We assume *strict* convexity and twice-differentiability to simplify the exposition.

with a new entrant by $\pi_I(\alpha)$, and the profit of the entrant by $\pi_E(\alpha)$, which we assume are differentiable functions of α. We let $\pi'_E(\alpha) > 0$, so that a higher α represents a policy that is more "protective" of the entrant in the sense that it raises the profit of the entrant in the period of entry. Less clear, however, is the overall effect of an increase in α on the incentive to innovate, since an increase in α also alters the value of becoming a continuing incumbent. We also denote by the differentiable function $\pi_m(\alpha)$ the profit of an incumbent that faces no competition in a period. (This may depend on α if the antitrust policy restricts an uncontested incumbent's behavior.) In Section III, when we consider specific applications, we show how these values can be derived from an underlying model of the product market.

We examine stationary Markov perfect equilibria of the infinite-horizon game using the dynamic programming approach. Let V_I denote the expected present discounted profit of an incumbent, and V_E that of a potential entrant (both evaluated at the beginning of a period). Then, since innovation occurs with probability ϕ, these values should satisfy

$$\text{(VI)} \quad V_I = \pi_m(\alpha) + \delta V_I + \phi[\pi_I(\alpha) - \pi_m(\alpha) + \delta(V_E - V_I)];$$

$$\text{(VE)} \quad V_E = \delta V_E + \phi[\pi_E(\alpha) + \delta(V_I - V_E)] - c(\phi).$$

Also, a potential entrant's choice of ϕ should maximize its expected discounted value. Letting $w \equiv \pi_E(\alpha) + \delta(V_I - V_E)$ denote the expected discounted benefit from becoming a successful innovator—what we shall call the *innovation prize*—the optimal innovation level given innovation prize w is

$$\Phi(w) = \arg\max_{\phi \in [0,1]} \{\phi w - c(\phi)\}.$$

Under our assumptions on the cost function $c(\cdot)$, $\Phi(w)$ can be described as the unique solution ϕ to the entrant's first-order condition $c'(\phi) = w$ for all $w \geq c'(0)$, and the Implicit Function Theorem implies that $\Phi(w)$ is a continuous and increasing function of w in this range. For $w < c'(0)$, $\Phi(w) = 0$. The graph of $\Phi(\cdot)$ gives us an "Innovation Supply" (IS) curve, as shown in Figure 1.

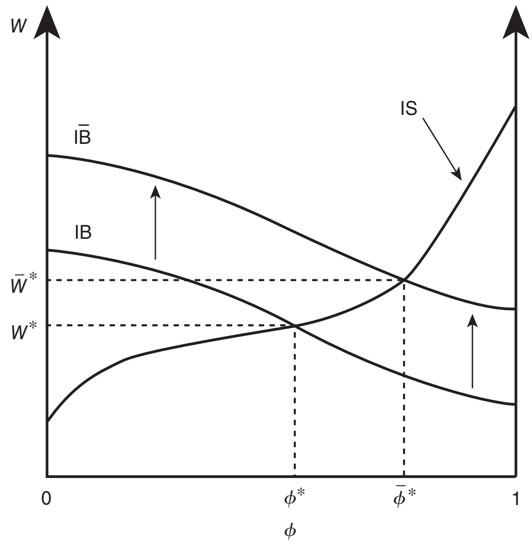

FIGURE 1. COMPARATIVE STATICS

Consider now the determinants of the innovation prize w. Subtracting (VE) from (VI), solving for $(V_I - V_E)$, and substituting its value into $w \equiv \pi_E + \delta(V_I - V_E)$ allows us to express the innovation prize as $w = W(\phi, \alpha)$, where

$$\text{(IB)} \quad W(\phi,\alpha) = \pi_E(\alpha) + \delta \times \frac{\left\{\begin{array}{l}\phi\pi_I(\alpha) + (1-\phi)\pi_m(\alpha) \\ \qquad - \phi\pi_E(\alpha) + c(\phi)\end{array}\right\}}{1 - \delta + 2\delta\phi}$$

$$= \frac{\left\{\begin{array}{l}[1-\delta(1-\phi)]\ \pi_E(\alpha) \\ + \delta[\phi\pi_I(\alpha) + (1-\phi)\ \pi_m(\alpha) + c(\phi)]\end{array}\right\}}{1 - \delta + 2\delta\phi}.$$

This equation defines the "Innovation Benefit" (IB)—the value of the innovation prize as a function of the innovation rate ϕ. Figure 1 graphs the IB curve along with the IS curve.

To simplify exposition in this section, we assume that the IB and IS curves intersect once, so there are unique equilibrium values (ϕ^*, w^*). In fact, existence of equilibrium is assured by the continuity of the functions $\Phi(\cdot)$ and $W(\cdot)$, while uniqueness is ensured in this simple model whenever $\pi_m(\alpha) \geq \pi_I(\alpha) + \pi_E(\alpha)$, so that industry profits are larger with an uncontested incumbent than when entry occurs (Jean Tirole (1988) calls this the "efficiency effect").[6, 7]

Note that the IS curve does not depend on α at all. As Figure 1 shows (and is established formally in the next section), if α shifts the IB curve up (down) at all values of ϕ, the equilibrium innovation rate increases (decreases). Differentiating $W(\phi, \alpha)$ with respect to α, we see that the protectiveness of antitrust policy increases (decreases) innovation if, for all $\phi \in [0,1]$,

$$(1) \quad \pi'_E(\alpha) + \delta\left[\frac{(1-\phi)\pi'_m(\alpha) + \phi\pi'_I(\alpha)}{1-\delta(1-\phi)}\right] \geq (\leq)\ 0.$$

Condition (1) indicates how to sort through the potentially conflicting effects of antitrust policy on innovation incentives that arise from the policy's dual effects on a successful entrant's initial profits and its returns from achieving incumbency. It shows that a change in policy encourages (discourages) innovation precisely when it raises (lowers) the incremental expected discounted profits over an innovation's lifetime. To see this, observe that the first term on the left side of (1) is the change in an entrant's profit in the period of entry due to the policy change. The second term equals the change in the value of a continuing incumbent (the numerator is the derivative of the flow of expected profits in each period of incumbency conditional on still being an incumbent; the denominator captures the "effective" discount rate, which includes the probability of displacement), and thus the change in the entrant's value once it is itself established as the incumbent.

In interpreting condition (1), it is helpful to think about the case in which the monopoly profit π_m is independent of the antitrust policy α, so that $\pi'_m(\alpha) = 0$. In this case, condition (1) tells us that innovation increases (decreases) if

$$(2) \quad \pi'_E(\alpha) + \left[\frac{\delta\phi}{1-\delta(1-\phi)}\right]\pi'_I(\alpha) \geq (\leq)\ 0.$$

Thus, innovation increases if a weighted sum of $\pi'_E(\alpha)$ and $\pi'_I(\alpha)$ increases, where the weight on $\pi'_E(\alpha)$ exceeds the weight on $\pi'_I(\alpha)$ due to discounting ($\delta < 1$). This implies that a more protective antitrust policy raises innovation whenever $\pi'_I(\alpha) + \pi'_E(\alpha) \geq 0$, that is, provided that an increase in α does not lower the joint profit of the entrant and the incumbent in the period of entry. Intuitively, observe that a successful innovator earns $\pi_E(\alpha)$ when he enters, and earns $\pi_I(\alpha)$ when he is displaced. A more protective antitrust policy that raises π_E and lowers π_I has a *front-loading effect*, effectively shifting profits forward in time. Since the later profits π_I are discounted, this front loading of profits increases the innovation prize, provided that the joint profit $\pi_I + \pi_E$ does not decrease.

Observe also that the weight on $\pi'_I(\alpha)$ is increasing in δ. Thus, the larger is δ, the more likely it is that a more protective policy reduces innovation. This is because with a larger δ the

[6] Indeed, suppose in negation we have two equilibrium innovation values, ϕ and ϕ', with $\phi' > \phi$. Since both are on the upward-sloping IS curve, equilibrium ϕ' has a larger innovation prize $w = \pi_E(\alpha) + \delta(V_I - V_E)$, and hence a larger value of $V_I - V_E$. Also, adding together (VI) and (VE), we see that equilibrium ϕ' has a lower value of $V_I + V_E$, since it has lower total market profits by the "efficiency effect," and has higher R&D costs incurred by the entrant. This implies that equilibrium ϕ' has a lower value of V_E. But this contradicts the fact that by (VE) and the entrant's R&D optimization, $(1-\delta)\ V_E = \max_{\phi\in[0,1]}[\phi w - c(\phi)]$, and so V_E must be larger when the innovation prize w is larger. When the efficiency effect holds, we can also show that the IB curve is downward sloping whenever it is above the IS curve. Note that the IB curve need not be downward sloping *everywhere*: e.g., for a high innovation level ϕ at which the marginal cost $c'(\phi)$ is very high, raising ϕ would reduce V_E by a lot, thus increasing the innovation prize.

[7] The condition says that a monopoly maximizes industry profits. The usual justification for the efficiency effect is that a monopolist can always replicate an entrant's behavior. However, this justification is not valid in our model, where a monopolist does not have access to the entrant's better product. At the end of Section IIIA, we discuss a case where this condition is violated.

discounted value of the profits in the period in which the entrant is displaced is greater. In the limit, as $\delta \to 1$, the amount by which the joint profit $\pi_E + \pi_I$ can be dissipated while still encouraging innovation converges to zero: the cost of a one dollar reduction in the value π_I that the entrant will receive when he is ultimately displaced is exactly equal to the gain from receiving a dollar more in the period in which he enters.

Condition (1) also provides insights into several specific settings of interest, as we now discuss.

A. *R&D-Deterring Activities*

We began by noting an inherent tension in the effect of more protective policies on innovation. Surprisingly, though, when we consider policies that restrict an incumbent's ability to deter entrants' R&D, this tension is often absent, at least once we hold the rate of innovation fixed, as (1) suggests we do.

To see why, consider the following stylized model: imagine that at the end of each period the firm with the leading technology can commit to some behavior $b \in \mathbb{R}$ that affects both its profit and an entrant's profit in the following period. Let these profits be given by functions $\hat{\pi}_m(b)$, $\hat{\pi}_I(b)$, and $\hat{\pi}_E(b)$, where $\hat{\pi}_E'(b) < 0$.

Consider a unique stationary Markov perfect equilibrium in which potential entrants' R&D in period t is a decreasing differentiable function $\phi^*(b_t)$ of the state variable b_t chosen in the previous period. Then, the equilibrium level of b absent any antitrust constraint, b^*, maximizes

$$\phi^*(b)[\hat{\pi}_I(b) + \delta V_E] + (1 - \phi^*(b))[\hat{\pi}_m(b) + \delta V_I], \tag{3}$$

and satisfies the first-order condition

$$[\phi\hat{\pi}_I'(b^*) + (1 - \phi)\hat{\pi}_m'(b^*)] + \phi^{*\prime}(b^*)\{[\hat{\pi}_I'(b^*) + \delta V_E] - [\hat{\pi}_m(b^*) + \delta V_I]\} = 0.$$

If it is valuable to the incumbent to prevent entry, so that $\{[\hat{\pi}_I(b^*) + \delta V_E] - [\hat{\pi}_m(b^*) + \delta V_I]\} \leq 0$, then $[\phi\hat{\pi}_I'(b^*) + (1 - \phi)\hat{\pi}_m'(b^*)] \leq 0$. But this implies that a small reduction in b, *holding ϕ fixed*, actually *raises* the profit of a continuing incumbent. If the antitrust policy α requires that $b \leq b^* - \alpha$, then for small α which lead to small reductions in b no tension arises: *both* terms in condition (1) are nonnegative.[8] We will see an example of this type in Section III.

The key insight here comes because of condition (1), which instructs us to think about profit effects holding the rate of innovation ϕ fixed. If (on the margin) an R&D-deterring activity involves a sacrifice in incumbent profit in return for a reduction in the probability of entry, then, *holding ϕ fixed*, both entrant and incumbent profits are increased by a slightly more protective policy, and so innovation increases. Interestingly, the condition that a restriction on an incumbent's action b raises his expected profit holding ϕ fixed is similar to the idea behind Janusz A. Ordover and Robert D. Willig's (1981) definition of predatory behavior.[9] Thus, prohibiting actions that are predatory in this sense raises the rate of innovation.

B. *Voluntary Deals*

Although we have motivated our analysis by discussing examples of exclusionary behaviors, our framework is not limited to such applications. Another sort of behavior to which we can apply our model is a voluntary deal between an incumbent and a new entrant. For example, an incumbent might license a new entrant's technology to serve his captive customers in the long-term contracting model we study in Section IIIA, or the two firms might collude in their pricing to consumers. With the innovation rate held fixed, such voluntary deals—by definition—raise both parties' payoffs.[10] By condition (1), such deals should therefore increase equilibrium innovation.

[8] Formally, taking $\pi_j(\alpha) \equiv \hat{\pi}_j(b^* - \alpha)$ for $j = m, I, E$, we have $\phi\hat{\pi}_I'(\alpha) + (1 - \phi)\hat{\pi}_m'(\alpha) \geq 0$.

[9] Ordover and Willig explicitly restrict their definition to actions that eliminate existing rivals. Here we extend their notion to actions that reduce the likelihood of entry.

[10] This conclusion depends on the fact that we have only two parties negotiating. With more than two active firms, profits for some or all parties may fall when voluntary deals are allowed (see Segal 1999).

C. *Market Growth*

Up to this point (and in the remainder of the paper), we have focused on a stationary setting. The front-loading feature of protective antitrust policies suggests, however, that the rate of market growth may alter the impact of antitrust policy on innovation. As an illustration, imagine that profits in period 1 are instead $\beta\pi_I(\alpha)$, $\beta\pi_E(\alpha)$, and $\beta\pi_m(\alpha)$ for $\beta \leq 1$, and are $\pi_I(\alpha)$, $\pi_E(\alpha)$, and $\pi_m(\alpha)$ from period 2 on. Following a derivation similar to that above, an increase in α will now increase (decrease) the period 1 innovation rate ϕ_1 if

$$\beta\pi'_E(\alpha) + \delta\left[\frac{(1-\phi)\pi'_m(\alpha) + \phi\pi'_I(\alpha)}{1-\delta(1-\phi)}\right] \geq (\leq)\ 0.$$

Thus, the greater is market growth (the lower is β), the less likely it is that period 1 innovation will increase. For example, when $\beta < 1$ and $\pi'_m(\alpha) = 0$, an increase in α may reduce innovation even when joint profits upon entry increase $[\pi'_I(\alpha) + \pi'_E(\alpha) \geq 0]$.

II. Comparative Statics for More General Innovation Supply and Benefit

In this section, we extend our comparative statics result to more general innovation supply and benefit settings. For one thing, we can generalize the innovation benefit by allowing all three of the profits π_I, π_E, and π_m to be affected by the equilibrium rate of innovation ϕ. (This may happen because the price at which consumers purchase a durable good or accept a long-term contract may depend on their expectation of the innovation rate, as in the example studied in Section IIIA.) Denoting these profits by $\pi_I(\alpha,\phi)$, $\pi_E(\alpha,\phi)$, and $\pi_m(\alpha,\phi)$, we see that the argument of Section I continues to hold if we reinterpret the derivatives in (1) as being partial derivatives with respect to α holding ϕ fixed.

We can also allow for alternative models of innovation supply, such as having more than one potential entrant engage in R&D, or even allowing free entry (i.e., infinitely many potential entrants). (We still do not allow the incumbent to conduct R&D; we relax this assumption in Section IV.) To do so, we define the industry's "innovation rate" ϕ as the probability that the incumbent technology is displaced with an innovation. For a symmetric industry, ϕ also determines the potential entrants' *individual* R&D investments (and, hence, their R&D cost $c(\phi)$) and the probability $u(\phi)$ that a *given* potential entrant becomes a new incumbent (i.e., moves "up"). The expected present discounted profits of an incumbent (V_I) and a potential entrant (V_E) can then be described by

$$\text{(VI*)}\qquad V_I = \pi_m(\alpha) + \delta V_I + \phi[\pi_I(\alpha) - \pi_m(\alpha) + \delta(V_E - V_I)],$$

$$\text{(VE*)}\qquad V_E = \delta V_E + u(\phi)[\pi_E(\alpha) + \delta(V_I - V_E)] - c(\phi).$$

Subtracting, we can express the innovation prize $w = \pi_E + \delta(V_I - V_E)$ with the following function:

$$\text{(IB*)}\qquad W(\phi,\alpha) = \frac{\left\{\begin{aligned}&[1-\delta(1-\phi)]\ \pi_E(\alpha,\phi)\\ &\quad + \delta[\phi\pi_I(\alpha,\phi)\\ &\quad + (1-\phi)\ \pi_m(\alpha,\phi) + c(\phi)]\end{aligned}\right\}}{1-\delta+\delta(\phi+u(\phi))}$$

As for the innovation supply curve, describing entrants' R&D for a given innovation prize, our comparative statics results hold as long as the curve is described by a nondecreasing continuous function $\Phi(w)$.[11] In a model of such generality, uniqueness of equilibrium can no

[11] In the working paper version of this article (Segal and Whinston 2005), we showed that the results hold for more general settings in which the innovation supply curve is described by a correspondence $\Phi(\cdot)$, provided that it (a) is nonempty- and convex-valued, (b) has a closed graph, and (c) has the property that any selection from it is nondecreasing. For example, if the entrant in the simple model has, instead, a linear R&D cost function $c(\phi) = c\phi$ with $c > 0$, we have $\Phi(w) = 0$ for $w < c$, $\Phi(w) = [0,\infty]$ if $w = c$, and

longer be assured, but we can still formulate comparative statics results for the "largest" and "smallest" equilibria, depicted in Figure 2 as $\underline{\phi}$ and $\overline{\phi}$, respectively, using the results of Paul Milgrom and John Roberts (1994). Those results imply that antitrust policy affects the largest and smallest equilibrium innovation rates in the same direction that it shifts the innovation benefit curve, which we can determine by partially differentiating $W(\phi, \alpha)$.

PROPOSITION 1: *If the innovation supply function $\Phi(\cdot)$ is continuous and nondecreasing, and the innovation benefit function $W(\phi, \alpha)$ is continuous in ϕ, then the largest and smallest equilibrium innovation rates exist, and both these rates are nondecreasing (nonincreasing) in α, the protectiveness of antitrust policy, if*

$$\frac{\partial \pi_E(\alpha,\phi)}{\partial \alpha} + \frac{\delta}{1-\delta(1-\phi)} \times \left[(1-\phi)\,\frac{\partial \pi_m(\alpha,\phi)}{\partial \alpha} + \phi \frac{\partial \pi_I(\alpha,\phi)}{\partial \alpha}\right] \geq (\leq)\, 0 \tag{4}$$

for all $\phi \in [0,1]$.

In what follows, we apply this proposition by saying that a change in policy "increases (decreases) innovation" whenever the largest and smallest equilibrium innovation rates exist and are both nondecreasing (nonincreasing) in α. When there is a unique equilibrium, this implies determinate comparative statics. If there are multiple equilibria, some of them may move in the direction opposite that predicted by Proposition 1. As suggested by Figure 2, however, the same local comparative statics must hold for any locally unique equilibrium $\phi(\alpha)$ at which we have *crossing from above*, i.e., for some interval $[\underline{\phi}, \overline{\phi}]$, $\Phi(W(\alpha,\phi)) \subset [\phi, 1]$ for $\phi \in [\underline{\phi}, \phi(\alpha)]$ and $\Phi(W(\alpha,\phi)) \subset [0,\phi]$ for $\phi \in [\phi(\alpha), \overline{\phi}]$, which is a necessary condition for (Lyapunov) stability.[12] Proposition 2 states this conclusion.

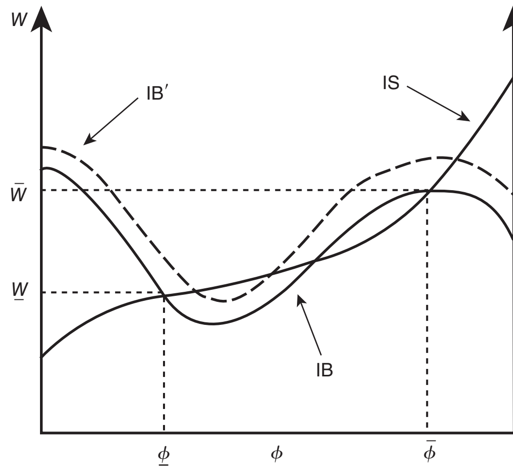

FIGURE 2. EQUILIBRIUM COMPARATIVE STATICS: THE GENERAL CASE

PROPOSITION 2: *Suppose that the innovation supply function $\Phi(\cdot)$ is continuous and nondecreasing, and the innovation benefit function $W(\phi, \alpha)$ is continuous in ϕ. Suppose, in addition, that for any $\alpha \in [\underline{\alpha}, \overline{\alpha}]$ there is a unique equilibrium innovation rate $\phi(\alpha)$ on an interval $[\underline{\phi}, \overline{\phi}]$ and that the IB curve crosses the IS curve from above on this interval. Then, $\phi(\alpha)$ is nondecreasing (nonincreasing) if condition (4) holds for all $\alpha \in [\underline{\alpha}, \overline{\alpha}]$ and $\phi \in [\underline{\phi}, \overline{\phi}]$.*

We now provide three examples in which the innovation supply function is continuous and nondecreasing, so that Propositions 1 and 2 determine the comparative statics effect of antitrust policy on the rate of innovation. All three examples also have the feature that $\Phi(0) = 0$, which we also assume from now on.

A. *More Than One Potential Entrant*

Suppose that there are $N > 1$ potential entrants in any period (in addition to the single incumbent). In the beginning of each period, each potential entrant i independently chooses its R&D rate $\psi_i \in [0,1]$; the cost of R&D is

$\Phi(w) = 1$ if $w > c$. This correspondence satisfies (a)–(c), so all of our results would apply in this setting.

[12] Specifically, suppose that the industry in period t adaptively expects the future innovation rate to be ϕ_{t-1}, and so chooses innovation rate $\phi_t \in \Phi(W(\alpha, \phi_{t-1}))$. If we have "crossing from below," then the dynamics starting at any disequilibrium $\phi \in [\underline{\phi}, \overline{\phi}]$ must leave the interval $[\underline{\phi}, \overline{\phi}]$, so that the equilibrium is not (Lyapunov) stable.

given by a strictly convex function $\gamma(\psi_i)$, with $\gamma(0) = 0$.[13] The R&D of a given potential entrant i yields a discovery with probability ψ_i (we assume that the discoveries are independently realized). We shall focus on symmetric equilibria, in which all potential entrants choose the same level of R&D, denoted by ψ. In this case, the probability that at least one of the N potential entrants makes a discovery is given by $\phi = [1 - (1 - \psi)^N]$. Thus, in a symmetric equilibrium with aggregate innovation rate ϕ, the potential entrants' individual R&D choices are $\psi_N(\phi) = 1 - (1 - \phi)^{1/N}$.

Among the potential entrants who make a discovery, only one may patent it. Denote by $r_N(\psi)$ the probability that a given potential entrant receives a patent, conditional on it making a discovery, when all other potential entrants are doing R&D at level ψ. We assume that $r_N(\cdot)$ is a strictly decreasing continuous function.[14] A potential entrant who receives a patent enters and competes with the incumbent in the present period, and becomes the incumbent in the next period, while the previous incumbent then becomes a potential entrant.

Note that in a symmetric equilibrium, the probability that a given entrant becomes the incumbent is $u(\phi) = \psi_N(\phi)\, r_N(\psi_N(\phi))$, and the entrant's R&D cost is $c(\phi) = \gamma(\psi_N(\phi))$. The innovation benefit curve is then given by substituting these expressions in (IB*) above.

As for innovation supply, the equilibrium individual innovation rate solves the following equilibrium condition for a given value of w:

$$\text{(5)} \quad \psi = \arg\max_{\psi' \in [0,1]} \{\psi' r_N(\psi)\, w - \gamma(\psi')\}.$$

Lemma A1 in the Appendix shows that this condition describes a unique equilibrium level of ψ, which is a continuous and nondecreasing function of w. As a result, the *aggregate* innovation supply $\Phi(w) = [1 - (1 - \Psi(w))^N]$ is also a continuous and nondecreasing function of w, so Propositions 1 and 2 apply to this model.

B. *Free Entry*

In some circumstances, it may be more appropriate to assume that there is free entry into R&D competition.[15] This assumption can be interpreted as a limiting case of a very large number N of potential entrants, each of whom engages in an infinitesimal amount of R&D while the aggregate innovation rate is positive. If a patent is awarded randomly to one of the successful innovators, an innovator's conditional probability of getting a patent can then be written using the expression in footnote 14 as

$$\begin{aligned} \bar{r}(\phi) &= \lim_{N\to\infty} r_N(\psi_N(\phi)) \\ &= \lim_{N\to\infty} \frac{\phi}{(1 - (1 - \phi)^{1/N})\, N} \\ &= -\frac{\phi}{\ln(1 - \phi)}, \end{aligned}$$

which is a continuous decreasing function of $\phi \in [0, 1]$. From (5), the first-order condition for each potential entrant to choose a positive infinitesimal level of R&D is $w\bar{r}(\phi) = \gamma'(0)$, which determines the innovation supply function $\Phi(w) = \bar{r}^{-1}(\gamma'(0)/w)$ for $w > \gamma'(0)$, and $\Phi(w) = 0$ for $w \le \gamma'(0)$. Since this is a continuous and nondecreasing function, the comparative statics is again described by Propositions 1 and 2, where in (IB*) we take $c(\phi) = u(\phi) = 0$ since each potential entrant in expectation does not incur any cost and has zero chance of innovating.

[13] In the working paper version of this article (Segal and Whinston 2005), we show that the same result holds for cost functions $\gamma(\cdot)$ that are only weakly convex, allowing, e.g., for linear innovation technology.

[14] When the patent is awarded randomly to one of the successful innovators,

$$\begin{aligned} r_N(\psi) &= \sum_{k=0}^{N-1} \left(\frac{1}{k+1}\right) \binom{N-1}{k} \psi^k (1 - \psi)^{N-1-k} \\ &= \frac{1 - (1 - \psi)^N}{\psi N}. \end{aligned}$$

[15] The fixed N model is the appropriate model when there are a limited number of firms with the capability of doing R&D in an industry (perhaps because of complementary assets they possess due to participation in related industries).

C. *Free Entry with a Randomly Arriving Idea*

In the models of O'Donoghue, Scotchmer, and Thisse (1998) and Fudenberg and Tirole (2000), there are infinitely many potential entrants, but in each period only one of them is randomly drawn to receive an "idea." This potential entrant observes a randomly drawn implementation cost γ and chooses whether to invest in implementing the innovation. If he does, he is certain to become the next incumbent.

Observe first that a potential entrant's optimal strategy takes the form of choosing a cost threshold $\bar{\gamma}$ below which to implement their idea. This is equivalent to choosing the probability of innovation $\phi = \Pr\{\gamma \leq \bar{\gamma}\} \equiv F(\bar{\gamma})$ at an expected cost of $\int_0^{F^{-1}(\phi)} \gamma\, dF(\gamma)$, which is a strictly convex function of ϕ provided that $F(\cdot)$ is strictly increasing. The innovation supply function is therefore determined just as in the single-firm model of Section I, so it is continuous and nondecreasing. For the innovation benefit curve, on the other hand, in (IB*) we take $c(\phi) = u(\phi) = 0$ since each potential entrant in expectation does not incur any cost and has zero chance of innovating. Once again, the comparative statics are described by Propositions 1 and 2.

III. Applications

In this section, we use the results of Section II to analyze two models in which an incumbent engages in activities designed to deter the R&D of potential entrants. In these applications, we derive the relevant profit functions from fundamentals, and there is a fully specified consumer side so that welfare analysis is possible. Both models are versions of the "quality ladder" models introduced in the recent literature on economic growth (e.g., Grossman and Helpman 1991; Aghion and Howitt 1992). Our results in this section hold for any continuous nondecreasing innovation supply function.

A. *Long-Term (Exclusive) Contracts*

We first consider a model in which the incumbent can sign consumers to long-term contracts. There are at least two firms and a continuum of infinitely lived consumers of measure 1 who may consume a nonstorable and nondurable good with production cost $k \geq 0$. R&D may improve the quality of this good, and consumers value "generation j" of the good at $v_j = v + j \cdot \Delta$. At any time t, one firm—the current "incumbent"—possesses a perfectly effective and infinitely lived patent on the latest generation product j_t. Likewise, at time t there is a patent holder for each of the previous generations of the product $(j_t - 1, j_t - 2, \ldots)$. We assume, as in Sections I and II, that at time t only firms other than the incumbent in the leading technology—the potential entrants—can invest in developing the generation $j_t + 1$ product. One implication of this assumption is that in each period t the holder of the patent on generation $j_t - 1$ is a firm other than the current incumbent, which holds the patent on the current leading generation j_t.

Suppose that in each period t, the incumbent can offer long-term contracts to a share b_{t+1} of consumers. The contracts specify a sale in period $t + 1$ at a price q_{t+1} to be paid upon delivery. (In our simple model, this is equivalent to an exclusive contract that prevents the consumer from buying from the entrant, subject to some irrelevant issues with the timing of payments.) The antitrust policy α restricts the proportion of consumers that can be offered long-term contracts not to exceed $1 - \alpha$. We assume that the production cost k exceeds the quality increment Δ, so that an entrant cannot profitably make a sale to a consumer who is bound to a long-term contract.

The timing in period t is:

- **Stage** t.1: Each potential entrant i observes the previous history, including the share of captured consumers B_t, and chooses its innovation rate ψ_{it}. Then innovation success is realized.
- **Stage** t.2: Firms name prices p_{it} to free consumers. We assume that firms do not name below-cost prices.
- **Stage** t.3: Free consumers accept/reject these offers.
- **Stage** t.4: The firm with the leading technology chooses to offer to a share $b_{t+1} \leq 1 - \alpha$ of consumers a period $t + 1$ sales contract at price q_{t+1} to be paid upon delivery. These offers are observable.
- **Stage** t.5: Consumers accept/reject these contract offers (each assumes that her decision

has no effect on the likelihood of future entry). The share of consumers accepting these contracts is $B_{t+1} \leq b_{t+1}$, and is observable. We assume throughout that consumers all accept if accepting is a continuation equilibrium (we do not allow consumers to coordinate).[16]

Observe that when $\alpha = 1$ so that no long-term contracts can be written, we have a simple model with Bertrand competition in each period between the leading firm and firms further down the ladder. In the Web Appendix (available at http://www.e-aer.org/data/dec07/20050026_app.pdf), we discuss this benchmark quality ladder model, and the distortions in the rate of innovation relative to the first-best level that arise. In general, innovation may be either insufficient or excessive because of "Schumpeterian" and "business stealing" effects.

Here we focus on Markov perfect equilibria, in which potential entrants in stage t.1 condition their innovation choices only on the current share of captive customers B_t, and in which choices at all other stages are stationary. (Note that since period t contracts expire at the end of that period, there is no relevant state variable affecting the contracting choice of the leading firm at stage t.4. At stage t.5, the maximal price at which consumers are willing to accept a long-term contract depends on the share of consumers offered that contract, b_{t+1}, but does not depend on the state at the start of the period, B_t.) In the Appendix we present a detailed formal discussion of these equilibria. Here we provide a more informal discussion aimed at conveying the key ideas and conclusions.

In a Markov perfect equilibrium, on the equilibrium path the state variable B_t is the same at the start of every period (and after any sequence of innovations). Thus, the value of being an incumbent, the value of being an entrant and the rate of innovation are also the same in every period on the equilibrium path. As a result, we can use our model of innovation in Section II to determine the equilibrium rate of innovation once we have determined the equilibrium profit flows.

Suppose that the share of customers signing long-term contracts on the equilibrium path is B^*. It is immediate that in any such equilibrium, the prices offered to free consumers in any period t are $k + \Delta$ by the firm with the leading technology j_t, which wins the sale, and k by the firm with technology $j_t - 1$. So the leading firm in a period, whether a continuing incumbent or an entrant, earns $(1 - B^*)\Delta$ from sales to free consumers.

Now, consider a consumer's decision of whether to accept a long-term contract. If in period t the expected probability of entry in period $t + 1$ is ϕ_{t+1}, a consumer who rejects the leading firm's long-term contract offer for period $t + 1$ anticipates getting the period t surplus level $v + (j_t - 1)\Delta - k$ plus an expected gain in surplus of $\phi_{t+1}\Delta$ due to the possibility of technological advancement in period $t + 1$. Thus, he will accept the contract if and only if the price q_{t+1} satisfies $v + j_t\Delta - q_{t+1} \geq v + (j_t - 1 + \phi_{t+1})\Delta - k$. Hence, the maximum price the incumbent can receive in a long-term contract is $q_{t+1} = k + (1 - \phi_{t+1})\Delta$, which leaves the consumer indifferent about signing. Therefore, in a steady-state equilibrium with innovation rate ϕ, the incumbent seller earns $B^*(1 - \phi)\Delta$ from the consumers who sign long-term contracts.

In the proof of Proposition 3 in the Appendix, we show that this model has a Markov perfect equilibrium in which the antitrust constraint binds, so that $B^* = 1 - \alpha$, and from now on we focus on such equilibria. Furthermore, the antitrust constraint must bind in *any* Markov perfect equilibrium as long as the innovation supply function $\Phi(\cdot)$ is strictly increasing.[17] To see why, observe that—ignoring any effects on the probability of entry—the leading firm is indifferent about signing up an extra consumer at the price $q_{t+1} = k + (1 - \phi_{t+1})\Delta$ when period $t + 1$'s innovation rate is ϕ_{t+1}: its period t expectation of the profit from a free consumer in period $t + 1$ is $(1 - \phi_{t+1})\Delta$, which exactly equals its expected profit from a long-term contract. However, if $\Phi(\cdot)$ is strictly increasing, then ϕ_{t+1} is strictly decreasing in B_{t+1}, because

[16] The leading firm could achieve this by, for example, committing to auction off the desired number of long-term contracts.

[17] This is ensured, for example, in the innovation supply model with $N \geq 1$ potential entrants whose R&D cost function is differentiable and satisfies the Inada conditions $c'(0) = 0$ and $c'(\psi) \rightarrow +\infty$ as $\psi \rightarrow 1$.

an increase in the share of captive consumers reduces the profits a successful entrant can collect in period $t + 1$ (we show this formally in the Appendix). Since the incumbent benefits from a lower probability of entry, he optimally signs up as many long-term customers as the antitrust constraint allows.

Putting these points together, we see that we can fit this model into our basic model by letting

$$(6)\quad \pi_m(\alpha,\phi) = \alpha\Delta + (1-\alpha)(1-\phi)\Delta;$$

$$\pi_I(\alpha,\phi) = (1-\alpha)(1-\phi)\Delta;$$

$$\pi_E(\alpha,\phi) = \alpha\Delta.$$

How does a change in the allowed share of long-term (exclusive) contracts, α, affect the rate of innovation? Observe, first, that in this model an increase in α does indeed raise π_E. More significantly, the expected profit of a continuing incumbent is

$$[\phi\pi_I(\alpha,\phi) + (1-\phi)\pi_m(\alpha,\phi)]$$

$$= (1-\phi)\alpha\Delta + (1-\alpha)(1-\phi)\Delta$$

$$= (1-\phi)\Delta.$$

Thus, holding ϕ fixed, the expected profit of a continuing incumbent is *independent of* α. The reason is that, holding ϕ fixed, it is a matter of indifference to the incumbent (and consumers) whether consumers accept a long-term contract. This implies that there is no conflict—holding ϕ fixed—between the effects of a more protective policy on entrant profit and on the expected profit of a continuing incumbent.[18] Since condition (4) is then necessarily satisfied, we have the following proposition.

PROPOSITION 3: *In the Markov perfect equilibria of our model of long-term (exclusive) contracts in which the restriction on the use of long-term (exclusive) contracts binds (which is true in all Markov perfect equilibria if the innovation supply function is strictly increasing), tightening this restriction increases innovation.*

Consider now the welfare effects of a once-and-for-all increase in the policy α in some period τ. We assume that this intervention occurs just after stage $\tau.1$.[19] The payoff effects of such a change begin in period $\tau + 1$. Note, first, that the increase raises consumer surplus: consumers are indifferent about signing exclusives when the innovation rate is held fixed, and an increase in the innovation rate delivers to free consumers higher-quality goods at the same prices. What about the sum of consumer surplus and current incumbent (i.e., the firm with the leading technology just after stage $\tau.1$) profits? There is no direct effect of the policy change on either consumers or the current incumbent: holding ϕ fixed, both are indifferent about whether long-term contracts are signed. What about the indirect effect caused by the increase in ϕ? Intuitively, an innovation in period $\tau + 1$ reallocates surplus $\alpha\Delta$ from the incumbent to period $\tau + 1$ consumers. However, in subsequent periods, an innovation confers an expected benefit Δ to consumers but at an expected cost to the incumbent that is less than Δ as long as the probability of future displacement is positive (i.e., $\phi > 0$).

Now, consider the effects on potential entrants. Observe that in each of the innovation supply examples discussed in Section II, $u(\Phi(w))w - c(\Phi(w))$ is nondecreasing in w, i.e., each potential entrant is weakly better off if the innovation prize increases. We refer to this as the *value monotonicity property*. Since we move along the upward-sloping IS curve when α increases, the increase in ϕ must be associated with an increase in w. So when the value monotonicity property is satisfied, each potential entrant becomes (weakly) better off, which also implies that aggregate surplus increases.

[18] The lack of tension that arises here has a different cause (the indifference property) than that in Section I's discussion of R&D-deterring activities. Here, the incumbent's deterrence choice is not interior, in contrast to the argument in Section I.

[19] We make this assumption to simplify the analysis. By doing so, the equilibrium innovation rate transitions immediately to its new steady-state value, and all payoff changes begin in period $\tau + 1$. If, instead, the intervention occurs at the start of the period, there would be a one-period transitory effect on the innovation rate, because in period τ the share of captive customers facing an entrant would be the level before the policy change, while the continuation values V_I and V_E starting in period $\tau + 1$ would be the levels in the new steady state.

Finally, what about the current incumbent? This turns out to be ambiguous: on one hand, the increase in ϕ speeds the incumbent's displacement. On the other hand, the value V_E that the incumbent receives when he is displaced may increase.[20] In sum:

PROPOSITION 4: *A once-and-for-all restriction on the share of long-term (exclusive) contracts that strictly increases innovation raises consumer surplus and (when the value monotonicity property is satisfied) weakly raises the values of potential entrants and increases aggregate surplus. The effect on the current incumbent's value is ambiguous.*

It is perhaps surprising that the welfare effect of an increase in α is necessarily positive, given that the equilibrium innovation rate may be above the first-best level due to business stealing (see the Web Appendix). Note, however, that long-term contracts involve an inefficiency, since when entry occurs the incumbent makes sales of an old technology to captive consumers. Proposition 4 tells us that even when an increase in the share of captive consumers brings a socially excessive innovation rate closer to the first-best level, the waste effect dominates and aggregate welfare is reduced.

As this suggests, a restriction on long-term (exclusive) contracting may lower aggregate surplus if the incumbent is able to renegotiate the long-term contracts once entry occurs, and sell captive consumers a better product that he procures (or licenses) from the entrant. To consider the easiest case, suppose that the incumbent can extract a captive consumer's full willingness to pay for the better product, equal to Δ, and splits this gain with the entrant. If $\theta \in [0,1]$ is the incumbent's share of this gain, then whenever the antitrust constraint binds, we can apply our framework by modifying the profit flows to be

$$\pi_m(\alpha,\phi) = \alpha\Delta + (1-\alpha)(1-\phi)\Delta; \tag{7}$$

$$\pi_I(\alpha,\phi) = (1-\alpha)(1-\phi)\Delta + \theta(1-\alpha)\Delta;$$

$$\pi_E(\alpha,\phi) = \alpha\Delta + (1-\theta)(1-\alpha)\Delta.$$

This reduces the positive impact of a more protective policy on the entrant, holding ϕ fixed, to $\partial\pi_E(\alpha,\phi)/\partial\alpha = \theta\Delta > 0$. It also makes negative the effect of raising α on a continuing incumbent's expected profit holding ϕ fixed:

$$\left[\phi\frac{\partial\pi_I(\alpha,\phi)}{\partial\alpha} + (1-\phi)\ \frac{\partial\pi_m(\alpha,\phi)}{\partial\alpha}\right] = -\theta\phi\Delta.$$

Nonetheless, since the entrant's gain exceeds the continuing incumbent's expected loss, the front-loading effect implies that an increase in α still increases innovation. However, because the best product is consumed by all consumers in each period, the effect on aggregate surplus is now determined solely by whether the innovation rate is too high or too low, which can go in either direction in general.

B. *Compatibility in a Network Industry*

We next consider a model of compatibility choices by a leading firm in an industry with network externalities. The model is patterned after Fudenberg and Tirole (2000), who studied limit pricing in a dynamic model.[21] Overlapping generations of consumers live for two periods and make purchases only when young. The good lasts for two periods and each generation is of unit mass. Thus, in each period there is an old generation of consumers alive who made purchases in the previous period, and a young generation who will buy today. The value of consumption in period t is $v(N) + j\Delta$ if a consumer consumes version j of the good and this good has a "network size" of N. We follow the convention in the network externalities literature and assume that consumers in each generation coordinate their purchases, acting as a single agent and purchasing from a single firm.

There are many ways in which compatibility is determined in actual markets. Here we focus on one that leads to a relatively simple model that fits our framework. We assume that each firm that offers its product to consumers chooses a price p and also a compatibility level $b \in [0,1]$

[20] This effect is absent when there is free entry, since then $V_E = 0$. In that case, the incumbent's value must fall.

[21] There are several key differences: we have only one type of consumer (so limit pricing is not possible), firms make compatibility choices, and patent protection lasts forever.

of this product with higher-quality products. Network benefits are determined as follows: the network size enjoyed by generation g of consumers who have bought version j is 2 if all consumers in a period consume version j, is 1 if the other existing generation of consumers consumes a higher quality good, and is $1 + b_{j-1}$ if the other existing generation of consumers consumes version $j - 1$, which has compatibility level b_{j-1}. (We assume that products more than one step removed in quality are incompatible with each other, which simplifies the analysis of off-equilibrium situations.) That is, while consumers of the higher-quality product j can benefit from the existence of consumers who consume the lower-quality product $j - 1$ (to the extent that version $j - 1$'s producer allows), the reverse is not true.[22, 23] The cost of producing a product with compatibility level b is $k(b)$. We let $\underline{k} \equiv \min_{b \in [0,1]} k(b)$.

In this model, the leading firm has an incentive to reduce its compatibility with a future entrant's product to reduce entrants' R&D. Our policy parameter $\alpha \in [0,1]$ puts a lower bound on the compatibility level b it can choose.[24]

We assume that the quality improvement in each step is large enough relative to the costs and benefits of compatibility: $\Delta > v(2) - v(1) + [k(0) - \underline{k}]/(1 + \delta)$. This assumption implies that the order of products' values (including network effects) is always the same as their technological ordering, and will allow us to construct an equilibrium in which the technological leader wins the competition for consumers in every subgame.

Formally, the timing in each period t is:

- **Stage** t.1: Each potential entrant i observes the previous history—including the compatibility choice of the firm that sold to consumers in the previous period, B_{t-1}, and that firm's lag behind the technological frontier—and chooses its innovation rate ψ_{it}. Then innovation success is realized.
- **Stage** t.2: Each firm i chooses its compatibility level $b_{it} \in [0,1]$ and names price p_{it} to young consumers. The leading firm must have $b_{it} \geq \alpha$.
- **Stage** t.3: Young consumers make their purchase decisions.

We focus on Markov perfect equilibria, in which agents condition their choices only on the payoff-relevant state. In stage t.1, entrants' R&D choices are contingent on the compatibility level B_{t-1} chosen by the firm that sold in the previous period, as well as its lag behind the technological frontier. Each firm's offer at stage t.2 is contingent on the past winner's technological standing and compatibility level B_{t-1}, on the firm's technological standing, and on whether entry occurred at stage t.1. A young consumer's decision at t.3 depends on all this, as well as on the offers made at stage t.2. We focus on Markov perfect equilibria in which the leading firm (whether a continuing incumbent or a new entrant) always wins the sales to young consumers. In addition, we focus on equilibria with the property that losing firms do not make offers on which they would lose money if they were accepted (this could be viewed as ruling out weakly dominated strategies).

In the Appendix, we show that, in any such equilibrium, the leading firm at stage t.2 has the same compatibility incentives regardless of whether it is a new entrant or an uncontested incumbent, and in the Markovian spirit we restrict it to make the same compatibility choice, which we denote b^*. Then, on the equilibrium path, the value of being an incumbent, the value of being an entrant, and the rate of innovation are the same in every period. As a result, we can use the framework of Section II to determine the equilibrium rate of innovation once we have determined the equilibrium profit flows.

Consider, first, the period t profit of a continuing incumbent, who has the leading quality, in a period without entry. Its relevant competitor is the firm with the next-highest quality product, the previous incumbent. The previous incumbent

[22] As an example, consider 386 and 486 chips: software designed for 386 machines also worked on 486 machines, but not the reverse.

[23] In essence, we assume that the higher-quality product can costlessly achieve as much backward compatibility as the lower- quality firm allows. Were we to allow the higher-quality firm a choice of whether it wants backward compatibility (at no cost), it would always choose the maximal possible level. Regarding the lower-quality product, our assumptions allow its producer to commit to a compatibility level. This may be thought of as a product design choice. For example, a patented interface may prevent a new entrant from achieving backward compatibility.

[24] We remark below on the effects of having all firms subject to this constraint.

cannot offer any compatibility benefits to the young consumers (since tomorrow's young consumers are certain to buy a higher-quality product), and also does not derive any future benefits from today's sales (as we argue in the Appendix). It therefore chooses the least-cost compatibility level, and offers young consumers a (lifetime) surplus of

$$(1 + \delta)\,[v(1) + (j_t - 1)\Delta] - \underline{k}.$$

The continuing incumbent can, instead, offer young consumers a (lifetime) surplus of $[v(2) + \delta G(\phi) + (1+\delta) j_t \Delta] - k(b^*)$, where

$$G(\phi) \equiv (1 - \phi)v(2) + \phi v(1)$$

is the future expected network benefit that a young consumer anticipates from the current leading product when the probability of entry tomorrow is ϕ. The continuing incumbent therefore earns

$$\pi_m = (1 + \delta)\,\Delta + [v(2) + \delta G(\phi) - k(b^*)] - [(1+\delta)v(1) - \underline{k}]$$

from its sales to young consumers.

Now, consider the profit of an entrant competing against an incumbent who was the highest quality firm prior to entry. The continuing incumbent (who is the entrant's relevant competitor) can offer no compatibility benefits in period $t + 1$ to period t's young consumers, and derives no future benefits from today's sales. It therefore chooses the least-cost compatibility level and offers young consumers a surplus of

$$v(2) + \delta v(1) + (1 + \delta)(j_t - 1)\Delta - \underline{k}.$$

The entrant can, instead, offer young consumers a (lifetime) surplus of $[v(1 + b^*) + \delta G(\phi) + (1+\delta) j_t \Delta] - k(b^*)$. The entrant will therefore earn

$$\pi_E = (1 + \delta)\,\Delta + [v(1 + b^*) + \delta G(\phi) - k(b^*)] - [v(2) + \delta v(1) - \underline{k}]$$

on the current consumers. Finally, a contested incumbent makes no sales and earns $\pi_I = 0$.

Now, we consider the effect of antitrust policy α on innovation. If, for a given α, there are equilibria with different compatibility levels, we focus on the ones with the highest compatibility level, which we denote $b^*(\alpha)$. In the Appendix, we show that this level $b^*(\alpha)$ is nondecreasing in α (in particular, if the antitrust constraint binds with equality, $b^*(\alpha) = \alpha$). Substituting into the profit expressions above, the equilibrium fits into our framework with the following profit functions:

$$\pi_m(\alpha,\phi) = (1+\delta)\Delta + [v(2) + \delta G(\phi) - k(b^*(\alpha))] - [(1+\delta)v(1) - \underline{k}]; \tag{8}$$

$$\pi_I(\alpha,\phi) = 0;$$

$$\pi_E(\alpha,\phi) = (1+\delta)\Delta + [v(1 + b^*(\alpha)) + \delta G(\phi) - k(b^*(\alpha))] - [v(2) + \delta v(1) - \underline{k}].$$

This model fits into the framework of R&D-deterring activities discussed in Section I. So introducing a protective policy that just slightly constrains the incumbent's compatibility choice necessarily increases innovation. More generally, suppose that $k(\cdot)$ is decreasing, so that it is costly to block compatibility by a higher-quality entrant. In this case, π_m and π_E both increase when α increases holding ϕ fixed, while π_I remains unchanged (see (8)). Since no tension arises in this case, Propositions 1 and 2 tell us that in this case a more protective policy always increases innovation.

PROPOSITION 5: *In our model of compatibility choice, when production costs are lower for more compatible products, raising the leading firm's minimum compatibility requirement increases innovation in a highest-compatibility equilibrium.*

When $k(\cdot)$ is not everywhere decreasing, however, increased protectiveness may instead lower innovation. For example, suppose that

$k(\cdot)$ is convex with its minimum at $\underline{b} \in (0,1)$. Examining (8), we see that increasing α above $\underline{b}$ may reduce innovation: when the antitrust constraint binds, it certainly lowers π_m and may even lower π_E [if $v(1+\alpha) - k(\alpha)$ falls].[25] In this case, forcing compatibility above the level that minimizes costs may reduce R&D because we may no longer be in a region where a continuing incumbent's profit is increased, holding ϕ fixed, by an increase in compatibility.

A full investigation of the welfare effects of a more protective policy in this model is beyond our scope here. Nevertheless, it is worth noting that in this model innovation can be excessive, even with only a single potential entrant (in contrast to our benchmark model in the Web Appendix). The reason is that an externality exists between the two generations of consumers in each period: when the young purchase an entrant's product, they leave old consumers with lower network benefits. Indeed, while young consumers have a benefit of $[\Delta - v(2) + v(1+\alpha)]$ in the first period that an entrant is in the market, the old consumers lose $[v(2) - v(1)]$. Thus, when $\Delta/(1-\delta) < [2v(2) - v(1) - v(1+\alpha)]$, an innovation lowers aggregate welfare, even ignoring R&D costs. Thus, it would not be surprising if a more protective policy that raises innovation could lower aggregate welfare here.

IV. Incumbent Innovation

The analysis above imposed the strong restriction that only potential entrants engage in R&D. Although useful for gaining insight, this assumption is not representative of most settings of interest. In this section, we explore how our conclusions are affected when incumbent firms may also engage in R&D.

Allowing incumbent firms to engage in R&D has the potential to complicate the analysis considerably. In particular, once we allow for incumbent investment, we need to introduce a state space to keep track of the incumbent's current lead over the potential entrants. In general, the rates of R&D investment by the incumbent and its challengers may be state dependent along the equilibrium path (see, for example, Aghion et al. 2001).

Here, we focus on two special cases in which R&D strategies are nonetheless stationary. Although clearly restrictive, these two models have the virtue of capturing two distinct motives for incumbent R&D: (a) preventing displacement by an entrant, and (b) increasing the flow of profits until displacement by increasing the lead over the previous incumbent.

A. *R&D to Prevent Displacement*

Suppose the incumbent can conduct R&D denoted by ϕ_I, while the potential entrants' aggregate R&D is denoted by ϕ_E. Similarly, the incumbent's and entrants' respective cost functions are $c_I(\phi_I)$ and $c_E(\phi_E)$ (we allow the cost functions to differ between the incumbent and the potential entrants). In this first model, we assume that if the leading quality level in period t is j_t, then quality level j_{t-1} is freely available to all potential producers. That is, it enters the public domain. Thus, the incumbent never has a lead greater than one step on the ladder. If so, the only reason for an incumbent to conduct R&D is to try to get the patent on the next innovation in cases in which at least one potential entrant has made a discovery—that is, to prevent its displacement.[26] With these assumptions, we need not keep track of any states, and there is a stationary equilibrium.

[25] If the policy, instead, constrained all firms and not just the leading firm, then we would replace $\underline{k}$ with $k(\alpha)$ in the profit expressions. In this case, π_m and π_E would be weakly increasing in α, so innovation would increase with α.

[26] In the usual sort of (Poisson) continuous-time model considered in the R&D literature (see, e.g., Lee and Wilde 1980; Reinganum 1989; Grossman and Helpman 1991), the probability of ties is zero, and so one might worry that our formulation here is dependent on a merely technical feature of the discrete-time setup. Indeed, in a continuous-time model, the incumbent would conduct no R&D. However, the usual continuous-time model relies on the implicit assumption that following an innovation, all firms reorient their R&D activity instantaneously to the next technology level. If, instead, we were to use a continuous-time model in which there is a fixed time period after a rival's success before which R&D for the next technology level cannot be successful, then we would get effects that parallel those in our discrete-time model (where the discount factor δ reflects how quickly R&D activity can be reoriented to the next technology level). Thus, our discrete-time formulation captures an arguably realistic feature of the economics of R&D.

To proceed, let $r_I(\phi_E)$ denote the probability that an incumbent who innovates preserves its incumbency. Since this will always happen if no entrant makes a discovery, $r_I(\phi_E) \geq (1 - \phi_E)$. The difference $[r_I(\phi_E) - (1 - \phi_E)]$ equals the probability that entrants have made a discovery and the innovating incumbent retains its incumbency. Letting $d(\phi_I, \phi_E) \equiv (1 - \phi_I)\phi_E + \phi_I(1 - r_I(\phi_E))$ denote the probability that the incumbent is displaced by an entrant and $u(\phi_I, \phi_E)$ denote the probability that a given entrant innovates and becomes the next incumbent, we can write the values V_I and V_E of the incumbent and a potential entrant as

(VI**)

$$V_I = \pi_m(\alpha, \phi) + \delta V_I + d(\phi) \times \{\pi_I(\alpha, \phi) - \pi_m(\alpha, \phi) - \delta(V_I - V_E)\} - c_I(\phi_I);$$

(VE**)

$$V_E = \delta V_E + u(\phi) \times [\pi_E(\alpha, \phi) - \delta(V_I - V_E)] - c_E(\phi_E),$$

where $\phi = (\phi_I, \phi_E)$. The incumbent will choose its R&D to solve

$$\phi_I \in \arg\max_{\psi_I \in [0,1]} \psi_I[r_I(\phi_E) - (1 - \phi_E)] \times w_I - c_I(\psi_I), \tag{9}$$

where

$$w_I \equiv \pi_m(\alpha, \phi) - \pi_I(\alpha, \phi) + \delta(V_I - V_E) \tag{10}$$

is the incumbent's expected gain from getting a patent conditional on one of the entrants making a discovery. As for the potential entrants, they are facing innovation prize

$$w_E = \pi_E(\alpha, \phi) + \delta(V_I - V_E), \tag{11}$$

and will choose $\phi_E = \Phi(r_E(\phi_I)w_E)$, where $r_E(\phi_I)$ is the probability that the incumbent is replaced conditional on an innovation by an entrant, and $\Phi(\cdot)$ is the entrants' innovation supply function.

The incumbent and entrant innovation probabilities ϕ_I and ϕ_E may differ for several reasons. First, their cost functions may differ. Second, there are many entrants but only one incumbent. Third, their innovation prizes differ: when $\pi_m \geq \pi_I + \pi_E$ (the efficiency effect), w_I is at least as large as the entrants' innovation prize w_E. Fourth, an incumbent finds innovating valuable only when it prevents an entrant from receiving the patent on the next generation product, while an entrant finds innovation worthwhile even if no other firm has made a discovery (this is commonly called the "replacement effect").

In some cases, the incumbent will, in fact, find it optimal to do no R&D. To illustrate, suppose the first three effects are absent: that is, suppose there is one entrant, its R&D cost function is identical to the incumbent's, and the efficiency effect is absent $(\pi_m = \pi_I + \pi_E)$. Assume, in addition, that a patent is randomly awarded among innovating firms, so that the entrant and incumbent have the same ability to secure a patent. In that case, $r_i(\phi_{-i}) = (1 - \phi_{-i}/2)$ for $i = I$, E, and the entrant's R&D level is

$$\phi_E \in \arg\max_{\psi_E \in [0,1]} \psi_E \, r_E(\phi_I) \, w_E - c_E(\psi_E). \tag{12}$$

Comparing (12) and (9), the entrant has a larger incentive to do R&D than the incumbent. If, in addition, the R&D cost functions have constant returns to scale or if there is free entry, so that the entrant in equilibrium just breaks even on its R&D efforts, the incumbent does no R&D in equilibrium. In such cases, our previous results apply directly.

To examine the effects of antitrust policy in the more general case in which both the incumbent and entrants do R&D, we solve (VI**) and (VE**) for $(V_I - V_E)$ and substitute it in (10) and (11). This yields (suppressing the arguments of functions)

$$w_I = \frac{1}{1 - \delta + \delta(d + u)} \times \{\pi_m - (1 - \delta)\pi_I + \delta u(\pi_m - \pi_E - \pi_I) + \delta(c_E - c_I)\}, \tag{13}$$

$$(14) \quad w_E = \frac{1}{1 - \delta + \delta(d + u)} \times \{\delta\pi_m + (1 - \delta)\pi_E - \delta d(\pi_m - \pi_E - \pi_I) + \delta(c_E - c_I)\}.$$

A change in policy can be thought of as having both direct and indirect effects on entrant and incumbent R&D. Direct effects capture the changes that would occur in each innovation rate if we held fixed the other innovation rate. Direct effects do not, on their own, determine the effects of a policy change on innovation rates, however, because a change in one innovation rate may have an effect on the other. Given limited space and the difficulty of characterizing in general these indirect effects, we focus on direct effects only in the rest of the section.

Using reasoning like that in Sections I and II, the direct effects on the incumbent's and entrants' innovation rates have the same sign as the effects on w_I and w_E, respectively, holding (ϕ_I, ϕ_E) fixed. Differentiation of (14) implies that the sign of the direct effect on ϕ_E is given by condition (4), but with the probability of displacement now being $d(\phi)$ rather than ϕ. So, the direct effects on entrant R&D of a more protective policy are similar to those identified before. The direct effect on incumbent R&D, on the other hand, can be either positive or negative. The following two examples illustrate these observations using extensions of the models in Section 3.

Example 1: Consider an extension of our long-term (exclusive) contracting model of Section IIIA that allows for incumbent innovation to prevent displacement. We now assume that a long-term contract is a commitment by the incumbent to deliver his best current product in the next period. Using similar reasoning as in Section IIIA, the profit functions in this case are the same as in (6), but with $d(\phi)$ replacing ϕ. Thus, the efficiency effect is zero: $\pi_m = \pi_I + \pi_E$. Moreover, an increase in α raises π_m and π_E, and lowers π_I. From (13) and (14), these facts imply that the direct effects of a more protective policy raise both ϕ_E and ϕ_I. Intuitively, by raising π_m and lowering π_I, a more protective policy increases an incumbent's incentive to avoid displacement.

Example 2: Now consider an extension of the compatibility model of Section IIIB that allows for incumbent innovation to prevent displacement. We assume that a firm's products that are one generation apart are fully compatible with one another. When an incumbent innovates and its previous leading product enters the public domain, however, the versions sold by other firms are not compatible with the incumbent's new or previous leading products. With these assumptions, a consumer's expected compatibility benefits from the leading and nonleading firms' products are the same as in Section IIIB. Therefore, the profit functions are the same as in (8), but with $G(d(\phi))$ replacing $G(\phi)$.

In this example, the efficiency effect $(\pi_m - \pi_I - \pi_E)$ is positive but decreasing in α, while π_m and π_E are increasing in α (π_I is zero). From (13) and (14), these facts imply that the direct effect on ϕ_E is positive, while the direct effect on ϕ_I is ambiguous. If u is low (it is zero with free entry), then the direct effect on ϕ_I is positive.

B. *R&D to Increase Profit Flows*

We next consider a model in which rivals do not get access to the second-best technology when the incumbent innovates. Thus, the incumbent can increase its flow of profits by innovating, until the time when it is displaced. Specifically, let s denote the number of steps that the incumbent is ahead of its nearest rival (this is our state variable). The variable s affects the incumbent's profit flow when entry does not occur, which we now denote by $\pi_m^s(\alpha, \phi)$ (it does not affect either π_I or π_E). We make two assumptions that imply that there is an equilibrium in which the R&D levels of the incumbent and potential entrants do not depend upon s. Specifically, we assume that $\pi_m^s(\alpha, \phi) = \bar{\pi}_m(\alpha, \phi) + s\tilde{\pi}_m(\alpha, \phi)$ (models with this feature are presented in the two examples below) and that an entrant gets the patent whenever at least one entrant has made a discovery.

Given these assumptions, we can construct a Markov perfect equilibrium of the model in which both innovation rates, ϕ_E and ϕ_I, as well as the value V_E, are stationary, while the value of an incumbent that is s steps ahead takes the form $V_I^s = \bar{V}_I + s\tilde{V}_I$, where $\tilde{V}_I$ is the incremental value to an incumbent of moving one step further ahead. Because the probability of displacement is unchanging and an incumbent's profits increase linearly in the size of its lead, this incremental value is independent of the incumbent's current lead and equals

$$\tilde{V} = \frac{\tilde{\pi}_m(\alpha,\phi)(1-\phi_E)}{1-\delta+\delta\phi_E}.$$

The incumbent's equilibrium innovation rate then satisfies

$$(15) \qquad \phi_I \in \arg\max_{\psi\in[0,1]} \psi\tilde{\pi}_m(\alpha,\phi) \times\left[\frac{1-\phi_E}{1-\delta+\delta\phi_E}\right] - c_I(\psi).$$

Condition (15) implies that the direct effect of antitrust policy on incumbent innovation in this equilibrium is determined solely by $\partial\tilde{\pi}_m(\alpha,\phi)/\partial\alpha$.

As for the entrants' innovation benefit function, it can be obtained by solving for $(V_I^1 - V_E)$ from the value equations and then writing (suppressing arguments of functions)

$$(16)$$

$$w_E = \pi_E + \delta(V_I^1 - V_E)$$

$$= \frac{\left\{\begin{aligned}&[1-\delta\ (1-\phi_E)]\pi_E + \delta(1-\phi_E)\overline{\pi}_m \\ &+\delta\phi_E\pi_I \\ &+\delta(1-\phi_E)\left[1+\frac{\phi_I}{1-\delta+\delta\phi_E}\right]\tilde{\pi}_m \\ &-\delta(c_I-c_E)\end{aligned}\right\}}{1-\delta+\delta(\phi_E+u(\phi_E))}.$$

This parallels the case without incumbent investment, except for the term multiplying $\tilde{\pi}_m$, which now accounts for the possibility of rising incumbent profits over time (the difference vanishes when $\phi_I = 0$). The direct effect of α on ϕ_E can be seen by differentiating (16) with respect to α.

Example 3: Consider an extension of the long-term (exclusive) contracts model to incumbent innovation to increase profit flows. To fit into our framework here, however, we change the timing of the payment in this contract, assuming, instead, that the payment is made when the contract is signed. In this case,[27, 28]

$$(17)$$

$$\begin{aligned}\pi_m^s(\alpha,\phi) &= \alpha s\Delta + (1-\alpha)[\delta(s+\phi_I)(1-\phi_E)\Delta \\ &\qquad -(1-\delta)k] \\ &= (1-\alpha)[\delta\phi_I(1-\phi_E)\Delta - (1-\delta)k] \\ &\qquad + s\Delta[\alpha+\delta(1-\alpha)(1-\phi_E)];\end{aligned}$$

$$\pi_I(\alpha,\phi) = -(1-\alpha)k;$$

$$\pi_E(\alpha,\phi) = \alpha\Delta + (1-\alpha)\delta[k+(1+\phi_I)(1-\phi_E)\Delta].$$

From (17), $\tilde{\pi}_m(\alpha,\phi) = \Delta[\alpha+\delta(1-\alpha)(1-\phi_E)]$, so the direct effect on ϕ_I of an increase in α is always positive. On the other hand, the direct effect on ϕ_E may now be either positive or negative. When ϕ_I is close to zero, the direct effect is the same as absent incumbent innovation, so it is necessarily positive. When ϕ_I is large, however, this conclusion can be reversed.

Example 4: Consider an extension of the compatibility model to incumbent innovation

[27] Observe that the price of a long-term contract, which is paid when signed in period t, is $q_t = \delta[k + (s + \phi_I)(1-\phi_E)\Delta]$. A continuing uncontested incumbent in period t sells to free consumers (for a profit of $\alpha s\Delta$), delivers on contracts written in period $t-1$, and writes new contracts for period $t+1$ delivery. An incumbent who faces new entry in period t delivers only on period $t-1$ contracts. An entrant in period t sells to free consumers (at a profit of $\alpha\Delta$), and writes new contracts for period $t+1$ deliveries.

[28] Note that with this change in the timing of payments, an increase in α that leaves more free consumers can lower π_E as defined here.

to increase profit flows. The profit of an uncontested incumbent who is s steps ahead is

$$\pi^s_m(\alpha,\phi) = (1+\delta)\ s\Delta + [v(2) + \delta G(\phi_E) - k(b^*(\alpha))] - [(1+\delta)v(1) - \underline{k}].$$

In this example, $\tilde{\pi}_m(\alpha,\phi) = (1+\delta)\ \Delta$, which is independent of α. Therefore, a more protective policy has no direct effect on incumbent innovation. From (16), the effect of a more protective policy on entrant innovation is exactly the same as in Section IIIB.

V. Other Types of Antitrust Policies

Up to this point, we have focused on policies that alter the profits that incumbents and entrants earn in competition with one another. In this section, we briefly consider two examples of policies that have different effects.

A. *Predatory Activities*

In some situations, antitrust may affect not only an entrant's profits in competition with the incumbent, but also the entrant's probability of survival. To focus solely on this effect, take π_I, π_E, and π_m, as fixed and suppose that a new entrant's probability of survival following its entry is $\lambda(\alpha)$, where $\lambda(\cdot)$ is increasing in α.

Now the innovation prize is

$$w = [\pi^E + \delta\lambda(\alpha)(V_I - V_E)]. \tag{18}$$

If $(V_I - V_E)$ were fixed, an increase in α would necessarily increase innovation. Now,

$$V_I = \pi_m + \delta V_I + \phi[\pi_I - \pi_m + \delta\lambda(\alpha)(V_E - V_I)], \tag{VI***}$$

$$V_E = \delta V_E + u(\phi)[\pi_E + \delta\lambda(\alpha)(V_I - V_E)] - c(\phi). \tag{VE***}$$

Subtracting (VE***) from (VI***), we can express the innovation prize with the function

$$W(\phi,\alpha) = \left\{\pi_E + \left(\frac{\delta\lambda(\alpha)}{1-\delta+\delta\lambda(\alpha)(\phi+u(\phi))}\right) \times [\phi\pi_I + (1-\phi)\pi_m - u(\phi)\pi_E + c(\phi)]\right\}. \tag{19}$$

The fraction $\delta\lambda(\alpha)/[1-\delta+\delta\lambda(\alpha)(\phi+u(\phi))]$ is increasing in α. Hence, provided that $(V_I - V_E)$ is positive, a more protective antitrust policy that raises the likelihood of entrant survival necessarily increases the innovation benefit.[29] Propositions 1 and 2 then tell us that innovation increases with this change in α.

In the working paper version (Segal and Whinston 2005), we illustrated this effect in a simple model of predatory pricing in which the entrant's probability of survival was an increasing continuous function $\lambda(\pi_E)$ of its first-period profit. This created an incentive for the incumbent to undercut a new entrant, and reduced equilibrium prices in the period of entry. The antitrust constraint α was a lower bound on the incumbent's price. An increase in this lower bound raised π_E, did not alter π_m and π_I (the latter because the incumbent makes no sales when entry occurs), and raised the entrant's probability of survival $\lambda(\pi_E)$. Both the profit flow and survival effects therefore increased the rate of innovation, which had ambiguous effects on aggregate welfare.

B. *Shifting Innovation Supply*

In some cases, incumbents may take actions that instead affect innovation supply. For example, an incumbent may buy up needed R&D inputs, thereby raising potential entrants' R&D costs. As another example, incumbents may engage in meritless patent litigation claiming that an entrant's innovation infringes its own

[29] For example, this will always be true whenever $V_E = 0$ (say, because of free entry) and π_m and π_I are nonnegative. Another sufficient condition is $\phi\pi_I + (1-\phi)\pi_m \geq u(\phi)\pi_E$ for all ϕ.

patent, raising the cost or lowering the probability of the entrant receiving a patent. Formally, we now denote the innovation supply correspondence by $\Phi(\cdot,\alpha)$. As the following propositions establish, increases in innovation supply lead to increases in innovation in the same senses as before (we omit the proofs, which are similar to those earlier).

PROPOSITION 6: *If the innovation supply function $\Phi(\cdot,\alpha)$ is continuous and nondecreasing, and the innovation benefit function $W(\phi,\alpha)$ is continuous in ϕ, then the largest and smallest equilibrium innovation rates exist, and both these rates are nondecreasing (nonincreasing) in α—the protectiveness of antitrust policy—if $\Phi(\cdot,\alpha)$ is nondecreasing (nonincreasing) in α for any $\phi \in [0,1]$.*

PROPOSITION 7: *Suppose that the innovation supply correspondence $\Phi(\cdot,\alpha)$ is continuous and nondecreasing, and the innovation benefit function $W(\phi,\alpha)$ is continuous in ϕ. Suppose, in addition, that for all $\alpha \in [\underline{\alpha},\overline{\alpha}]$ there is a unique equilibrium innovation rate $\phi(\alpha)$ on an interval $[\underline{\phi},\overline{\phi}]$ and that the IB curve crosses the IS curve from above on this interval. Then, $\phi(\alpha)$ is nondecreasing (nonincreasing) if $\Phi(\cdot,\alpha)$ is nondecreasing (nonincreasing) in α for all $\alpha \in [\underline{\alpha},\overline{\alpha}]$ and $\phi \in [\underline{\phi},\overline{\phi}]$.*

Thus, rightward (leftward) shifts of the innovation supply correspondence cause the rate of innovation to increase (decrease) in every stable equilibrium. Returning to the two examples mentioned above, these incumbent behaviors shift both innovation benefit and supply. If an uncontested incumbent overbuys needed R&D inputs at the end of each period, this raises potential entrants' R&D costs (shifting the IS curve leftward) and lowers π_m. Meritless patent litigation, on the other hand, shifts the IS curve leftward and lowers both π_I and π_E. In both cases, restricting the behavior increases innovation.

VI. Conclusion

In this article, we have studied the effects of antitrust policies in industries in which innovation is central to competitive outcomes. In general, a tension arises in discerning the effects of antitrust policy on innovation. On the one hand, limiting incumbent behaviors that reduce the initial profit of entrants increases the incentives for R&D. But since these same limitations affect a successful entrant once it in turn becomes an incumbent, they could actually reduce innovation incentives. Our results show how to disentangle these two effects, and we illustrate their implications for a number of antitrust policies. Interestingly, once one looks at the effects on entrant and incumbent profits holding the rate of innovation fixed—as our comparative statics results instruct us to do—limitations on R&D-deterring activities may involve no tension at all, as both entrant and incumbent profits increase, holding the rate of innovation fixed.

Throughout we have focused on stationary equilibria in stationary models. While this stationarity allows us to focus in a clean way on the fundamental tension in these policies, and also facilitates our formal analysis, it may be important for some of our results. For example, observe that we have focused throughout on the level of innovation once an existing incumbent is established. Our results do not apply to the *first* firm in a market. That firm faces no incumbent and therefore derives no initial benefit from a protective policy. Protective policies therefore necessarily slow the creation of entirely new innovation ladders.

Finally, the tension between effects on entrant and continuing incumbent profits that is our focus also arises in other policy settings, most notably intellectual property (IP) protection. While the traditional tension considered in patent policy is between the (presumed) welfare benefit of creating greater ex ante R&D incentives and the welfare cost of greater ex post monopoly pricing distortions, the tension we focus on suggests instead that there are conflicting effects of IP protection on the incentives to innovate: it reduces profits for new innovators (who may produce infringing innovations) but increases profits for continuing incumbents (by delaying their replacement).

While our stationary model cannot study the effects of patent duration, the most commonly studied IP policy (because patents lead to nonstationary market states), versions of our model can be applied to other types of IP policies. For example, we have implicitly assumed throughout that a new innovator not only can get a patent, but also is free to build on the previous leading technology without paying anything to its owner. In a sense, the previous leading technology must be licensed for free. Suppose, instead, that the new innovator must pay a fee f to license the previous technology. In our model, this fee raises π_I by f and lowers π_E by f. By our front loading result, an increase in f must decrease the rate of innovation.

A tension similar to that in our model is at the heart of recent work by O'Donoghue, Scotchmer, and Thisse (1998), Hunt (2004), and Hugo Hopenhayn, Gerard Llobet, and Matthew F. Mitchell (2006) on the effects of a novelty requirement (leading breadth), which requires that patentable innovations be at least a certain amount better than the existing state of the art. Such rules directly prevent small innovations from coming to market. They also indirectly affect the rate of innovation by changing innovation incentives. This incentive effect involves a tension similar to that in our model: a more stringent novelty requirement reduces the profits from innovation in some ways (by reducing the likelihood of finding a patentable innovation) and increases it in others (by slowing an incumbent's replacement). A general result in such settings, for example, is that innovation always increases with a slight strengthening of the novelty requirement when the novelty requirement is initially low. This is because small innovations bring little profit to a new innovator, but can destroy a great deal of continuing incumbent profit.

In the working paper version of this article (Segal and Whinston 2005), we studied an extension of our long-term (exclusive) contracting model in which the size of an innovation's improvement over the existing technology is random. By reducing the profitability of implementing an innovation, long-term contracts prevent small innovations from entering the market, just as does a novelty requirement. This effect can lead innovation to decrease when antitrust policy restricts long-term contracting, in contrast to Proposition 3. We then introduced a novelty requirement and asked how optimal antitrust and intellectual property policies interact. In the model, a novelty rule is the better way to keep small innovations out of the market because, unlike long-term contracts, it does not create inefficient production. Indeed, once a novelty requirement is available, aggregate surplus is greatest when long-term contracts are prohibited, just as in Proposition 4. However, novelty rules require courts to determine the degree of improvement brought by new innovations, while enforcement of antitrust policy does not.[30]

Exploring further the effect of the tension we identify for IP policy, and also the interactions among the various policies that can affect innovation, seems a fruitful area for further research.

[30] A leading breadth policy could in principle also be implemented indirectly, by requiring an innovator to pay a fee to gain access to the market, as in Hopenhayn, Llobet, and Mitchell (2006).

Appendix: Proofs

PROOF OF PROPOSITION 1:

The equilibrium innovation rates are the fixed points of the composite function $f(\cdot,\alpha)\equiv\Phi(W(\cdot,\alpha))$, which is continuous under our assumptions. Furthermore, if $W(\phi,\alpha)$ is nondecreasing (nonincreasing) in α, then so is $f(\phi,\alpha)$. Theorem 1 of Milgrom and Roberts (1994) then establishes that the largest and smallest fixed points of $f(\cdot,\alpha)$ exist and are nondecreasing (nonincreasing) in α.

PROOF OF PROPOSITION 2:

Let

$$\chi(\phi) = \begin{cases} \underline{\phi} \text{ for } \phi \in [0,\underline{\phi}) \\ \phi \text{ for } \phi \in [\underline{\phi},\overline{\phi}]. \\ \overline{\phi} \text{ for } \phi \in (\overline{\phi},1] \end{cases}$$

Crossing from above implies that $\phi(\alpha)$ is the unique fixed point of the continuous composite function $f(\alpha,\cdot)\equiv\chi(\Phi(W(\alpha,\cdot)))$ on the interval $[\underline{\phi},\overline{\phi}]$. The proof of Proposition 1 then yields the result.

LEMMA A1: *The model with $N > 1$ entrants and innovation benefit w has a unique symmetric equilibrium innovation choice $\Psi(w)$, which is continuous and nondecreasing in w.*

PROOF:

Let $B(a) = \arg\max_{\psi' \in [0,1]} [a\psi' - \gamma(\psi')]$. Note that $B(a)$ is single-valued (by strict convexity of $\gamma(\cdot)$), nondecreasing by the Monotone Selection Theorem of Milgrom and Chris Shannon (1994), and continuous (by Berge's Theorem of the Maximum). The symmetric equilibrium R&D rates are then described as fixed points of the function $\sigma(w,\cdot) \equiv B(wr_N(\cdot))$ on the $[0,1]$ interval. Since $r_N(\cdot)$ is strictly decreasing and continuous, $\sigma(w,\cdot)$ is nonincreasing and continuous, and so $\sigma(w,\cdot)$ has a unique fixed point, which we denote by $\Psi(w)$. Since $B(\cdot)$ is nondecreasing, $\sigma(w,\psi)$ is nondecreasing in w, and therefore by Theorem 1 of Milgrom and Roberts (1994), the fixed point $\Psi(w)$ is nondecreasing in w. Furthermore, since $\sigma(w,\psi)$ is a continuous function, it can be seen (using Berge's Theorem of the Maximum) that $\Psi(\cdot)$ is also a continuous function.

PROOF OF PROPOSITION 3:

The payoff-relevant state at the start of period t is the pair (B_t, q_t) describing the share of consumers bound to long-term contracts and the price specified in their contracts. Markovian strategies specify: (a) an entrants' R&D choice as a function of the state, $\{\psi_i(B_t)\}$, which leads to an aggregate innovation rate $\phi^*(B_t)$ (note that entrants do not care about the price q_t specified in the incumbent's long-term contracts, only the share of consumers bound to them); (b) the leading firm's long-term contract choices, (b^*, q^*), which are independent of the state at the beginning of the period (since contracts are of one-period duration); (c) the consumers' long-term contract acceptance rule, which specifies for each possible share of consumers who are offered a long-term contract, b_{t+1}, the maximum price at which all consumers who received an offer accept, $\bar{q}^*(b_{t+1})$;[31] and (d) price offers to, and acceptance decisions by, free consumers (as noted in the text, these are just of the standard Bertrand variety).

The value functions giving beginning-of-the-period values as a function of the state (B_t, q_t) can then be written as

$$V_I(B_t, q_t) = \bar{V}_I(B_t) + [q_t - \bar{q}^*(B_t)]\ B_t,$$

where

$$\text{(A1)}\quad \bar{V}_I(B_t) = [\bar{q}^*(B_t) - k]\ B_t + [1 - \phi^*(B_t)]\Delta(1 - B_t) + \delta\{[1 - \phi^*(B_t)]\ \bar{V}_I(B^*) + \phi^*(B_t)\ \bar{V}_E(B^*)\},$$

and

$$V_E(B_t, q_t) = \bar{V}_E(B_t),$$

where

$$\text{(A2)}\quad \bar{V}_E(B_t) = \delta\bar{V}_E(B^*) + u(\phi^*(B_t))\{(1 - B_t)\Delta + \delta[\bar{V}_I(B^*) - \bar{V}_E(B^*)]\} - c(\phi^*(B_t)).$$

Offers to, and acceptance decisions by, free consumers are of the standard Bertrand variety, as noted above. Consider, next, the long-term contract acceptance rule $\bar{q}^*(b_{t+1})$. In the equilibrium, if all consumers among the b_{t+1} share of consumers offered a long-term contract accept, the next period's innovation rate will be $\phi^*(b_{t+1})$. Since each consumer views this rate as unaffected by his

[31] Given our assumption that consumers are small (take the probability of entry as independent of their own acceptance decision), whenever the leading firm offers a price $q \leq \bar{q}^*(b_{t+1})$ to share b_{t+1} of the consumers, it is a best response for a consumer to accept, given that all other consumers do. Thus, the continuation equilibrium in that case is that all of the b_{t+1} consumers accept (given our selection of the all-accept continuation equilibrium whenever it exists). If, instead, $q > \bar{q}^*(b_{t+1})$, then no consumers accept.

own decision, he is indifferent about accepting if $v + j_t\Delta - q_{t+1} \geq v + [j_t - 1 + \phi^*(b_{t+1})]\Delta - k$. Thus, $\bar{q}^*(b_{t+1}) = k + [1 - \phi^*(b_{t+1})]\Delta$. So, if the leading firm wishes to sign a share B_{t+1} of consumers, it offers them the price $\bar{q}^*(B_{t+1})$.

The optimality of R&D choices by entrants implies that the aggregate innovation rate in period t in state (B_t, q_t) satisfies

$$\phi^*(B_t) = \Phi((1 - B_t)\Delta + \delta[\bar{V}_I(B^*) - \bar{V}_E(B^*)]). \tag{A3}$$

Observe that $\phi^*(\cdot)$ is nonincreasing in B_t, and it is strictly decreasing if $\Phi(\cdot)$ is strictly increasing. Now, consider the leading firm's choice of the share of customers to sign to long-term contracts, B^*. This choice must satisfy

$$\begin{aligned} B^* &\in \arg\max_{B \leq 1-\alpha} \bar{V}_I(B) \\ &= \arg\max_{B \leq 1-\alpha} [1 - \phi^*(B)]\Delta B + [1 - \phi^*(B)]\Delta(1 - B) + \delta\{[1 - \phi^*(B)]\bar{V}_I(B^*) + \phi^*(B)\bar{V}_E(B^*)\} \\ &= \arg\max_{B \leq 1-\alpha} [\Delta + \delta\bar{V}_I(B^*)] - \phi^*(B)\{\Delta + \delta[\bar{V}_I(B^*) - \bar{V}_E(B^*)]\}. \end{aligned}$$

Since $\phi^*(\cdot)$ is nonincreasing, $\bar{V}_I(\cdot)$ is nondecreasing in B, provided that $\Delta + \delta[\bar{V}_I(B^*) - \bar{V}_E(B^*)] \geq 0$. Note that in an equilibrium we must have

$$\Delta + \delta[\bar{V}_I(B^*) - \bar{V}_E(B^*)] \geq \alpha\Delta + \delta[\bar{V}_I(B^*) - \bar{V}_E(B^*)] = w > 0,$$

where w is the equilibrium innovation prize. (Supposing in negation that $w \leq 0$, the equilibrium innovation rate would have to be $\Phi(w) \leq \Phi(0) = 0$, but then (A1) and (A2) would yield equilibrium values $\bar{V}_E(B^*) = 0$ and $\bar{V}_I(B^*) > 0$, and the equilibrium innovation prize w would be strictly positive, a contradiction.) Thus, $\bar{V}_I(B)$ is nondecreasing in B, so $B^* = 1 - \alpha$ is an optimal choice. Thus, we restrict attention to equilibria in which $B^* = 1 - \alpha$. Furthermore, if $\Phi(\cdot)$ is strictly increasing, then so is $\bar{V}_I(\cdot)$, and then *any* Markov perfect equilibrium must have $B^* = 1 - \alpha$. Finally, substituting $B_t = 1 - \alpha$, $\phi^* = \phi^*(1 - \alpha)$, $(1 - \phi^*)\Delta = \bar{q}^*(B_t)$, $V_I = \bar{V}_I(1 - \alpha)$, and $V_E = \bar{V}_E(1 - \alpha)$ into (A1), (A2), and (A3), we can fit this model into the framework of the stylized model of Section II, with , π_E, π_I, and π_m, given by (6). We can then can apply Proposition 1 as in the text to show that an increase in α raises the rate of innovation.

PROOF OF PROPOSITION 4:

We consider, in turn, the change in the payoffs of entrants, the current incumbent, consumers, and the current incumbent plus consumers.

Potential Entrants: If ϕ increases, then w must have increased. Using (VE*), we see that

$$(1 - \delta)\, V_E = u(\phi)w - c(\phi),$$

which implies that a potential entrant's value V_E has weakly increased if the value monotonicity property holds.

Sum of Current Incumbent and Consumers: We first compute a lower bound on the value change of the current incumbent (the firm with the leading technology just after stage τ.1). A policy change just after stage τ.1 changes the current incumbent's profits only beginning in the next period. From equation (VI*) we see that we can write

$$\begin{aligned} (1 - \delta + \delta\phi)V_I &= [(1 - \phi)\pi_m + \phi\pi_I] + \delta\phi V_E \\ &= (1 - \phi)\Delta + \delta\phi V_E. \end{aligned} \tag{A4}$$

Again, since V_E has weakly increased, a lower bound on the change in the current incumbent's value V_I starting at time $\tau + 1$ is the change in

$$\text{(A5)} \qquad \frac{(1-\phi)\Delta}{(1-\delta+\delta\phi)}.$$

Now, consider the consumers. Consumer welfare does not change until period $\tau + 1$ either. Since every consumer is always indifferent between signing an exclusive and being free, we can derive consumer welfare from period $\tau + 1$ on by assuming that all consumers are free. Thus, consumer welfare starting in period $\tau + 1$ is

$$\text{(A6)} \quad \left[(v_{j\tau}-k-\Delta)+\phi\frac{\Delta}{1-\delta}\right]+\delta\left[(v_{j\tau}-k-\Delta)+\phi\frac{\Delta}{1-\delta}\right]+\cdots=\left[\frac{(v_{j\tau}-k-\Delta)}{1-\delta}+\phi\frac{\Delta}{(1-\delta)^2}\right],$$

where $v_{j\tau}$ is the value of the quality of the leading good at the end of period τ. This establishes that consumers are better off, since ϕ increases. Now adding (A5) and (A6), a lower bound on the change in the sum of consumer plus current incumbent payoffs is given by the change in

$$\frac{(1-\phi)\Delta}{(1-\delta+\delta\phi)}+\phi\,\frac{\Delta}{(1-\delta)^2},$$

which is increasing in ϕ.

Current Incumbent: Finally, consider the current incumbent. If there is free entry so that $V_E = 0$, then (A4) implies that the incumbent is worse off since ϕ increases. On the other hand, suppose that there is a single potential entrant ($N = 1$) and that ϕ is fixed at some $\overline{\phi}$ (e.g., $c(\cdot)$ is finite only at $\overline{\phi}$). Then, (A4) implies that the incumbent is better off since V_E increases.

PROOF OF PROPOSITION 5:

For simplicity of exposition, we assume in the following proof that only the two highest-quality firms at stage $t.2$ are able to make offers to young consumers. It is straightforward to show that there is an equilibrium in the actual game in which no firms below the second-highest ever make offers and, more generally, that the offers of firms below the second highest are irrelevant to competition in any Markov perfect equilibrium in which the leading firm always wins and a losing firm never makes an offer that it would want to have refused. In any MPE from this class, all of the strategies are the same except for those of firms below the second-highest firm. We also assume that there is a unique least-cost compatibility level $\underline{b} \in \arg\min_{b\in[0,1]} k(b)$. This, too, is inessential, but simplifies the exposition.

In a Markov perfect equilibrium, the payoff-relevant state at the start of period t can be described as the triple $(B_t^e, B_t^1, B_t^2) \in [0,1]^3$, which denotes, respectively, the network sizes beyond one available to a potential entrant, the current leading firm, and the current second-highest firm. (Thus, $B_t^i = 1$ if the previous generation of consumers, the current old consumers, bought firm i's product, $B_t^i = B_{t-1}$, if it bought a product one step lower on the technological ladder, and $B_t^i = 0$ otherwise). The R&D strategy of a potential entrant is described by a function $\psi(B^e, B^1, B^2)$, which translates into an aggregate innovation probability $\phi(B^e, B^1, B^2)$ as described in Section II. The payoff-relevant state after stage $t.1$ is $(\overline{B}^1, \overline{B}^2) \in [0,1]^2$, which denotes, respectively, the network sizes available to the current leading firm and the current second-highest firm. If entry occurred at stage $t.1$, so that the leading firm at the start of stage $t.2$ is the new entrant, then $\overline{B}_t^1 = B_t^e$ and $\overline{B}_t^2 = B_t^1$. Otherwise, the leading firm at the start of stage $t.2$ is the continuing incumbent, so that $\overline{B}_t^1 = B_t^1$ and $\overline{B}_t^2 = B_t^2$. The leading firm's pricing and compatibility strategy is given by functions $b(\overline{B}^1, \overline{B}^2)$ and $p(\overline{B}^1, \overline{B}^2)$. The second-highest firm's strategy is given by $\underline{b}(\overline{B}^1, \overline{B}^2)$ and $\underline{p}(\overline{B}^1, \overline{B}^2)$. The consumers' acceptance decision at stage $t.3$ depends on the state $(\overline{B}^1, \overline{B}^2)$ and the price and compatibility offers made by the firms.

We construct an equilibrium in which the leading firm always wins in each state (in the end we show that it indeed has the incentive to do so). Working backward, in stage t.3, the young consumers accept an offer that offers them the highest expected surplus. Moreover, in equilibrium, those consumers must be indifferent between the leading firm's offer that they accept and the offer of the second-highest firm (otherwise the winning firm could deviate to a slightly higher price and still win). Given our assumption that losing firms do not make offers that they would prefer not to have accepted, the second-highest firm must make an offer that it is indifferent to having accepted. Finally, the second-highest firm must choose a compatibility level that maximizes the joint surplus it shares with the young consumers (otherwise, it could do better by changing its compatibility level and winning by giving young consumers a slightly greater surplus). To derive this compatibility choice, note that the second-highest firm wins next period only if it innovates, in which case its product will be two steps removed from, and therefore incompatible with, its current product. Therefore, if current young consumers deviate to buy from the second-highest firm, they will not expect any future compatibility benefits, nor will the sale benefit the second-highest firm in the future regardless of the good's compatibility level. Thus, the second-highest firm offers the least-cost compatibility level $\underline{b}$ at a price equal to its cost $\underline{k}$ (that is, $\underline{b}(\overline{B}^1,\overline{B}^2) = \underline{b}$ and $\underline{p}(\overline{B}^1,\overline{B}^2) = \underline{k}$ for all $(\overline{B}^1,\overline{B}^2)$).

The leading firm wins by matching the surplus offered by the second-highest quality firm. Thus, its price-compatibility offer (p,b) satisfies

$$p = \hat{p}(\overline{B}^1_t,\overline{B}^2_t,b) \equiv (1+\delta)\Delta + [v(1+\overline{B}^1_t) + \delta G(\phi(b,1,0))] - [v(1+\overline{B}^2_t) + \delta v(1)] + \underline{k}.$$

The leading firm's chosen compatibility level b will maximize its expected discounted profits, $[\hat{p}(\overline{B}^1_t,\overline{B}^2_t,b) - k(b) + \delta V_I(b,1,0)]$, subject to the antitrust constraint, which means that regardless of $(\overline{B}^1_t,\overline{B}^2_t)$, it chooses

$$\text{(A7)} \qquad b^* \in \arg\max_{b\geq\alpha} [\delta G(\phi(b,1,0)) - k(b) + \delta V_I(b,1,0)].$$

Thus, the leading firm's strategy at stage t.2 is $p(\overline{B}^1,\overline{B}^2) = \hat{p}(\overline{B}^1,\overline{B}^2,b^*)$ and $b(\overline{B}^1,\overline{B}^2) = b^*$. (In the Markovian spirit, we assume that firms that have the same compatibility incentives make the same choice in the event of indifference.)

Now, go back to the beginning of period t with state (B^e_t, B^1_t, B^2_t). If there is entry, we proceed to stage t.2 with state $(\overline{B}^1_t,\overline{B}^2_t) = (B^e_t, B^1_t)$, and otherwise we proceed with state $(\overline{B}^1_t,\overline{B}^2_t) = (B^1_t, B^2_t)$. Thus, the firms' values at the beginning of period t are

$$\begin{aligned}\text{(A8)} \quad V_E(B^e_t,B^1_t,B^2_t) = {} & \delta V_E(b^*,1,0) \\ & + u(\phi(B^e_t,B^1_t,B^2_t))\,[\hat{p}(B^e_t,B^1_t,b^*) - k(b^*) \\ & + \delta(V_I(b^*,1,0) - V_E(b^*,1,0))] - c(\phi(B^e_t,B^1_t,B^2_t);\end{aligned}$$

$$\begin{aligned}\text{(A9)} \quad V_I(B^e_t,B^1_t,B^2_t) = {} & (1-\phi(B^e_t,B^1_t,B^2_t))[\hat{p}(B^1_t,B^2_t,b^*) - k(b^*) + \delta V_I(b^*,1,0)] \\ & + \phi(B^e_t,B^1_t,B^2_t)\,\delta V_E(b^*,1,0).\end{aligned}$$

As for the innovation probability in state (B^e, B^1, B^2), it will be

$$\text{(A10)} \qquad \phi(B^e,B^1,B^2) = \Phi(\hat{p}(B^e,B^1,b^*) - k(b^*) + \delta[V_I(b^*,1,0) - V_E(b^*,1,0)]).$$

We now argue that the leading firm indeed prefers to win at stage t.2 in every state $(\overline{B}^1_t,\overline{B}^2_t)$. If it deviates from the strategy above and does not win, and instead the second-highest firm wins, the leading firm receives zero profits in the present period, and its continuation value is $V_I(0, \underline{b}, 1)$. We show that the leading firm does better than that by winning with a compatibility level of 0 at a price just below $\hat{p}(\bar{B}_t^1, \bar{B}_t^2, 0)$, which implies that it prefers to win with its optimal compatibility offer of $b(\bar{B}_t^1, \bar{B}_t^2) = b^*$ and $p(\bar{B}_t^1, \bar{B}_t^2) = \hat{p}(\bar{B}_t^1, \bar{B}_t^2, b^*)$. This follows from the following five observations:

(a) The leading firm does better in the present period if it wins with this zero-compatibility offer than it does by losing, since for any $(\bar{B}_t^1, \bar{B}_t^2)$,

$$\hat{p}(\bar{B}_t^1, \bar{B}_t^2, 0) - k(0) \geq (1+\delta)\Delta - [k(0) - \underline{k}] - (1+\delta)[v(2) - v(1)] > 0$$

under our assumption that $\Delta > v(2) - v(1) + [k(0) - \underline{k}]/(1+\delta)$.

(b) The leading firm's continuation payoff beginning in period $t+1$ is independent of whether it wins in period t when entry occurs in period $t+1$.

(c) When entry does *not* occur in period $t+1$, the leading firm's continuation payoff beginning in period $t+1$ is higher if it wins in period t with this zero-compatibility offer than if it loses in period t. To see this, note that the firm's profits in period $t+1$ are increased, since $\hat{p}(1,0,b^*) > \hat{p}(\underline{b}, 1, b^*)$, and its continuation payoff starting in period $t+2$ is $V_I(b^*, 1, 0)$ regardless of whether it wins or loses in period t.

(d) By (A10), the probability of entry in period $t+1$ if the leading firm wins with this zero-compatibility offer in period t is weakly lower than if it loses in period t, since $\hat{p}(0,1,b^*) \leq \hat{p}(0,0,b^*)$ and the innovation supply function $\Phi(\cdot)$ is nondecreasing.

(e) Following winning with this zero-compatibility offer, the leading firm prefers no entry to entry. To see this, note that

$$\hat{p}(1,0,b^*) - k(b^*) + \delta[V_I(b^*,1,0) - V_E(b^*,1,0)] \geq \hat{p}(b^*,1,b^*) - k(b^*)$$
$$+ \delta[V_I(b^*,1,0) - V_E(b^*,1,0)] = w > 0,$$

where the left side of this inequality is the incumbent's continuation payoff without entry, less its continuation payoff with entry, and w is the entrant's equilibrium innovation prize, which must be strictly positive. (Supposing in negation that $w \leq 0$, the equilibrium innovation rate would have to be $\Phi(w) \leq \Phi(0) = 0$, but then (A8) and (A9) would yield equilibrium values $V_E(b^*,1,0) = 0$ and $V_I(b^*,1,0) = (1+\delta)^{-1}[\hat{p}(1,0,b^*) - k(b^*)] \geq (1+\delta)^{-1}\{[\hat{p}(1,0,0) - k(0)] + \delta[V_I(0,1,0) - V_I(b^*,1,0)]\} > 0$, where the weak inequality follows from (A7) and the strict inequality follows since $V_I(0,1,0) \geq V_I(b^*,1,0)$, and, from observation (a), $\hat{p}(1,0,0) - k(0) > 0$. This would imply that the innovation prize w would be strictly positive, a contradiction.)

Finally, it follows from the above that, on the equilibrium path, the state at the start of every period is the same and we focus on the equilibrium state $(B^e, B^1, B^2) = (b^*(\alpha), 1, 0)$, where $b^*(\alpha)$ is the highest equilibrium compatibility level. Substituting this state in expressions (A9) and (A8) yields $V_I = V_I(b^*(\alpha), 1, 0)$ and $V_E = V_E(b^*(\alpha), 1, 0)$, which satisfy equations (VI*) and (VE*) in Section II for the profit functions $\pi_m(\phi, \alpha)$, $\pi_I(\phi, \alpha)$, $\pi_E(\phi, \alpha)$ described in (8), and the equilibrium innovation rate $\phi = \phi(b^*(\alpha), 1, 0) = \Phi(\pi_E(\phi, \alpha) + \delta(V_I - V_E))$. Furthermore, the highest equilibrium compatibility choice $b^*(\alpha)$ is nondecreasing in the antitrust constraint α. Indeed, if $b^*(\alpha') \geq b^*(\alpha'')$ for $\alpha' < \alpha''$, then we must have $b^*(\alpha') \geq \alpha'' > \alpha'$. But, then, an equilibrium for parameter value α' would remain an equilibrium for parameter value α'', and so $b^*(\alpha'') = b^*(\alpha')$. The result then follows from the argument in the text.

REFERENCES

Aghion, Philippe, Christopher Harris, Peter Howitt, and John Vickers. 2001. "Competition, Imitation and Growth with Step-by-Step Innovation." *Review of Economic Studies*, 68(3): 467–92.

Aghion, Philippe, and Peter Howitt. 1992. "A Model of Growth through Creative Destruction." *Econometrica*, 60(2): 323–51.

Evans, David S., and Richard Schmalensee. 2002. "Some Economic Aspects of Antitrust Analysis in Dynamically Competitive Industries." In *Innovation Policy and the Economy* Volume 2, ed. Adam B. Jaffe, Josh Lerner, and Scott Stern, 1–49. Cambridge, MA, and London: MIT Press for the National Bureau of Economic Research.

Fisher, Franklin M., and Daniel L. Rubinfeld. 2000. "*United States v. Microsoft*: An Economic Analysis." In *Did Microsoft Harm Consumers? Two Opposing Views*, ed. David S. Evans, Franklin M. Fisher, and Richard Schmalensee, 1–44. Washington, DC: AEI-Brookings Joint Center for Regulatory Studies.

Fudenberg, Drew, and Jean Tirole. 2000. "Pricing a Network Good to Deter Entry." *Journal of Industrial Economics*, 48(4): 373–90.

Gilbert, Richard J., and Michael L. Katz. 2001. "An Economist's Guide to *U.S. v. Microsoft*." *Journal of Economic Perspectives*, 15(2): 25–44.

Gilbert, Richard J., and Steven C. Sunshine. 1995. "Incorporating Dynamic Efficiency Concerns in Merger Analysis: The Use of Innovation Markets." *Antitrust Law Journal*, 63: 569–601.

Grossman, Gene M., and Elhanan Helpman. 1991. "Quality Ladders in the Theory of Growth." *Review of Economic Studies*, 58(1): 43–61.

Hopenhayn, Hugo, Gerard Llobet, and Matthew Mitchell. 2006. "Rewarding Sequential Innovators: Prizes, Patents, and Buyouts." *Journal of Political Economy*, 114(6): 1041–68.

Hunt, Robert M. 2004. "Patentability, Industry Structure, and Innovation." *Journal of Industrial Economics*, 52(3): 401–25.

Lee, Tom, and Louis L. Wilde. 1980. "Market Structure and Innovation: A Reformulation." *Quarterly Journal of Economics*, 94(2): 429–36.

Loury, Glenn C. 1979. "Market Structure and Innovation." *Quarterly Journal of Economics*, 93(3): 395–410.

Milgrom, Paul, and John Roberts. 1994. "Comparing Equilibria." *American Economic Review*, 84(3): 441–59.

Milgrom, Paul, and Chris Shannon. 1994. "Monotone Comparative Statics." *Econometrica*, 62(1): 157–80.

O'Donoghue, Ted, Suzanne Scotchmer, and Jacques-Francois Thisse. 1998. "Patent Breadth, Patent Life, and the Pace of Technological Progress." *Journal of Management and Strategy*, 7: 1-32.

Ordover, Janusz A., and Robert D. Willig. 1981. "An Economic Definition of Predation: Pricing and Product Innovation." *Yale Law Journal*, 91: 8–53.

Quirmbach, Herman C. 1993. "R&D: Competition, Risk, and Performance." *RAND Journal of Economics*, 24(2): 157–97.

Reinganum, Jennifer F. 1989. "The Timing of Innovation: Research, Development, and Diffusion." In *Handbook of Industrial Organization* Volume 1, ed. Richard Schmalensee and Robert D. Willig, 849–908. Amsterdam: Elsevier Science, North-Holland.

Segal, Ilya. 1999. "Contracting with Externalities." *Quarterly Journal of Economics*, 114(2): 337–88.

Segal, Ilya, and Michael Whinston. 2005. "Antitrust in Innovative Industries." National Bureau of Economic Research Working Paper 11525.

Tirole, Jean. 1988. *The Theory of Industrial Organization*. Cambridge, MA: MIT Press.

Whinston, Michael D. 2001. "Exclusivity and Tying in *U.S. v. Microsoft*: What We Know, and Don't Know." *Journal of Economic Perspectives*, 15(2): 63–80.