

Rewarding Duopoly Innovators: The Price of Exclusivity*

Hugo Hopenhayn[†] and Matthew Mitchell[‡]

March 20, 2012

Abstract

We study an environment where a duopoly develops innovations that build on one another and compete. Since the quality of innovations is unobserved, rewards take the form of rights to produce the resulting products. There is a tradeoff between encouraging one firm to work on its innovations by granting it promised rights, and the fact that those rights deteriorate the rights of its competitors. We study constrained efficient allocations and show that they in monopolization: eventually one firm is promised nearly everything, and the competitor is almost completely ignored. This occurs because backloading rewards is an efficient incentive device. We interpret our results in three ways. First, if the firms are able to contract ex ante, our allocations can be interpreted as a patent pooling arrangement between the firms. Second, if (no contract ex ante, policy, state dependent) We argue that the backloading motive is different from existing ones in the literature. It both explains further the state-dependent protection results computed in [1], generating heterogeneity in patent protection in the absence of heterogeneity in innovation opportunities. Finally, one can interpret the optimal evolution of the duopoly as competition “for the market,” where backloading implies that optimal competition for the market eventually leads to nearly permanent monopolization by one firm.

*We thank participants at numerous seminars and conferences for helpful comments. This paper was previously titled “Optimal Patent Policy with Recurrent Innovators.”

[†]UCLA.

[‡]University of Toronto.

1 Introduction

This paper studies the optimal reward structure for a sequence of innovations generated by a duopoly of firms. The appropriate reward for innovation has long been considered an important issue by economists. In describing the benefits of patents as a reward mechanism, [13] wrote that patents are an effective reward “because the reward conferred by it depends upon the invention’s being found useful, and the greater the usefulness, the greater the reward.” This paper builds on that general principle, that incentives dictate that rewards take the form of a stake in the profitability of the innovation. When innovations compete, however, there is a trade-off: rewarding one innovation decreases the rewards to competing innovations. Often these competing innovations come from a few firms that interact repeatedly. We address the question of what forms of rewards best elicits innovation from those firms. Microsoft, for instance, argued that its strong position in the market was part of a sound policy in supporting innovation; we address the trade-off between potential benefits from rewarding one firm greatly, and the cost that imposes in terms of lost innovation from other firms. We show how the oligopoly context influences optimal reward for innovation. We find that rewards are backloaded, in the sense that the same opportunity receives greater rewards when it comes on the heels of other innovations by the same firm. This backloading leads to a market which is eventually strongly skewed toward one of the competing firms.

Our approach is to study constrained efficient allocations in a dynamic environment where innovations from the competing firms build on one another. Moral hazard precludes rewarding with a cash prize; instead, rewards must be earned through the allocation of preferential treatment, such as a patent in the product market. This is the sense in which our work follows the ideas of Mill. Because of the cumulative nature of the research, allowing one innovator to profit in the product market necessarily restricts what can be offered to the other innovator, since they compete in this common market. As a result, rights are scarce. Our optimal allocations allocate rights in the product market in order to encourage innovation. We study an environment where there are no static monopoly pricing distortions, so that we can focus on the trade off between rewarding innovations from the different innovators.

We show that the optimal allocation leads to identical opportunities being treated differently depending on the innovator’s past history of contributions. The allocations admit a variety of interpretations. The first is normative.

Recent papers on optimal patents, beginning from [17], stress that inherent heterogeneity across different innovation opportunities may lead to different rewards for different innovators. Here there is no heterogeneity built into the structure; all innovation opportunities are technologically identical. The rewards, however, are history dependent, leading to ex post heterogeneity in the reward for different innovations, and as a result, heterogeneity in the degree of innovation over time. This idea is familiar from recent work by [1]. In that paper, the authors use a growth theory structure similar to [2] and consider policies that change depending on the quality differential between the firm’s most recent innovations. [1] compute numerically the best policies within a particular class and show that they are backloaded, in the sense that firms that succeed repeatedly get increasing protection. Our paper considers a more general class of policies in a more abstract environment. We also generate strongly backloaded policies. Although our structure does not nest theirs, the intuition about backloading that we develop is new and applies to their environment. It therefore helps develop further understanding of their numerical results.

Although a normative interpretation of our results is possible, or allocations are identical to the ones that the firms would agree to were they to sign complete licensing contracts ex ante, in a patent pooling arrangements. The contract allows for growth of the patent pool as new innovations are developed. We show that at every point in time the pool’s profits are assigned entirely to one of the two firms. The contract specifies under what conditions the pools proceeds should switch from one firm to the other.

An alternative view of a policymakers job, then, is to ensure that outcomes mimic the ex ante contract design, in cases where complete ex ante contracting may not be feasible. Since the contract’s role is to decide when profits should shift between the innovating firms, one can interpret this regulation as regulating competition “for the market.”¹ [6] argue that competition for the market is as important as competition in the market. Our model focuses entirely on competition for the market; there is effectively no cost of market structures at a particular point in time. Our backloading result provides an interesting interpretation of optimal competition for the market. Backloading generates incentives to innovate by promising stronger and stronger monopolization as a firm generates more innovations. In the

¹In this sense our paper is related to the larger set of papers on regulation and innovation, for instance in [18] and [7].

long run, backloading implies that the optimal regulation for the market extinguishes itself: the promise to one firm is essentially so great that its competitor is effectively foreclosed. Although such foreclosure is costly to the planner, the promise of foreclosure motivates innovation in the interim.

Counterintuitively, monopolization occurs even when the marginal benefit of preferential treatment for the monopolizing firm is zero at monopoly. The motive for backloading is different from the standard backloading intuition from [4] and [12]. In those papers, backloading is beneficial because not only does it generate strong incentives late, but also because it generates strong incentives early, as the agent works hard to reach the point where the backloaded incentives kick in. Here we introduce a different motive for backloading which is tied to the trade off inherent in rewarding innovators with competing property rights. Since the planner can provide protection for a given agent at a given time for multiple innovations, it is useful to push preference to the later period where it is applicable to more than one innovation. This leads to protection whose duration increases with additional successes.

Our model can be described concisely. Two innovating firms randomly receive opportunities to generate social value (innovations) at independent rates. The innovations build on one another. Our approach is to study (information constrained) optimal allocations. The planner rewards the innovators by offering, at each instant, a given innovator the opportunity to profit from some subset of past innovations. The planner, then, needs to decide how to determine the allocation of rights at each state and history. This is potentially a very complicated problem; our results use a structure that is amenable to recursive methods. The planner faces a sort of extensive-intensive tradeoff: the more the planner promises rights to one firm, the more that firm takes advantage of its opportunities (the intensive margin), but the less rights are available to reward other firms (the extensive margin).

In order to focus on the dynamic tradeoff between rewarding different firms, the bulk of our results focus on the case where there is no static distortions. Under such an environment, the optimal allocation pools all innovations and gives the rights to those innovations to a single entity. We show this formally in section 3. Section 4 characterizes optimal policies when there is full exclusion at each point in time. The optimal policy trades off the extensive and intensive margins. As a particular innovator has more successes, his innovations are promised more and more. This varying state of promised rights generates the heterogeneity in otherwise homogenous ideas

and innovators: after numerous successes, a firm is favored and therefore does more with each opportunity. We show that this leads inevitably to states where the planner takes almost no advantage of innovations other than one firm that has been particularly successful; rights eventually enter a state where nearly the entire future has been promised to one firm, and the other firm is (nearly) completely foreclosed, its ideas virtually unused.

As a function of the promises the planner has made, the profits themselves evolves in a stark way. The firm with the greater duration promise gets the profits from all the cumulative innovations. As a result, when the future promise is skewed sufficiently toward one firm, even an arrival of an idea by the competitor leaves the leader with greater promised rights. Firms with sufficiently low promise get no immediate profits from their innovation; they are required to put them into a “pool” from which, initially, only the firm with the greater promise profits. The lagging firm’s payoff to generating the innovation is that the promise becomes less skewed, moving the state closer to its favor, where it will gain rights to all of the pooled patents. The interpretation most in keeping with the traditional literature on patent policy is to imagine this being the result of a patent policy that affords a patent with the power to exclude all other innovations, but which is only granted to followers after a sufficient collection of innovations are developed. Alternatively, one can interpret this as reflecting stark rules for the flow of profits from the pool, or as regulatory treatment that strongly favors one firm until another firm generates a sufficient collection of innovations.

To show that compete exclusion is not fundamental to our results, we consider a second regime where complete exclusion is not available to the planner. We call this alternative regime "incomplete exclusion rights"; the planner can grant a firm preferential treatment at a single instant for any innovations from a sequence of innovations by that firm that have arrived consecutively, without an intervening innovation by the other firm. In other words, whenever the competitor is granted some preferential treatment, the incumbent leading firm loses all rights to preferential treatment. This is comparable to a standard notion of patents used in models of cumulative innovation, where a leading firm can maintain market position by filing for new patents, until a competing improvement makes a new firm the market leader. The patents compete rather than being pooled.

We show that under the exclusive rights regime, we still get backloading; rewards increase with successive innovations by the incumbent. Moreover, the chance of a competitors innovation being implemented (and the incum-

bent being ousted) is declining with successes by the incumbent. This is familiar from the state dependent policies that [1] compute as optimal policies in the step-by-step model they study. We show the sense in which this structure is driven by a backloading incentive present in their paper as well.

We interpret this as simple patent system without licensing, and show that it can be decentralized through a system of non-infringing patents with an associated fee. The twist is that, since the optimal allocation forecloses the market, the most recent purchaser of a patent also has the right to pay an additional fee which disallows any more patents to be filed by the competition. With this decentralization, the authority need not observe anything, or ask for any reports; it simply allocates rights to anyone who pays the appropriate fees. Until the foreclosure fee is paid, the patent authority offers patents that are narrow, in the sense that they offer no rights to exclude other innovations; they simply give the innovator the right to solely market their own innovation. The foreclosure fee broadens the patent so that it excludes all future work.

Our paper links recent literature on the role of information constraints in generating particular features of the optimal reward structure with the literature that studies protection in particular growth theory contexts. In addition to [17], papers in the former category include [?], who also generate a menu of patents for different types of innovations, and [10], where optimal policy is a menu of lengths and breadths. [9] and [14] apply these methods to dynamic environments, based on the quality ladder structure in [16]. In those papers the set of innovators is large, so there are never repeat innovators; they therefore can not address the issues of oligopoly, state dependent rewards, and the evolution of market structure that we study here. Like the model of [1], our paper allows us to study repeat innovators and their treatment as a function of their history of innovations. Unlike papers in this spirit like [11], [19], and [5], we do not consider the role of market signals in generating out optimal allocations.

2 Competition and Innovation

2.1 Static Competition

The quality ladder structure follows the one explored in the patent literature in papers such as [16] and [9]. We begin with a model of static oligopoly

competition in a quality ladder. Suppose a collection of firms sells products of various quality levels. A single consumer either takes an outside option (normalized to zero) or purchases one physical unit of quality q , choosing which variety in order to maximize $q - p$, where p is the price paid for the quality q variety.² There are no costs of production. We take competition to be Bertrand, so that the leading edge product is always the one sold in equilibrium, and the social surplus at any point in time is the quality q either in the form of profits for the firm selling the leading edge product, or as consumer surplus if $p < q$. Industry profits are summarized by the difference between the quality levels of the highest and second highest quality level that is sold.

Our model abstracts from static monopoly costs. The implication of such costs is important, but this framework has little to say about them beyond what is known. What is new here is the role that skewed rights between the firms impacts the firms incentives to innovate, which we focus on here by abstracting from other considerations that are surely relevant in practice, but has been well studied elsewhere.

2.2 Innovation

Suppose that a firm has an opportunity to generate an innovation. We term this opportunity an idea. The innovation will be an improvement of size Δ upon the highest quality product currently available; this is the sense in which the firms are working on a common agenda. Ideas can be turned into innovations of size Δ in exchange for research cost $c(\Delta)$. Here Δ (and therefore $c(\Delta)$) is the non-verifiable feature that necessitates patents, in keeping with Mill's comments. It is well known in the literature (and very intuitive) that, if the degree of innovation cannot be verified by the planner, the innovators cannot be rewarded with transfers, since they could claim the transfer and not pay the costs of innovation.

The planner can use the promise of rights to sell in a market to successfully induce effort.³ The leading edge product (i.e. the most recent innovation) sells for a price equal to its quality differential over the other firm's best

²As is usual in this sort of model, in the event of a tie, the higher quality product is chosen.

³See Hopenhayn, et al. (2006) for more on this issue in the cumulative context, in a quality ladder model like this one.

product, since that trailing product will be provided at cost (zero) in the pricing game, and the leader charges the quality differential.

As in [9] we study policies that are described by contingent rights for a given innovation. This takes the form of a mapping from innovations to rights holders.⁴ Formally, for a list of arrivals of ideas $M = \{1, 2, \dots, M\}$, the planner chooses at each time t set $P_t^i \subseteq M$ of innovations for which the agent i receives exclusive rights. Define by \bar{P}_t^i the greatest element in the set. This grant of rights can be made contingent on the entire history; solving that history dependent problem is the topic of the next section. For now we focus on the implications of the granting of rights for an innovator's investment decision.

Any choice of rights implies profits for each firm in the amount B_t^i , where $B_t^i = \max\{\bar{P}_t^i - \bar{P}_t^j, 0\}$. One can interpret this as the amount of breadth of protection afforded that innovator relative to the existing innovations. This is the set of future instants during which the current size of the current innovation impacts the firm's profits. In all such instants, the firm's profits will be higher by Δ units if it makes an innovation of that size in the current innovation. Therefore when innovator i chooses Δ for innovation m , he solves

$$\Delta(d) = \arg \max_{\Delta} d\Delta - c(\Delta)$$

where

$$d = E \int e^{-rt} I\{\bar{P}_t^j < m\} dt$$

The features of the contract, for the purposes of the investment decision, can be summarized by the planner's promise of *expected discounted length of time* $d \geq 0$ during which the innovator will be given preferential treatment for an innovation made under that idea. This simplification is a key feature that allows the complete contingent rights contract to be tractable in a recursive way we introduce below.

Since the value of d , which we term the *duration* of preference, is in present discounted terms, it might come in many ways, for instance, a T period patent (where preference is guaranteed for all T periods) would have $d = (1 - e^{-rT})/r$. We use the language of duration to describe recursively how

⁴Without loss of generality we let preference for a given innovation be allocated to a single innovator at any instant; this is without loss because changing preference over time can effectively "split" preference for a given innovation across innovators. Such a split will, in addition, not be optimal.

the optimal policy proceeds, considering arbitrary duration policies, which may be contingent on future arrivals as well as the passage of time. A patent that offered T periods of protection for sure, followed by T' units of additional protection with probability $1/2$ would have $d = (1 - e^{-rT})/r + \frac{1}{2}e^{-rT}(1 - e^{-rT'})/r$. Since the planner can choose a preference policy at every instant, this duration can be delivered in any contingent way, evolving over time or with later arrivals of innovators, and with the identity of the innovator that arrives with an idea. Of course, since (discounted) time is not unbounded, the maximum possible promise of sure preferential treatment forever is $1/r$. This dynamic budget constraint of the planner's incentive tool is the key feature of the model.

Since the innovation will permanently increase the highest quality, every innovation yields $\Delta/r - c(\Delta)$ additional units of present discounted social surplus. We therefore can denote the benefit from the allocation of d units of duration as $R(d) = \Delta(d)/r - c(\Delta(d))$. Note that

$$\begin{aligned} R'(d) &= \Delta'(d)/r - c'(\Delta(d))\Delta'(d) \\ &= \Delta'(d)(1/r - d) \end{aligned}$$

Where the second line uses the fact that $c'(\Delta) = d$ by the agents FOC.

Now by the implicit function theorem it must be the case that

$$\Delta'(d) = \frac{1}{c''(\Delta)}$$

We then have that

$$R''(d) = \Delta''(d)(1/r - d) - \Delta'(d)$$

In order for R to be concave, then, we need the third derivative of c to be smaller than some positive bound.

Our model has no static distortions, so that $R(d)$ is maximized at $1/r$, we discuss in section 6 the possibility that $R(d)$ is not maximized at $1/r$, which can be interpreted as static costs of monopoly. Focusing on the case with no static distortions, however, is interesting for at least two reasons. First, it highlights the role of the dynamic force that we study, namely the scarcity of rights when competition is for the market only, without any other source of inefficiency. Further, the work of [8] suggests that a planner who allocates patent rights together with the ability to regulate the strength of

preference per period (for instance through patent breadth or direct price controls) will choose, in many circumstances, a long, narrow patent in the single innovation context.

3 Optimality of Complete Exclusion Rights

We now turn to studying the optimal allocation of duration across histories of arrivals. In this section we introduce the structure of the arrival of ideas, describe the full planner's problem, and show that it can be solved in a relatively simple way, in order to set the stage for the key results of the next section where we characterize the solution.

There is continuous time and an infinite horizon. We focus on the case where there are two agents (which we sometimes call firms or innovators) and a principal (or planner). Below we discuss extension to more firms, including an explicit extension for a specific case described below. Each firm receives ideas with independent Poisson arrival rate λ .⁵ Although arrivals are taken to be technologically identical, in order to highlight the endogenous emergence of heterogeneous treatment without heterogeneity in arrivals, there is a sense in which the model can be thought of as delivering small and large innovations; an unusually large innovation can be thought of as the rare arrival of several consecutive ideas in a short period of time to one innovator.

Since the planner cannot observe inputs or outputs, he cannot simply pay for the effort. The planner can, however, choose market structure, and, in turn, deliver market duration that generates profits and in turn innovation in line with the analysis of the previous section.⁶ We assume that the planner's history dependent market structure choice is made with full commitment at time zero. When an idea arrives, the contract prescribes future instants during which the innovator will be assigned rights that include the innovation

⁵Although we abstract from different λ across the innovators, nothing changes if λ differs across agents or differ for the firms based on which one had the last idea. We discuss this in the extensions section.

⁶This is not the only one that can be mapped into our structure. For instance, the planner could be a firm selling a product, where customers can be induced to buy with promises of future good treatment. An airline can reward today's purchase with promises of future good seat assignments available only to frequent flyers. These rewards are scarce; only a limited number of customers can be allocated the good seat assignment. The firm trades off using current promises to encourage sales against the fact that current promises restrict the possibility of later promises to other potential customers

stemming from that idea. From the prior section, the planner's expected discounted future payoff from a promise of d units of time of future preference to an arriving idea as $R(d)$. On the other hand, duration is limited by the planner's prior promises of duration to previous innovators. In particular, while the planner can effectively deliver duration to more than one innovation of one innovator, he faces a choice between rewarding one innovator or another. Therefore the stock of available rewards available to a particular innovator is curtailed by the amount that is promised to the competition.

To see how planner's choices translate into an evolution of promised duration for innovators, suppose that the planner offers the most recent innovator rights to all past innovations until the next arrival. Denote by \hat{d} the duration this offers to the new duration. When the other firm has an arrival, duration therefore drops to $1/r - \hat{d}$. Therefore \hat{d} solves

$$r\hat{d} = 1 + \lambda(1/r - \hat{d} - \hat{d})$$

or $\hat{d} = \frac{r+\lambda}{r(r+2\lambda)} > 1/2r$. The planner can offer every arrival this duration. If the planner offers any innovator more for one of its arrivals, however, it will curtail the ability of the planner to reward the other innovator. We study this trade off recursively in the next section.

3.1 Dynamic Program

At any instant, the planning problem is summarized by an outstanding duration promised to each of the innovators for prior work; the planner's "stock" of available rights to offer is determined by these values. Since one innovator can be granted rights simultaneously for multiple prior innovations (by setting the rights of the competitor low enough), one can think of this as the largest promise that is owed across all prior innovations; all other promises can be kept with a fraction of d , since a given innovator can be granted preference for multiple innovations at once.⁷

If an innovation by firm 1 arrives, the planner offers preferential treatment for that new innovation for duration d_1^n . It continues preferential treatment for the innovator's previous innovation (or innovations), which are owed d , for duration d_1^c . The planner will then enter the next instant with promise equal to the maximum of d_1^n and d_1^c , since the outstanding duration that

⁷It will turn out that all past innovations with any promise will get identical promises in the optimal allocation.

cannot be allocated to other firms is the larger of those promises. We will argue below that optimally $d_1^n = d_1^c$, and therefore we will eventually just use d_1 to denote the new promise. If innovator two has the next idea, then innovator 1's duration becomes d_2 . We keep track of the duration promise to the two firms by d and \underline{d} ; we show below that it is sufficient to track only one. In the interim, we speak generically about duration as d ; everything is symmetric across the innovators, so all statements apply equally to \underline{d} . In order to make everything completely symmetric, we refer to the promise to firm two in the event that firm one arrives by \underline{d}_2 , and so on.

In addition to duration promises, the planner must also decide how to allocate duration in intervening periods. In particular, the planner allocates a fraction x of the next dt instants to innovator one if no idea arrives. Although preference is the fundamental choice the planner makes, our study of the problem focuses on the promises of duration that the planner makes, and uses x and future promises to ensure past promises are kept.

If nothing arrives, the planner may change the duration promise by \dot{d} . We include the possibility for completeness; it will turn out that the optimal policy will have $\dot{d} = 0$, so the planner never uses the option to make changes after no arrival of an idea takes place. The dynamic program is, then,

$$rV(d, \underline{d}) = \max_{\substack{d_1^n, d_1^c, d_2, \dot{d}, x \\ \underline{d}_1, \underline{d}_1^n, \underline{d}_2, \underline{\dot{d}}, \underline{x}}} \left\{ \begin{array}{l} \lambda(R(d_1^n) + V(\max\{d_1^n, d_1^c\}, \underline{d}_2) - V(d, \underline{d})) + \\ \lambda(R(\underline{d}_1^n) + V(d_2, \max\{\underline{d}_1^n, \underline{d}_1^c\}) - V(d, \underline{d})) + \\ V_1(d, \underline{d})\dot{d} + V_2(d, \underline{d})\underline{\dot{d}} \end{array} \right\} \quad (1)$$

$$s.t. \quad (2)$$

$$rd = x + \lambda(d_1^c - d) + \lambda(d_2 - d) + \dot{d} \quad (3)$$

$$r\underline{d} = \underline{x} + \lambda(\underline{d}_1^c - \underline{d}) + \lambda(\underline{d}_2 - \underline{d}) + \underline{\dot{d}} \quad (4)$$

The first line of the maximand is the case where the current innovator, promised d for prior innovations, arrives with a new idea. The second line is the case where the competitor arrives with an idea. The final line is when nothing arrives. There are also the domain constraints:

$$0 \leq \max\{d_1^n, d_1^c\} + \underline{d}_2 \leq 1/r$$

$$0 \leq d_2 + \max\{\underline{d}_1^n, \underline{d}_1^c\} \leq 1/r$$

$$0 \leq x + \underline{x} \leq 1$$

The constraints in (1) guarantee that the planner actually does deliver d and is critical to understanding the problem. Given a current promise d , the

fraction x of the current period is allocated to the first innovator. Unless the constraint is not binding, it is clearly optimal for $x + \underline{x} = 1$, since in that case duration is scarce and should not be thrown away. The innovator gains $d_1^c - d$ if the innovator comes up with a new idea, and moves to d_2 if the other innovator has an idea and is implemented. This constraint also shows a key difference between this model and one with a sequence of innovators, as studied in Hopenhayn, et al (2006). In both models, duration promises to the current innovator make the PK constraint tighter in the future. In simple terms, increasing duration today makes the planner less able to make promises to other agents in the future. However, to the extent that future innovations come from the same source, greater duration does not preclude future innovations, and therefore is not making the PK constraint tighter in the future in those states. This impact of duration on the tightness of the PK constraint is formally the fundamental difference of this problem from ones with innovators who never recur.

Since greater d only makes the feasible set of possible choices of d_1 and d_2 smaller, it is immediate that $V(d, \underline{d})$ is weakly decreasing in each argument. This in turn implies that d_1^c can always be taken to be at least as big as d_1^n ; if d_1^c were less, raising it and offsetting the increase by lowering d to maintain promise keeping always does at least as well, and strictly better if V is strictly decreasing. Similarly, for $d_1^c > d_1^n$, reducing d_1^c at the margin is identical to increasing \dot{d} , and therefore we can let $d_1^c = d_1^n \equiv d_1$. However, in the modified program where $d_1^c = d_1^n \equiv d_1$ the envelope condition is⁸

$$V_1(d, \underline{d}) + \frac{1}{r + 2\lambda} V_{11}(d, \underline{d}) \dot{d} = \mu(d, \underline{d})$$

where $\mu(d)$ is the Lagrange multiplier on the PK constraint for d . This coincides with the first order condition for \dot{d}

$$V_1(d, \underline{d}) = \mu(d, \underline{d})$$

when $\dot{d} = 0$. We therefore have the following lemma.

Lemma 1. *Suppose V is concave. Then $d_1^c = d_1^n$ and $\dot{d} = 0$*

The interpretation of $d_1^c = d_1^n$ is that new arrivals always extend duration promises for all prior inventions for the incumbent. If the incumbent had

⁸Subscripts denote derivatives.

invented a drug that treats a given disease, and then came up with an improvement that treats the disease somewhat more effectively, it both obtains preferential treatment for the improved drug for d_1 , and from the point of the improvement gets d_1 units of preference for the basic treatment as well. The intuition for why this is optimal is identical to the reason why there is no statutory limit to preferential treatment ($\dot{d} = 0$): there is no benefit to lowering past duration promises, given that you are offering d_1^n to the innovator for his new innovation, and therefore this is an efficient time to deliver duration to satisfy the outstanding promise of d on the initial innovation. The fact that the planner's payoff from the improvement is independent of the treatment of the basic drug is crucial to that logic. The planner is better off delivering duration always contingent on an arrival of the competitor, so as to deliver the most duration to the competitor if he is next to arrive.

This logic implies that the planner is always offering rights to all of an innovator's innovations if he is granting rights to any of the innovator's innovations. Since protecting innovations of the other innovator has no impact, one can take this to mean that any point where an innovator is given exclusive rights to any innovation, that innovator is given exclusive rights to the entire history of innovations. Because this feature of the optimal allocation is strong, we also study below a weaker form, when protection ends whenever a firm's competitor arrives with an innovation. Here, the leading edge product is able to exclude all competing products (except the outside good with quality normalized to zero). One might imagine that older product infringe on newer products, but not vice-versa. In [16] this is referred to as "lagging breadth." Then at any time that the firm is the market leader, it profits from all of its past improvements (and all of the competitors, but this is a transfer and does not impact incentives). This maximizes the reward delivered to the innovator for past innovations.

We now verify that V is in fact concave. If it is, then imposing the earlier results we have a simplified problem

$$rV(d, \underline{d}) = \max_{\substack{d_1, d_2, x \\ \underline{d}_1, \underline{d}_2, \underline{x}}} \left\{ \begin{array}{l} \lambda (R(d_1) + V(d_1, \underline{d}_2) - V(d, \underline{d})) + \\ \lambda (R(\underline{d}_1) + V(d_2, \underline{d}_1) - V(d, \underline{d})) + \end{array} \right\} \quad (5)$$

$$\begin{aligned} rd &= x + \lambda(d_1 - d) + \lambda(d_2 - d) \\ r\underline{d} &= \underline{x} + \lambda(\underline{d}_1 - \underline{d}) + \lambda(\underline{d}_2 - \underline{d}) \end{aligned} \quad (6)$$

We can now verify that the value function described by the simplified

problem is concave.

Lemma 2. *V is concave*

Proof. The Bellman equation can be rewritten as

$$V(d, \bar{d}) = \frac{1}{r} \frac{\lambda}{r + 2\lambda} \max(R(d_1) + R(\underline{d}_1) + V(d_1, \underline{d}_2) + V(\underline{d}_2, \underline{d}_1))$$

From this we can see immediately that the Bellman operator maps concave functions into concave functions, since the convex combination of choices for two states (d, \bar{d}) is feasible at the convex combination of the states, and delivers more when V on the right is concave. \square

Next, we make the final step in simplifying the problem. We argue that for any value of the state (d, \bar{d}) , it must be the case that $d + \bar{d} = 1/r$. Intuitively, if there were only one innovation, the planner would like to offer it $1/r$; as a result, given the many ideas that will arrive, the planner never “wastes” any instants.

Lemma 3. $d + \bar{d} = 1/r$

Proof. Suppose $d + \bar{d} < 1/r$. Since both d and \bar{d} cannot be greater than 1, it must be the case that $x + \bar{x} = 1$, since, if either duration is less than 1 the corresponding x should be increased. Since this applies at all instants, it must always be the case that $x + \bar{x} = 1$ and as a result all instants are promised to one of the two innovators, that is, $d + \bar{d} = 1/r$ \square

This further simplifies the problem: we can study it in terms of duration d , with $\bar{d} = 1/r - d$ at all dates. We are now ready to characterize the planner’s allocation by studying that problem.

4 Evolution of Complete Exclusion Rights

4.1 Dynamic Program

Since all time is allocated to one innovator or the other, the planner’s problem can be written as

$$rV(d) = \max_{d_1, d_2, x} \left\{ \begin{array}{l} \lambda (R(d_1) + V(d_1) - V(d)) + \\ \lambda (R(\frac{1}{r} - d_2) + V(d_2) - V(d)) \end{array} \right\}$$

s.t.

$$rd = x + \lambda(d_1 - d) + \lambda(d_2 - d)$$

Since the problem is symmetric, we generally focus our discussion on the shape of V in the set $[1/2r, 1/r]$. We first study when the promise keeping constraint binds, which gives some basic insight into the shape of V . This question is analogous to the question of when x is strictly between zero and one, since from the first order condition for x it is clear that x could not be interior unless the promise keeping constraint were not binding.

4.2 Characterization

Since V is globally concave and symmetric, it is maximized at $1/2r$. This is intuitive: when duration promise is identical to the two agents, you can treat the agents identically upon the next arrival, setting $d_1 = 1/r - d_2$, which is best since R is concave. Note that having the agents treated identically requires

$$\begin{aligned} rd &= x + \lambda(1/r - d_2 - d) + \lambda(d_2 - d) \\ x &= (r + 2\lambda)d - \lambda/r \end{aligned}$$

Therefore an identical result can be accomplished with x between zero and one if $d \in [1/r - \hat{d}, \hat{d}]$. Intuitively, in this case, the planner can deliver any asymmetric preference by using x , leaving the balance of the duration promise identical across agents when the next innovation arrives, and allowing $d_1 = 1/r - d_2$. As a result it is immediate that

Lemma 4. $V(d)$ is constant in the range of $[1/r - \hat{d}, \hat{d}]$

This range is the one where the promise keeping constraint does not bind. Clearly, outside of this range the planner can no longer have $d_1 = 1/r - d_2$, and therefore value must be lower, since concavity in R dictates losses when the next arrivals are treated differently. It is clear that it is never optimal to choose a point in the interior of the flat portion, since raising the current innovator's promise has no cost. The following lemma shows that the planner must go even further.⁹

Proposition 5. $d_1(\hat{d}) > \hat{d}$

For duration promises in excess of \hat{d} , the first order condition for d_1 and concavity of V shows that duration is an increasing sequence for any consecutive ideas by innovator 1. An increasing sequence on an interval must

⁹Proofs of results from this point forward are contained in the Appendix.

converge, and of course by the first order condition for d_1 it cannot converge to $d < 1/r$, where $R' > 0$. Therefore, sequences of arrivals by firm one get arbitrarily close to a promised duration of $1/r$:

Corollary 6. *For all $d < 1/r$ and $\pi < 1$ there exists a T such that duration is greater than d with at least probability π .*

The implication of this result is that the allocations eventually have near monopolization, in the sense that eventually the system evolves to a point where one firm is promised almost the entire future. Duration rises and falls with arrivals by the two firms; the two firms engage in a "tug of war" for duration.

An interesting feature is the evolution of x . Starting from duration in the middle region where V is maximized, ideas by innovator 1 move duration up, and the promise keeping constraint binds. As a result, $x = 1$. Note that the intervening period between innovations is never split; for any duration $d > \hat{d}$, a sequence of innovations by the "trailing" innovator promised $1/r - d$ falls with every innovation by the trailing firm, but if d is high enough, innovations by the trailing firm may at first do not change x , so long as the duration promise falls but remains above \hat{d} . The trailing firm only receives rights when a sufficient number of innovations by it have moved duration below $1/r - \hat{d}$, i.e. past the decreasing portion of the value function. This conforms to the idea that trailing firms need to make sufficient progress before their innovations are deemed to "not infringe" on the current leader's patent. Here, during the period of infringement, the leader maintains rights to all innovations, including the ones being invented by the laggard firm, as if it has the ability to costly license the infringing ideas. The payoff to the trailing firm is the eventual ability to sell a product that embodies the entire history of ideas, once their duration promise is sufficiently high.

An alternative interpretation of the optimal allocation is not as patent policy alone, but as favorable treatment from a regulator more generally. Suppose favorable treatment allows the firm to reap all the benefits of innovations from any firm, for instance by the incumbent firm negotiating licensing contracts that extract full surplus. Here the optimal policy uses such favorable treatment as an incentive device. The regulator favors the leader until the laggard has had sufficient innovations to be the new leader, at which time the regulator shifts its favorable treatment to that firm. Laggard innovators innovate for eventual favorable treatment. In this sense the optimal allocation has a strong sense that it entails competition for the market as an

incentive device; this competition for the market, however, eventually leads to near permanent monopolization by some firm.

The environment we study introduces a natural desire by the planner to backload rewards, which is what leads to the monopolization. If the planner waits to provide a given firm preference, it can provide that firm preference for more innovations, since innovations are constantly arriving, and a given firm can be allocated preference for multiple innovations at a point in time. As a result, the planner waits to award the preference until the firm has a preponderance of the duration promise. Backloading of rewards is similar to the quantitative result in [1]. They stress the usual backloading motive which they term "trickle down incentives:" rewards that come when firms succeed repeatedly are useful both after several successes (when the backloaded reward arises) and earlier, when firms attempt to reach the stage where backloaded rewards arise. This is the usual backloading of incentives intuition from the literature on dynamic contracts. Our model generates backloading for a different reason, that rights can be granted to one innovator for two innovations, but not to two innovators for one innovation each, at the same point in time. This is fundamentally related to the idea that the innovations compete with one another, which is at the heart of the idea that rewards through rights are in conflict in such cases. Although our model does not nest the one used by [1], it is similar enough that the same force is likely at work in their numerical results.

An alternate interpretation is as an ex ante licensing contract signed between the two firms. An interesting feature is that the optimal contract at no point in time shares the profits generated by the joint research, in the sense of splitting the profits between the firms; the licensing always takes the form of dynamic splitting, where one firm is rewarded for their work by having a longer time during which they profit completely. The contract relies on being able to pre-specify the extreme rights (x being either zero or one) as a function of history that boils down to the state d . The policy is a sort of duopoly patent pool, where the firms pool their patents and profits flow to one of the pool members based on their relative contribution, measured through d .

So far we have studied exclusive rights, which are different from the ones in many papers including [1]. One might be concerned that complete exclusion rights are overly broad, in a way that might naturally lead to excessive monopoly power. Moreover, they require extreme shifts of rights that may be hard to implement. In the next section we consider a restricted class of

policies where complete exclusion is not available to the planner, so that the policies look more like ones studied elsewhere, and show that many of the same results are preserved. Moreover, we show that the resulting allocations can be easily decentralized through payments to the planner.

5 Incomplete Exclusion Rights

5.1 Definition

The previous section involved a policy where the planner kept promises of complete exclusion to innovators, sometimes only delivering that promise at dates far in the future. Further, it required that the planner be able to “turn off” rights later on: when the innovator received rights, it was absolute, including for innovations previously receiving rights for the competition. The planner delivers an airtight promise not only that the leading firm will be an exclusive producer of not only the leading edge product, but also that it will not face even the threat of competition from any product previously marketed. In the language of patents, the earlier innovator’s patent infringes on a patent that comes later. This is contrary to the usual notion of infringement on *prior* art. In this section we restrict the planner to allocate rights in a way that does not involve any of these features. In particular, we force the planner into two restrictions. First, if the planner is to give a particular innovator exclusive rights to produce a particular product, it must deliver those rights at the moment of the innovation’s arrival, and not later. This corresponds to a patent right that must be either granted or refused for the innovation. Further, the planner cannot ever exclude an innovator from producing something which it is at one point had exclusive rights to produce. In patent language, a patent can never infringe on a patent that comes after it. In the quality ladder model, this only requires that the leading firm is concerned about competition from prior patented (and marketed) products; laggard products still sell nothing. Competition from existing products is inevitable, as it is difficult to reduce rights over time.

As a result, the leading edge firm always faces competition from the last innovation of the laggard firm. If a new leader emerges, the laggard will never be able to profit from any prior idea. In the language of the dynamic program (5), the assumption maintained in this section is:

Assumption: (Incomplete Exclusion Rights) If $\underline{d}_1 > 0$, then $d_2 = 0$. If

$d_2 > 0$ then $d_1 = 0$.

This implies that, at any time, only one innovator has a promise.¹⁰ In other words, whenever a new idea is awarded some preference, all prior claims to market leadership by innovators other than the current one are set to zero, as here. Exclusivity is natural in the patent context, in that it mirrors the sort of market structures that are assumed by many dynamic models of patents such as [16]. It maps into a patent right that is non-infringing on past rights, but does not encompass the prior rights. That is, every patent that is granted is non-infringing on every other patent, both before and after the arrival of a given improvement. Firms have the exclusive right to produce quality levels they "invent," but no right to exclude previous products invented by others.

Note that this assumption effectively rules out any allocations with anything that looks like meaningful licensing. Given the moral hazard problem, rewards can only take the form of duration, so licensing would mean a share of profits; sales of patents are effectively ruled out because it would give firms an incentive to underinvest and sell. Since under the incomplete exclusion assumption duration is either-or across firms, allocations where one firm shares with the other are impossible.

In the complete exclusion case the optimal policy was a very sophisticated dynamic contract, where laggards innovate because of the promise of future ownership of all innovations after sufficient success. Effectively there are always two rights holders; a leader who profits, and a laggard who hopes to profit in the future. In this section we show that, under incomplete exclusion rights, we get the same sort of backloading as in the last section, but in a way that requires a much less complicated history dependence. In particular we show that the allocation can be decentralized through a simple sale of rights.

In the exclusive case, the planner has less duration to allocate, since rights can be given to a smaller number of innovations (by assumption) at any instant: only the most recent string of innovation by a given firm can be granted rights, and not ones that came before the last innovation by the other firm. Therefore there is, in this environment, a natural interpretation of $d = 1/(r + \lambda)$ that is similar to \hat{d} : all future innovations are implemented, meaning that duration for the current incumbent is defined by "until the next

¹⁰In Hopenhayn et al. (2006), such a patent system is defined to be exclusive. We avoid that terminology to avoid confusion with the rights described in the last section, where at any point in time one innovator has exclusive rights to the entire ladder.

idea of the other firm arrives," which in discounted terms is $1/(r + \lambda)$. As a result, duration $d \leq 1/(r + \lambda)$ can be delivered without excluding anything, and there is scarcity in duration since $1/(r + \lambda) < 1/r$. We now turn to optimally choosing under the incomplete exclusion rights assumption.

Incomplete exclusion rights allows for very simple extension to more than two firms, which we discuss in more detail below. In short, the laggard firm can be taken to be a collection of outside firms. If one wishes to make their arrival probability different from the leader (for instance, proportional to their number, for instance twice the leader's arrival rate if there are three firms and therefore two laggards), nothing formally needs to be changed. We discuss the details in the extensions below; we maintain a single λ in this section for notational convenience.

5.2 Dynamic Program

Because only one innovator has a duration promise, we can write the planner's problem recursively as a function of that duration promise d . When that firm comes with another innovation, it gets a revised promise d_1 . When the outside firm has an idea, the planner must decide whether or not to implement it. Duration promises to the incumbent greater than $1/(r + \lambda)$ require some exclusion; we call the planner's current probability of implementing the outsider p . If implemented ($p > 0$) the new innovation is promised duration d_2 . In this case, by exclusivity, the innovator who entered the instant with promise d has their duration adjusted to zero, and we track the new duration promise d_2 .

The dynamic program is therefore

$$\begin{aligned}
 rV(d) &= \max_{d_1, d_2, p} \left\{ \begin{array}{l} \lambda(R(d_1) + V(d_1) - V(d)) + \\ \lambda p(R(d_2) + V(d_2) - V(d)) \end{array} \right\} & (7) \\
 & \text{s.t.} \\
 rd &= 1 + \lambda(d_1 - d) - \lambda pd
 \end{aligned}$$

where d , d_1 , and d_2 can be taken to lie in $[1/(r + \lambda), 1/r]$, since if $d < 1/(r + \lambda)$, the PK constraint does not bind, and you can set $x < 1$. Therefore the value function is independent of d in this range, and we can restrict attention to the domain where all durations are in the range $[1/(r + \lambda), 1/r]$.¹¹ From

¹¹It is easy to show that this is a self generating property of the value function, and

the promise keeping constraint we see the tradeoff between offering rewards d_1 to the current innovator, and implementing outsiders: the greater you reward the insider through a promise of d_1 for their next innovation, the fewer outsiders are implemented.

5.3 Characterization

First, we show that if the current promise involves any exclusion ($d > 1/(r + \lambda)$), then the current offer to the incumbent firm if he arrives with an idea, $d_1(d)$, either involves the optimal level of innovation for that idea in isolation, $d_1(d) = 1$, or involves as little current exclusion as is consistent with promise keeping, $p(d) = 1$.

Proposition 7. *Suppose $d > 1/(r + \lambda)$. Then either $d_1(d) = 1$, or $p(d) = 1$.*

The characterization Lemma (7) shows the sense in which duration is backloaded: it is backloaded maximally, eventually giving every innovation of the leader's duration forever, and never implementing an outsider's innovation. The proof makes clear the reason for the backloading, which differs from standard theories of backloaded incentives. Since the planner is committed to d , he is committed to a fixed amount of exclusions of the non-incumbent firm. When those exclusions occur is welfare neutral, in the sense that all of the exclusions cost the planner missing out on a new incumbent starting with d_2 . When the planner implements duration d by excluding arrivals of the outside firm later, he raises the duration promise for all intervening arrivals by the incumbent, though, which raises the incumbents level of innovation. Intuitively, for a given d , every implemented non-incumbent is "bad luck" for the incumbent. The planner resolves this bad luck as soon as possible by making every "unfortunate" (for the incumbent) idea of the non-incumbent end the incumbents duration early on. Whenever the incumbent gets an arrival, then, the "good luck" for the incumbent is large: he has avoided a state where all his competitors ideas are implemented, and therefore gets the maximal duration increase $d_1(d)$ that the planner could have offered and maintained the promise of d . This makes the incumbent respond to an arrival of an idea with the maximal effort.

Backloading takes the strong form of monopolization in finite time.

therefore must be true of $V(d)$ which is a fixed point of the Bellman operator.

Corollary 8. *Suppose $d_2 > 1/(r + \lambda)$. Then for any probability $\pi < 1$ there exists T such that d^* has been achieved in no more than T periods with probability π .*

If new market leaders are granted any promise of exclusion of their competitor, it is optimal to completely monopolize the industry (in the sense that one firm is promised $1/r$, and the other is never implemented) in finite time. The question, then, is whether d_2 exceeds $1/(r + \lambda)$, so that duration ever enters this region. We show next that the answer is always yes, and therefore we have a complete characterization of the dynamics of d : increasing duration with arrivals by the incumbent to $1/r$, periodically resetting to d_2 when an idea is implemented by the non-incumbent.

Lemma 9. $d_2 > 1/(r + \lambda)$.

The Lemma shows that new incumbents are promised some exclusions of their competitors. The earlier results show that these exclusions are maximally backloaded, which generates duration promises that climb, with positive probability, to $d = 1/r$.

The intuition behind offering some exclusions to new incumbents is related to backloading. If the planner were forced to grant exclusions that generate d_2 immediately (i.e. for ideas that arrive immediately after the change in incumbency), then exclusions would not be beneficial and d_2 would be exactly $1/(r + \lambda)$. The reason is concavity of R : immediate exclusions cost $R(d_2)$ in un-implemented projects, but generate $R'(d_2)$. The latter is always smaller when R is strictly concave, but would be identical if R were linear. However, backloading leaves the cost of exclusions the same, but increases their benefit: exclusions far in the future generate benefits for all of the incumbents ideas that arrive in sequence in the meantime. For linear R , this extra benefit of exclusions shows immediately that backloaded exclusions are beneficial; by continuity they must be beneficial for R that are nearly linear. The proof extends this logic to show that for any concave R , a small amount of exclusions, sufficiently backloaded, is beneficial to the planner.

5.4 Implementation: Selling Exclusivity

In this section we show that the optimal policies under the incomplete exclusion rights restriction can be implemented without complicated contracts, but rather with fees. To do so we focus on the special case without the single

n such that p is interior. In that case there are two levels of exclusion: none (where $p = 1$) and complete exclusion after N arrivals. There will in turn be two fees. First is a typical “patent fee,” paid by a laggard who wants to become a new leader, and entitling the holder to a non-infringing patent on their innovation that lasts forever, but offers no “forward” protection: it effectively ends when another innovator pays the same fee and takes the lead. Second is an “exclusive rights” fee, which can be paid at any time by the current leader (i.e. the most recent payer of the patent fee), and which changes rights in one fundamental way: it disallows the competition from ever being granted a patent in the future. In essence it generates infinite forward breadth, which forecloses the market, since licensing is effectively impossible given moral hazard and the assumption of incomplete exclusion.

The nature of the implementation is that, given the patent fee f and the exclusive rights fee t , the planner need not observe anything; the fees themselves screen arrivals of innovations and implements exclusivity after N consecutive innovations by an incumbent. The way in which screening of the former is obtained is standard: profits from taking over leadership just covers the patent fee f , and therefore anyone without an innovation does not find it worthwhile to claim to have an innovation when they do not. Screening on the number of consecutive innovations works because, the greater is the number of innovations that an innovator is profiting from, the more value the exclusivity has. Therefore the fee needs to be set sufficiently high that only an innovator with N arrivals will be willing to pay the fee.

The complication is that, if an innovator planned to foreclose the market early, they could also over-innovate for each arrival in anticipation. The proposition below shows that the fees can be set such that this strategy is never profitable.

Proposition 10. *There exists f and t that decentralizes the optimal allocation. That is, laggards pay f only upon receiving an idea and leaders pay t after exactly N arrivals of ideas.*

One can interpret t as the “price of exclusivity.” It is an additional fee paid to make the firm never face competition from additional innovations.

6 Extensions and Discussion

6.1 Different arrival rates or more than two firms

A straightforward change to the model is to allow the firms to have different arrival rates λ for the two firms. Further, since the relevant arrival rates in the incomplete exclusion case are the leader firm and the laggard firm, this allows the very natural assumption that the firm with the last innovation may have a higher arrival rate for the next innovation. Nothing qualitatively changes about the results; we still have maximal backloading. This also allows a natural interpretation as N firms with symmetric arrival rates, where the leader has arrival rate λ and the laggard(s) have arrival rate $(N - 1)\lambda$, in proportion to their relative numbers.

Handling more than two firms in the complete exclusion case is conceptually similar to the case considered above, but is more challenging because now there are N promises, which even if all time is allocated, leads to a state variable with $N - 1$ dimensions. The basic economic intuition, that an arrival raises that firm's duration promise at the expense of the others, and that firms innovate sometimes for rewards that come later, must remain.

6.2 Market Structures with Static Distortions or Laggard Profits

Our assumption that the promise of preferential treatment for a given innovation is sufficient for computing social benefit from that innovation is important, and has strong implications. It implies that the return to offering preferential treatment for a given innovation does not depend on the way the firm's other innovations are being treated. For instance, the firm's incentive to innovate is determined entirely by the duration promise for the given innovation, and not what the firm's promises of preferential treatment are for other innovations. This does not imply that the firm can only profit while it gets preference, it simply implies that all units of time must generate profits only based on the preferential treatment, and profits do not depend on future units are allocated across firms. In the same spirit, the assumption implies that any social costs from distortions generated by the promise are, again, independent of the promises made to the firm's other innovations. In other words, the impact of innovation on profitability and on social welfare is not a function of the promises made for past innovations, or on the fu-

ture promises that might be made for future innovations. This assumption is also essential for the recursive solution we study: without it, one could not compute the return, let alone the optimal policy, without knowing at any point in time the two firms' complete portfolio of promises, making the state variable potentially expand without bound as time progresses.

We can, however, modify our structure to model two important features that the benchmark model does not include: static distortions from monopoly, and the possibility that laggard innovations generate profits. We can interpret the shape of R as directly making statements about the product market where patent rights are granted. In the analysis $R'(1/r) = 0$, so that there were no static distortions, since the allocation of the entire future (the period that the innovation will be enjoyed) makes the agent's incentives perfectly aligned with the planners, and maximizes social surplus. Less duration means less than efficient innovation, which is where the tension arises: duration is scarce relative to the amount needed to induce efficient research effort. To get efficient innovation on one innovation, the planner would need to preclude future (valuable) innovations. One could proceed with alternative market structures, leading to alternative $R(d)$ functions. Let $d^* = \arg \max_{0 \leq d \leq 1/r} R(d)$. The value of d^* is analogous to the (discounted) optimal patent length in a static model like [3] or [15]; since the market is driven by only one innovation, that innovation is granted duration d^* in order to maximize the planner's value in the space of rewards by product market treatment.

If $d^* < \frac{r+\lambda}{r(r+2\lambda)}$, it is immediate that the planner can implement every innovation at the Arrow-Nordhaus duration d^* , since the planner can always provide less than what is delivered under the plan that delivers \hat{d} . So the model only has a dynamic tradeoff if $d^* > \hat{d}$, to ensure that the planner faces a scarcity in market time. A simple description of this assumption is that if the planner offers the current innovator preferential treatment until the next arrival of the competitor, the protection is still insufficient relative to the Arrow-Nordhaus patent.

In this case all the formal results can be restated; backloading takes the form, in the exclusive rights case, or rising duration to no higher than d^* (rather than one); similarly, in the incomplete exclusion case, the planner generates increasing duration for the leader up to a maximum of the static-optimum, d^* .

One can generalize the example further so that profits when the firm is not preferred are not zero, but just less than when the firm is preferred. This

might be due to services it provides for the leading edge provider, in order to make the innovations work efficiently. Such an environment would mean that

$$\Delta(d) = \arg \max_{\Delta} d\Delta + \gamma(1/r - d)\Delta - c(\Delta)$$

where $\gamma < 1$ reflects the idea that the loss of exclusivity lowers the ability of the firm to profit from the innovation. In a sense γ in this example is inversely related to the scarcity the planner faces; when $\gamma = 1$ the firm gets the entire future for any innovation, and therefore there is no scarcity. If $\gamma = 0$ the firm can only profit when it holds the promise it was granted at the time of innovation.

7 Conclusions

We have characterized the solution to the problem facing a planner who must allocate rights to production across two firms who can use those rights to make profits, and in turn are encouraged to innovate by the provision of the rights. It allows us to address the question of what distribution of rights arises from planner's solution, and in particular how much the market becomes "concentrated." The planner, because he can allocate rights to a single firm for multiple innovations at any point in time, backloads rewards, giving the firm with the preponderance of the future promises an exclusive right to all of the current profits. The optimal policy we study leads to monopoly, in the sense that one firm is excluded even though it is getting useful ideas. We show that these basic results hold even if the planner is forced to use a restricted set of policies where rights are always granted immediately for any innovation that is implemented. In that case, the optimal allocation can be decentralized through a simple set of patent fees: one for a patent with no forward breadth, and an additional fee that gives the innovator infinite future breadth. One can interpret the results as casting light on regulatory policies designed to foster competition "for the market." When the state dependence of rights is combined with a dynamic model of competition for the market, competition dies out in the long run.

Appendix

Proof of Proposition 5:

Proof. Since it is clear that $d_1(\hat{d})$ can never be less than \hat{d} , we focus on the case where $d_1(\hat{d}) = \hat{d}$. Since promise keeping does not bind, this implies that $d_2(\hat{d}) = 1/r - \hat{d}$. In that case, the system just oscillates between \hat{d} and $1/r - \hat{d}$; the planners payoff is

$$V(\hat{d}) = \frac{2\lambda}{r}R(\hat{d})$$

We show that in this case that V is differentiable at \hat{d} , implying that $V'(\hat{d}) = 0$ since V is flat to the left of \hat{d} , which means that the first order condition

$$R'(d_1) = -V'(d_1) + \mu(d)$$

cannot be satisfied if $d_1 = d = \hat{d}$, since the envelope condition would then imply

$$\begin{aligned} R'(d_1) &= -V'(d_1) + V'(d) \\ &= 0 \end{aligned}$$

To show that V is differentiable at \hat{d} , we describe a differentiable function \tilde{V} that is below V near \hat{d} . Since V is concave, the existence of such a function implies that V is differentiable.

To construct \tilde{V} , suppose the planner delivers duration away from \hat{d} by ε units by giving firm one extra duration at all future points when the other firm has the most recent innovation (and $x = 1$ when firm one has the most recent innovation). This implies that all innovations by firm 1 receive $\hat{d} + \frac{r}{\lambda}\varepsilon$, and all innovations by firm 2 receive $\hat{d} - \frac{r}{\lambda}\varepsilon$. Therefore the planner's payoff

$$\tilde{V}(\hat{d} + \varepsilon) = \frac{\lambda}{r}R(\hat{d} + \frac{r}{\lambda}\varepsilon) + \frac{\lambda}{r}R(\hat{d} - \frac{r}{\lambda}\varepsilon)$$

Under the maintained assumption that $V(\hat{d}) = \frac{2\lambda}{r}R(\hat{d})$, \tilde{V} is a differentiable function equal to V at \hat{d} . Since it is feasible choice for the planner, must be less than the payoff V from the optimal policy. But therefore V is differentiable, implying that $d_1(\hat{d})$ must exceed \hat{d} , and contradicting that $V(\hat{d}) = \frac{2\lambda}{r}R(\hat{d})$. \square

Proof of Proposition 7:

Proof. If $d = 1/r$, then $p = 1$ and $d_1 = 1/r$ are immediate from promise keeping. Therefore we focus on the case where $d < 1/r$. Suppose that $p < 1$ and $d_1(d) < 1/r$.

Denote by $d_1^t(d)$ the duration promise, starting from d , after t consecutive arrivals of the incumbent, starting from a promise of d . Then $d_1(d) = d_1^1(d)$. Moreover, let $p^t(d)$ be the probability of implementing the entrant after t consecutive arrivals by the incumbent; we have assumed, to a contradiction, that $p^0(d) = p < 1$.

Denote by τ the smallest positive integer such that $p^\tau(d) > 0$. Since $d < 1/r$, it must be the case that τ is finite. In words, τ is the number of consecutive arrivals by the incumbent before $p > 0$. Note that following the promise keeping constraint, duration is weakly falling with these arrivals, so $d_1^t(d) < 1/r$ for $t \leq \tau$.

We show that this policy cannot be optimal by considering the following variation. We increase $p(d)$ by ε . In order to maintain promise keeping, we must lower the implementation of entrants elsewhere. We do this at the node that follows τ consecutive arrivals by the incumbent. The initial policy dictates that the entrant, if it arrives, be implemented with probability $p^\tau(d)$. We lower this probability by $(r + \lambda)^\tau \varepsilon$; with that probability we do not implement the entrant, but instead keep the incumbent, but increment the incumbent's duration to the initial duration promise d from τ arrivals hence.

Since we are, in some states, offering the incumbent d (rather than nothing) after τ consecutive arrivals, we are of course increasing $d^\tau(d)$. This increases the duration promises for all $d^t(d)$ for $t > 1$ in turn. Since $d^t(d) < 1/r$ for $t > 1$, all of these are improvements to welfare.

We now claim that the change maintains initial promise keeping of d , and has no additional impact on welfare. That it maintains d is by construction: the exclusions after an entrants arrival following τ arrivals by the incumbent increase duration by

$$\frac{1}{(r + \lambda)^{\tau+1}} d (r + \lambda)^\tau \varepsilon$$

The first term is the discounting until τ arrivals by the incumbent, followed by one by the entrant; d is the gained duration in this state; and $(r + \lambda)^{\tau+1} \varepsilon$ is the probability that the duration is granted. This simplifies to $d(r + \lambda)^{-1} \varepsilon$, which is exactly the amount of duration that is lost when an additional ε

probability of losing d at the first node is lost, after arrival by the entrant.

To see that it has no additional impact on welfare, note that there are two other changes induced by the policy. First, with (discounted) probability $(r + \lambda)^{-1}\varepsilon$, we implement immediate arrivals by the entrant, and the planner moves from state d to state d_2 , earning $R(d_2)$ from the entrant. However, with the identical (discounted) probability, an entrant who would have been given d_2 is not implemented, and in this case the state transit to d instead of moving to d_2 . Since these happen with equal discounted probabilities, these changes exactly cancel, and therefore the modification is a strict improvement in welfare. \square

Proof of Lemma 9:

Proof. Consider a small exclusion, only if the entrant arrives immediately after $\tau - 1$ arrivals by the incumbent. As in the proof of, in this event of exclusion, set the incumbents duration back to $1/(r + \lambda)$, so it has no change in continuation utility (in other words, this is a one shot exclusion). This generates losses due to exclusions of

$$\left(\frac{1}{\lambda + r}\right)^\tau R(1/(r + \lambda))$$

It generates gains for the τ periods from the beginning to the $\tau - 1$ arrival of incumbent ideas. For instance, the $\tau - 1$ arrival by an incumbent has duration increased at rate

$$\frac{1}{\lambda + r} \frac{1}{\lambda + r}$$

since, if the next arrival is by the entrant (the first term) then there is an increase of $\frac{1}{\lambda+r}$ for every unit of exclusion. This generates gain

$$\left(\frac{1}{\lambda + r}\right)^{\tau-1} R'(1/(r + \lambda)) \left(\frac{1}{\lambda + r}\right)^2$$

where the first term is the discounting until the duration promise increases, the second term is the gain from the increased duration, times the rate at which duration is increasing.

It is easy to verify that all of the benefit terms simplify to the same $\left(\frac{1}{\lambda+r}\right)^{\tau+1} R'(1/(r + \lambda))$; earlier terms get less duration increase by a factor of

$\frac{1}{\lambda+r}$, but happen sooner by the same factor. Therefore the total gain is

$$\tau \left(\frac{1}{\lambda+r} \right)^{\tau+1} R'(1/(r+\lambda))$$

and therefore there is an improvement if

$$\begin{aligned} \tau \left(\frac{1}{\lambda+r} \right)^{\tau+1} R'(1/(r+\lambda)) &> \left(\frac{1}{\lambda+r} \right)^{\tau} R(1/(r+\lambda)) \\ \tau &> \frac{R(1/(r+\lambda))}{\frac{1}{\lambda+r} R'(1/(r+\lambda))} \end{aligned}$$

□

Proof of Proposition 10

Denote by W_N the value of an outside firm, upon receiving an idea and paying f , if he excludes after N arrivals. We set $f = W_N$. This implies that the value of being the laggard firm is zero. We need to show that, first, we can set t so that it is optimal to pay t after N arrivals. Then it is immediate that a laggard firm without an idea does not find it worthwhile to pay f , and one who does is indifferent between paying f and not.

Consider the deviation of paying t after n arrivals, after the payment of f . The value of this plan is denoted W_n for arbitrary n . For any such deviation strategy we have associated durations when n steps from foreclosure. They can be solved recursively from

$$rd_n = 1 - \lambda d_n + \lambda (d_{n-1} - d_n)$$

with $d_0 = \frac{1}{r}$. Denote $\beta = \lambda \int e^{-(r+2\lambda)t} dt = \frac{\lambda}{r+2\lambda}$. The recursion implies

$$d_n = \frac{(1 - \beta^n)}{1 - \beta} + \frac{\beta^n}{r}.$$

We can divide up rewards into profits and expected payment of fees.

$$W_n = v_n - \beta^n t$$

The profits from selling follow a simple recursion:

$$v_{n+1} = \pi_{n+1} + \beta v_n$$

where $\pi_n = \max_x d_n x - c(x)$. This recursion decomposes the reward from innovating into two parts. First, there is the expected profits from selling the current increment. Since that innovation lasts for d_n units of time, its profits are π_n . If the incumbent gets the next idea (embodied in the discounting by β), they will face the same problem as any innovator $n - 1$ steps from foreclosure, except for the profits they make from the increment they generated from earlier innovations.

We want to show that an arbitrary n (in particular N) can be made the maximum of W_n by appropriate choice of t . We start by showing that we can make it a local maximum, i.e.

$$\begin{aligned} v_n - \beta^n t &\geq v_{n+1} - \beta^{n+1} t \\ v_n - \beta^n t &\geq v_{n-1} - \beta^{n-1} t \end{aligned}$$

Note that the first is equivalent to

$$\beta^n(1 - \beta)t \leq (1 - \beta)v_n - \pi_{n+1}$$

and the second is equivalent to

$$\beta^n(1 - \beta)t \geq \beta((1 - \beta)v_{n-1} - \pi_n)$$

Note that $v_n > \beta v_{n-1}$ and $\pi_n > \pi_{n+1}$, so these two can always be satisfied simultaneously for appropriate t .

The next two claims verify that any local maximum is also a global one. This completes the proof.

Claim 11. If $v_n - \beta^n t \geq v_{n-1} - \beta^{n-1} t$ then $v_{n-1} - \beta^{n-1} t \geq v_{n-2} - \beta^{n-2} t$.

Proof. Rewrite the first inequality as $\pi_n + \beta v_{n-1} - \beta^n t \geq \pi_{n-1} + \beta v_{n-2} - \beta^{n-1} t$. Observing that $\pi_n < \pi_{n-1}$ it follows that $\beta v_{n-1} - \beta^n t \geq \beta v_{n-2} - \beta^{n-1} t$. Dividing through by β we get $v_{n-1} - \beta^{n-1} t \geq v_{n-2} - \beta^{n-2} t$. \square

Claim 12. $v_n - \beta^n t \geq v_{n+1} - \beta^{n+1} t$ implies $v_{n+1} - \beta^{n+1} t \geq v_{n+2} - \beta^{n+2} t$.

Proof. Multiply the first inequality by β and substituting on the left hand side βv_n by $v_{n+1} - \pi_{n+1}$ and on the right hand side βv_{n+1} by $v_{n+2} - \pi_{n+2}$ gives:

$$v_{n+1} - \pi_{n+1} - \beta^{n+1} t \geq v_{n+2} - \pi_{n+2} - \beta^{n+2} t.$$

Observing that $\pi_{n+1} > \pi_{n+2}$ this implies that $v_{n+1} - \beta^{n+1} t \geq v_{n+2} - \beta^{n+2} t$. \square

References

- [1] Daron Acemoglu and Ufuk Akcigit. State-dependent intellectual property rights policy. NBER Working Papers 12775, National Bureau of Economic Research, Inc, December 2006.
- [2] Howitt P Aghion P, Harris C and Vickers J. Competition , imitation and growth with step-by-step innovation. *Review of Economic Studies*, 68(3):467–492, 2001.
- [3] K.J. Arrow. Economic welfare and the allocation of resources for invention. In National Bureau of Economic Research conference report, editor, *The Rate and Direction of Inventive Activity*, pages 619–25. Princeton, 1962.
- [4] Gary S. Becker and George J. Stigler. Law enforcement, malfeasance, and compensation of enforcers. *The Journal of Legal Studies*, 3:1–18, 1974.
- [5] V.V. Chari, Mikhail Golosov, and Aleh Tsyviski. Prizes and patents: using market signals to provide incentives for innovations. Working Papers 673, Federal Reserve Bank of Minneapolis, 2009.
- [6] D.S. Evans and R. Schmalensee. Some economic aspects of antitrust analysis in dynamically competitive industries. In *Innovation Policy and the Economy, Vol. 2*. NBER and MIT Press, 2002.
- [7] Joshua S. Gans. Negotiating for the market. 2010.
- [8] Richard Gilbert and Carl Shapiro. Optimal patent length and breadth. *RAND Journal of Economics*, 21(1):106–112, Spring 1990.
- [9] Hugo Hopenhayn, Gerard Llobet, and Matthew F. Mitchell. Rewarding sequential innovators: Prizes, patents, and buyouts. *Journal of Political Economy*, 114(6):1041–1068, December 2006.
- [10] Hugo A Hopenhayn and Matthew F. Mitchell. Innovation variety and patent breadth. *RAND Journal of Economics*, 32(1):152–66, Spring 2001.
- [11] Michael Kremer. Patent buyouts: A mechanism for encouraging innovation. *Quarterly Journal of Economics*, 113(4):1137–1167, 1998.

- [12] E. Lazear. Agency, earnings profiles, productivity, and hours restrictions. *The American Economic Review*, 71:606–620, 1981.
- [13] John Stuart Mill. *Principles of Political Economy with some of their Applications to Social Philosophy*. Longmans, Green and Co., ed. 1909, London, 1848.
- [14] Matthew Mitchell and Yuzhe Zhang. Shared rights and technological progress. 2011, 2011.
- [15] William D. Nordhaus. *Invention, Growth, and Welfare: A Theoretical Treatment of Technological Change*, chapter 5. Cambridge, Massachusetts, 1969.
- [16] Ted O’Donoghue, Suzanne Scotchmer, and Jacques-François Thisse. Patent breadth, patent life, and the pace of technological progress. *Journal of Economics & Management Strategy*, 7(1):1–32, 03 1998.
- [17] Suzanne Scotchmer. On the optimality of the patent renewal system. *RAND Journal of Economics*, 30(2):181–196, Summer 1999.
- [18] Ilya Segal and Michael D. Whinston. Antitrust in innovative industries. *America*, 97:1703–1728, 2007.
- [19] E. Glen E. Weyl and Jean Tirole. Materialistic genius and market power: Uncovering the best innovations. Unpublished, 2010.