Rewarding Duopoly Innovators: The Price of Exclusivity*

Hugo Hopenhayn†and Matthew Mitchell‡

September 18, 2013

Abstract

We study an environment where a duopoly develops innovations that build on one another. The quality of innovations is private information, so optimal 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 result in near-permanent 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, but for different reasons than in the literature on backloaded incentives. We interpret our results in three ways. First, if one thinks of our allocations as offering policy guidance for intellectual property, then protection is state-dependent as in [1], generating heterogeneity in patent protection in the absence of heterogeneity in innovation opportunities. We show how this protection can be simply implemented by allowing firms to buy additional protection. Second, if the firms are able to contract ex ante, our allocations can be interpreted as a patent pooling arrangement between the firms, where one firm comes to dominate pool membership. Finally, one can interpret the optimal

*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.

evolution of the duopoly as competition "for the market," where static competition is not the key form of competition. Backloading implies that optimal competition for the market eventually leads to nearly permanent monopolization by one firm – competition for the market eventually extinguishes itself.

1 Introduction

This paper studies the optimal reward structure for a sequence of innovations generated by a small number 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 find that optimal 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 tradition of Mill and others. 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.

In our model, firms contribute sequentially to a quality ladder. In order to make this contribution a firm must have have an *idea* and make an *investment*. Ideas are random draws that arrive exogenously according to independent Poisson processes. Investment turns this idea into a step increase on the quality ladder. The size of this step is a function of both the quality of the idea and the level of investment. 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.

Following the literature on sequential innovation, we assume that only one firm can profit from sales of the good representing this ladder at any point in time, while potential competition of others might limit its price. A *patent mechanism* allocates these rights to sell as a function of the history, in an incentive compatible way. Effectively, the mechanism grants the selected leader the right to exclude from competition any product that lags the current one by less than a certain number of steps. As an example, in the case of *forward exclusion* rights -considered by most of the previous papers- this interval includes only past contributions of the current leader, since no firm is excluded from competing with its own previous product developments.

A leading firm will profit from a past contribution to the ladder only when that step is contained in the exclusion interval. The marginal value of investment for a given step looks forward to all future instances where this step will be included. The larger is this set, the higher the returns to investment and consequently the larger such step increase. This simple observation suggests two properties of an optimal patent mechanism that we prove in the paper. First, the exclusion interval should be *maximal*, because the larger the exclusion interval the more steps that are incentivized. In absence of any other restrictions this interval should include all steps, i.e. all other firms should be excluded from competing *in the market* even with their past product developments while the selected one captures all surplus. [1]

The second implication concerns a bias towards innovators that were more productive in the past, holding a lower bar for the quality of ideas to be implemented, a handicap in the competition *for the market;* identical opportunities are treated differently depending on the innovator's past history of contributions. While this unequal treatment is *ex-post* inefficient, it turns

[1] With the added restriction of forward exclusion, this set should include all past *continuous* contributions of the leading firm.

out to be *ex-ante* optimal as it incentivizes a larger set of prior innovation steps. In the extreme, this form of dominance will lead eventually to excluding forever from the patent race all but the luckiest innovator. Although such foreclosure is costly to the planner, the promise of foreclosure motivates innovation in the interim.

Our results are more stark when restricted to the case of forward exclusion -i.e. where firms are allowed to compete in the market with their own developed products- and there is no heterogeneity of ideas, which we consider first in the paper. In this setting, the optimal patent policy displays a bang-bang property: there is no exclusion from new ideas coming from other innovators (i.e. no forward protection) until the incumbent reaches a threshold of n consecutive innovations, at which point all outsiders are excluded permanently from the innovation race.

Counter-intuitively, total exclusion occurs even when the marginal benefit of preferential treatment for the remaining 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 exerts effort to increase the probability of reaching 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. In our setting, the probability of reaching the backloaded state are fixed, since they are tied to the exogenous arrival of ideas. The planner backloads rewards because a unit of time allocated after several ideas by the incumbent provides incentive to increase quality on all of those innovations that will contribute to cumulative quality. This complementarity between number of innovations and the reward that the planner provides leads to protection whose duration increases with additional successes.

We derive some additional implications for the *industrial organization* of innovation. Because the optimal patent mechanism exploits dynamic contracting opportunities, welfare increases when recurrence is more frequent. This occurs, for instance, when inventive ability is more concentrated: holding fixed the total arrival of ideas, welfare decreases with the number of innovators. Along the same lines, we show that welfare is higher when there is positive correlation in the arrival of ideas, where recent innovators are more likely to have follow on innovations than outside firms. Our results also suggest that in an environment where innovators differ in their *inventiveness* (measured by the rate of arrival of ideas), those that are slower will face a negative handicap and be held to higher standards for implementation.

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 types of innovations. In our setting, even if there is no heterogeneity built into the structure, the rewards are history dependent, leading to ex post heterogeneity in the reward for different innovations. This idea is related to 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. They 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 applies to their environment. It therefore helps develop further understanding of their numerical results. While they stress "trickle down incentives" in the spirit of [4] and [12], we show that backloading in the patent example can come from the fact that later rewards use scarce market time to reward more innovations from a given innovator, since innovators compete with one another but not with themselves.

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.

If instead of competing the firms can always jointly monopolize the product market, the planner can generate even better incentives. Under this structure, the optimal 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 arrangement, since the total surplus maximized by the planner is identical to the firms joint profits. The model can therefore be used to study the optimal dynamic patent pooling arrangement between the firms. 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. The results are qualitatively similar to the case where innovations compete; rewards are backloaded in the sense that one firm is nearly-permanently excluded in the long run. An alternative view of a policymakers job, in this case, 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."[2]

As a function of the promises the planner has made, the profits themselves evolve 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.

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

[2]In this sense our paper is related to the larger set of papers on regulation and innovation, for instance in [18] and [7]. [6] argue that competition for the market is as important as competition in the market.

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.

The paper is organized as follows. In Section 2 we describe preferences, the production technology, and the innovation technology. In Section 3 we study optimal allocations when the planner is constrained to assign competing patent rights, without licensing. This environment most closely resembles existing papers in the patent literature and demonstrates our basic results, as well as a simple decentralization of the optimal allocation. In Section 4 we consider the case where that restriction is lifted, and interpret the results as a dynamic patent pool or competition for the market. We consider the impact of static market distortion in Section 5, and then conclude.

2 Innovation and Competition

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[3] either takes an outside option (normalized to zero) or purchases one physical unit of the good of quality q that maximizes $q-p$, where p is the price paid for that variety.[4] 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

[3]Or unit mass of identical consumers.

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

selling the leading edge product, or as consumer surplus if $p < q$. The firm with the highest quality earns profits equal to the difference between the quality level of the highest and second highest quality level that is sold.

Our model abstracts from static monopoly costs. Focusing on the case with no static distortions 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. What is new here is the role that skewed rights between the firms impacts the firms incentives to innovate. In section 5 we describe how the model can be modified to incorporate static distortions.

2.2 Innovation

There is continuous time and an infinite horizon, with the future discounted at the rate r. We focus on the case where there are two agents (which we call firms or innovators) and a principal (or planner). Below we discuss extension to more firms. At any point in time a firm can be in either of two states: either it is endowed with a new *idea* or not. New ideas allow to produce a new quality level -an innovation- as described below. [5]

The innovation will be an improvement of size Δ upon the highest quality product currently available. Ideas can be turned (instantaneously) into innovations of size Δ in exchange for research cost $c(\Delta, \theta)$. The cost type θ is drawn from a continuous distribution $F(\theta)$. The stochastic process for the arrival of ideas and investment choices determine a path of arrivals $T_i = \{t_{i1}, t_{i2}, ...\}$ for player i and corresponding types $\Theta_i = \{\theta_{i1}, \theta_{i2}, ...\}$ and investment choices $\Gamma_i = \{\Delta_{i1}, \Delta_{i2}, ...\}$ .Denote by $T_i\,(t)$, $\Theta_i(t)$ and $\Gamma_i\,(t)$ the restrictions of these paths to arrivals prior to (and including) time t. Let $T\,(t) = T_1\,(t) \cup T_2\,(t)$, $\Theta(t) = \Theta_1(t) \cup \Theta_2\,(t)$ and $\Gamma\,(t) = \Gamma_1\,(t) \cup \Gamma_2\,(t)$. Correspondingly, at time t the sequence of innovations determines a frontier product $q\,(t) = \sum_{\Gamma(t)} \Delta_{in}$. Given that all consumers are identical, efficiency would require that this be the only product sold at time t.

[5] Although 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.

2.2.1 Patent policy

A *patent policy* $\{\bar{\tau}_1(t), \bar{\tau}_2(t)\} \epsilon T(t) \times T(t)$ prescribes at time t the latest innovation each player is allowed to use. This patent policy can be identified also with exclusion rights comprising all those innovations that the leader (the firm entitled to the frontier product) can exclude the other firm from using. In the above case, the exclusion set is $\xi(t, \bar{\tau}_1, \bar{\tau}_2) = \{\tau \epsilon T(t) | \bar{\tau}_2 < \tau \leq \bar{\tau}_1\}$. Profits are determined by Bertrand competition as above. For example, assuming $\bar{\tau}_1 > \tau_2$ the price and profit flow obtained by the first innovator is $q(\bar{\tau}_1) - q(\bar{\tau}_2) = \sum_{\xi(t,\bar{\tau}_1,\bar{\tau}_2)} \Delta_\tau$. A *patent mechanism* prescribes a patent policy at time t as a function of the history of arrivals, i.e. a function from $T(t) \times \Theta(t)$ to the class of patent policies described above. The restriction to policies that are functions of the arrivals and not the actual innovations (the set $\Gamma(t)$) is common to most of the literature on sequential innovation and is the fundamental source of moral hazard. The patent mechanism implies that (1) only one firm can earn profits at any instant; (2) profits are granted through an exclusive right to exploit a subset of the given innovations and (3) these rights can only depend on the sequence of ideas but not the size of the innovations. Our implementation discussed in Section 3.3.2 shows that the planner need not observe the arrival of ideas.

2.2.2 Incentives for innovation

We now derive the incentives for innovation from a given patent mechanism. Take for example an idea obtained by player one at time τ_i. Player one will be able to profit from this innovation every period where $\bar{\tau}_2(t) < \tau_i \leq \bar{\tau}_1(t)$ In every such period, player one will receive a flow of payments Δ_i for this particular innovation. Letting

$$d_i = E \int_{\bar{\tau}_2(t) < \tau_i \leq \bar{\tau}_1(t)} e^{-r(t-\tau_i)} dt$$

this gives player one an expected discounted value from this innovation equal to $\Delta_i d_i$. Because of the strong separability of the innovations' contributions to quality, this is all the relevant information a firm needs to know in order to choose optimally its investment to transform ideas into innovations.

We now derive the incentives for innovation from a given patent mechanism. Take for example an idea obtained by player one at time τ_i. For any product $q \epsilon Q = \{q(t) | t \epsilon T\}$ we will say that $\tau_i \succ (resp. \prec) q$ if innovation i comes after (resp. before) this product, *i.e.* $\tau_i < (resp. >) t^{-1}(q)$.

Player one will be able to profit from this innovation every period where $q_2(t) \prec \tau_i \preceq q_1(t)$. In every such period, player one will receive a flow of payments Δ_i for this particular innovation. Letting

$$d_i = E \int_{q_2(t) \prec \tau_i \preceq q_2} e^{-r(t-\tau_i)} dt$$

this gives player one an expected discounted value from this innovation equal to $\Delta_i d_i$. Because of the strong separability of the innovations' contributions to quality, this is all the relevant information a firm needs to know in order to choose optimally its investment to transform ideas into innovations.

When innovator i chooses Δ for innovation m, he solves

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

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. We assume that higher θ corresponds to a greater impact of duration on innovation: $\Delta_{12} > 0$. This assumption corresponds to a third derivative condition on c, which is satisfied, for instance, $c(\Delta, \theta) = \Delta^{\alpha}/\theta + k$ with $\alpha > 1$.

A particular duration d for a given innovation can be granted in many different ways. For example,, 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 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 the allocation of rights at every instant, this duration can evolve deterministically over time, change with later arrivals of innovators, and with the identity and type of those innovators that arrives with an idea. Because discounted time is not unbounded, the maximum possible promise of sure preferential treatment forever is $1/r$. This bound highlight the fact that rights are a limited resource that must be allocated in the most efficient way to provide incentives for innovation by firms.

The strong separability of innovations highlighted before carries over to the calculation of social value. Since any innovation Δ contributes this increment from it inception into the infinite future, it yields $\Delta/r - c(\Delta, \theta)$ additional units of present discounted social surplus.The social contribution of allocating d units of duration to an idea θ is then $R(d, \theta) = \Delta(d, \theta)/r - c(\Delta(d, \theta), \theta)$ Note that the derivative with respect to duration is

$$\begin{aligned} R_1(d, \theta) &= \Delta_1(d, \theta)/r - c_1(\Delta(d, \theta), \theta)\Delta_1(d, \theta) \\ &= \Delta_1(d, \theta)(1/r - d) \geq 0 \end{aligned}$$

where the second line uses the fact that $c_1(\Delta, \theta) = d$ by the innovator's FOC. Since convexity of costs implies that $\Delta_1(d, \theta) > 0$, the return is strictly increasing in duration except where $d = 1/r$. Our model has no static distortions, so that $R(d)$ is maximized at $1/r$.[6] Under the sorting condition $\Delta_{12} > 0$, the marginal return to duration increases in type, $R_{12} > 0$.

The planner trades off allocation of duration to the current innovator, which generates value via R, against the cost of that promised duration in limiting future duration left to allocate to the other innovator. The stock of available rewards available to a particular innovator is curtailed by the amount that is promised to the competition. We now turn to studying the optimal allocation of duration across histories.

3 Forward Exclusion Rights

3.1 Policy Space

In this section we consider a restricted set of contingent rights that is closest in spirit to a typical view of a patent in many dynamic models of patents such as [16]. In particular, the planner offers innovators exclusive rights to market their own innovation, and potentially excludes some (small) improvements to the patented work. We call this protection *forward exclusion rights*, and it is equivalent to forward patent breadth in the language of [16].

Let $q_i(t)$ denote the product that this policy entitles innovator i to sell. (More formally, letting $\tau_i(t)$ denote the last time player i innovated before t, $q_i(t) = q(\tau_i(t))$.) Assuming for the sake of argument that player 1 is the

[6] Section 5 generalizes to the possibility that $R(d)$ is not maximized at $1/r$, which can be interpreted as static costs of monopoly.

one with the last innovation, $q_2(t) < q_1(t) = q(t)$, Bertrand competition will entitle players profit flows $\pi_1(t) = q_1(t) - q_2(t)$ and $\pi_2(t) = 0$. It is consistent a patent right that is non-infringing on past rights, but does not encompass the prior rights. Firms have the exclusive right to produce quality levels they "invent," but no right to exclude previous products invented by others. In consequence, the most recent patent holder makes profits equal to the gap between his quality and the quality level of the most recent innovation by his rival.[7]

It is instructive, before considering the recursive structure of the problem, to think about the limits of the planner's powers to offer exclusion rights. On the one hand, the planner could offer "full exclusion" to the most recent innovator, excluding all of the competition's future ideas. Such duration would imply $d = 1/r$ for all of the incumbents innovations forever, but get nothing from the other firm. On the other hand, if the planner promised no forward exclusion of future innovations, each innovation would receive duration $d = 1/(r+\lambda)$, i.e. the discounted time until the next arrival by the other innovator. Duration $d \leq 1/(r+\lambda)$ can be delivered without excluding anything, and there is scarcity in duration since $1/(r+\lambda) < 1/r$.

3.2 Dynamic Program

We now turn to the optimal choice of duration. We solve the problem recursively. The planner has an outstanding promise of d as a result of commitments made previously to the current incumbent. If this incumbent receives a new arrival θ ,the planner grants an updated duration $d_1(\theta)$. Alternatively, if the non-incumbent receives an arrival θ, the planner can either choose not to implement the innovation (offering it duration zero, and maintaining the current incumbent's position), or to implement it by offering some $d_0(\theta)$, making the outsider the new incumbent.[8] It is immediate that, among excluded ideas, the planner always chooses to exclude the worst ideas; therefore the planner has some cutoff rule $\bar{\theta}$ below which ideas of the non-incumbent

[7] In Hopenhayn et al. (2006), a patent system with a single rights holder at any instant is defined to be exclusive. We avoid that terminology to avoid confusion with the rights described in the next section, where at any point in time one innovator has exclusive rights to the entire ladder.

[8] In the appendix we verify that the planner neither wants to adjust the incumbents promise when an arrival occurs that is not good enough to be implemented, nor at instants when nothing arrives.

are excluded. The planner's problem is

$$
\begin{aligned}
rV(d) &= \max_{d_1(\theta),d_2(\theta),\bar\theta} \left\{ \begin{array}{c} \lambda\left(\int(R(d_1(\theta),\theta)+V(d_1(\theta))-V(d)\right)dF(\theta)+ \\ \lambda\int_{\bar\theta}^{\infty}\left(R(d_2(\theta),\theta)+V(d_2(\theta))-V(d)\right)dF(\theta) \end{array} \right\} \qquad (1) \\
& \quad s.t. \\
rd &= 1+\lambda\left(\int d_1(\theta)dF(\theta)-d\right)-\lambda(1-F(\bar\theta))d
\end{aligned}
$$

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 planner always could have offered the prior innovator more time without increasing exclusions, and therefore increase welfare.

This most recent incumbent has *forward* exclusions that guarantee a duration d for his state-of-the-art product. Future arrivals of the other firm may be excluded in order to deliver d.[9] The constraint in (1), which we call promise keeping (PK), guarantees that the planner actually does deliver d and is critical to understanding the problem. Given a current promise d, the planner delivers intervening instants until the next arrival. If the next arrival is by the incumbent, his duration promise increments to d_1; if the next arrival is by the other firm, the incumbent is supplanted with probability $1 - F(\bar\theta)$, i.e. if the arrival is good enough to not be excluded. The constraint simply says that, in expectation, the duration offered delivers the promise of d. This constraint is what allows us to solve the fully history dependent problem of allocating rewards in a tractable, recursive way.

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 formal difference of this problem from ones with innovators who never recur. From the promise keeping constraint we see the tradeoff between offering rewards d_1 to the current innovator, and implementing outsiders: the greater you

[9]For now, the interpretation is that no further contracting, such as licensing, is possible. In the next section we allow for the possibility that the innovators pool their patented work and share the returns according to a prespecified rule that can be interpreted as a rich ex ante licensing agreement. None of the key features of the model or the results differ.

reward the insider through a promise of d_1 for their next innovation, the fewer outsiders are implemented.

We characterize the solution to the dynamic programming problem using first order conditions. First we verify that the value function is concave and differentiable.[10]

Lemma 1. *$V(d)$ is concave and differentiable on* $[1/(r+\lambda), 1/r]$.

The first order condition for $d_1(\theta)$ is

$$R_1(d_1, \theta) + V'(d_1) = \mu$$

where μ is the Lagrange multiplier on the PK constraint. The envelope condition for the value function implies that

$$\mu = V'(d)$$

Substituting for μ in the first order condition immediately implies that the policy function $d_1(\theta, d)$ (i.e. the choice of duration as a function of the type and current state) has the following properties.

Lemma 2. *$d_1(\theta, d)$ is weakly increasing in θ and d. Moreover, $d_1(\theta, d) > d$.*

Proof. Immediate from $R_{12} > 0$ and the fact that $R_1(d_1, \theta) + V'(d_1) = V'(d)$. □

Better arrivals are granted more duration and new promised duration is increasing in the outstanding promise. It also follows immediately that the planner never excludes innovations by the incumbent, since there is no trade-off between rewarding new innovations by the incumbent and keeping the promise on old innovations. In this sense the incumbent is immediately favored relative to the outsider, who faces some exclusion. The second part of the Lemma implies that duration rises with each new arrival of the incumbent, so with a series of consecutive arrivals d cannot converge to a level less than $1/r$. Therefore 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. A firm with many consecutive successes will come to dominate the market, nearly forever. Correspondingly, almost all ideas of the other firm will not be developed.

[10]When omitted in the text, proofs are contained in the appendices.

In the next section we show that, unless heterogeneity is sufficient, the monopolization of the market happens in finite time. We accomplish this by studying the case where there is no heterogeneity in θ so all innovations have the same cost function. It further shows the sense in which the model generates heterogeneity in innovation size endogenously; all ideas are identical, but differential treatment leads to heterogeneous improvements. We develop the result in a way that explains further the mechanism behind the backloading in the model.

3.3 Identical Ideas

3.3.1 Backloading

In this section we consider a simplified version of the prior section where all ideas are identical, and we therefore suppress the θ variable. Because we no longer have the ability to convexify across types, we require that $R(d)$ is concave.[11] The absence of heterogeneity also implies that in order to deliver on the promise d, the planner may need to mix between implementing and not implementing a new idea from the non-incumbent. Denoting the probability of implementation of an idea by the non-incumbent firm by p the dynamic program is

$$\begin{aligned} rV(d) &= \max_{d_1,d_2,p} \left\{ \begin{array}{c} \lambda\left(R(d_1)+V(d_1)-V(d)\right)+ \\ \lambda p\left(R(d_2)+V(d_2)-V(d)\right) \end{array} \right\} \qquad (2) \\ & \quad s.t. \\ rd &= 1+\lambda(d_1-d)-\lambda pd \end{aligned}$$

Lemma 3. *There exists $\widehat{d} \le \frac{1}{r}$ such that $p=1$ if $d \le \widehat{d}$, is strictly decreasing and positive when $\widehat{d} < d < \frac{1}{r}$ and equal to zero when $d=\frac{1}{r}$.*

---

[11] 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.

Proof. First note that the choice of d_2 is independent of d and satisfies $R'(d_2) + V'(d_2) = 0$. As for the choice of p it is either interior, in which case $R(d_2) + V(d_2) - (V(d) - \mu d) = 0$ (where μ is the multiplier of the promiss-keeping constraint) or equal to zero or one. By the envelope condition $V'(d) = \mu$, so the term in brackets equals $V(d) - V'(d)\, d$. This term is increasing in d and can only be constant in a region where $V'(d)$ is constant. From the first order conditions for d_1 it follows that when $\mu\,(= V'(d))$ is constant, d_1 is also constant and hence p must decreas with d. Also note that the fact that p is decreasing implies that when zero, it must be the case that $d = \frac{1}{r}$. Hence there exist $\widehat{d}$ such that $p = 1$ if $d \leq \widehat{d}$, is strictly decreasing and positive when $\widehat{d} < d < \frac{1}{r}$ and equal to zero when $d = \frac{1}{r}$. □

Proposition 4. *If* $d_1(d) < 1/r$ *then* $p(d) = 1$ *and if* $p(d) < 1$ *then* $d_1(d) = \frac{1}{r}$.

This proposition provides a bang-bang result: either there is **no** exclusion for the outsider or there is **maximal exclusion** (upon one more arrival of the incumbent.) The intution for this result can be given with aid of Figure 1. Consider an incumbent with promise d. In the figure, $p_0 = p(d)$ and $p_1 = p(d_1(d))$. If the outsider has an innovation, with probability p_0 it is implemented and the duration promise goes to zero. If the incumbent gets an idea before the next time an outsider is implemented, the duration promise updates to $d_1(d)$ and the planner revises the probability of an outsider being implemented to p_1. The intuitition behind this result is as follows. Suppose towards a contradiction $p_1 > 0$ and $p_0 < 1$. From Lemma 3 it follows immediately that $d_1 < \frac{1}{r}$. Now consider a variation where p_0 and d_1 are increased slightly maintaining the original value d. This has no impact on the initial incentives for investment as d is unchanged while it increases the incentives to innovate upon the next arrival as d_1 is increased. It is also neutral on the expected discounted time of implementation of the outsider as d is unchanged, leading to the same intital expected exclusion cost as the original plan and thus increasing total value. In conjunction with the previous Lemma, this argument shows that there can be at most one period where $0 < p < 1$ and the Proposition easily follows. The proof makes clear the reason for the backloading. 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

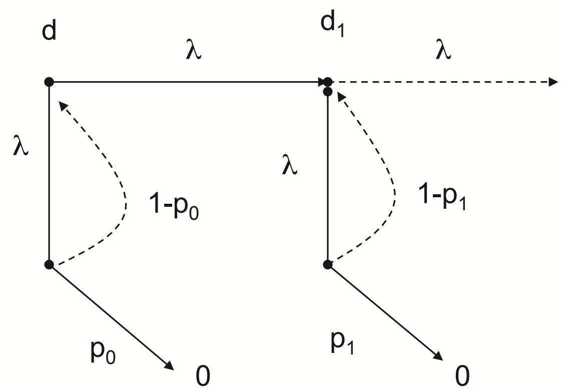

Figure 1: Optimal exclusions

the duration promise for all intervening arrivals by the incumbent, though, which raises the incumbents level of innovation. In other words, a scarce unit of time allocated later rewards both early and late innovations.

.

The above result can be also used to easily construct the path for duration starting from any initial value d. Note that for any given p, the promise keeping constraint gives a linear relationship between d_1 and d. Figure 2 shows this relationship for $p = 0$ and $p = 1$. First note that for $p = 1$, $\underline{d} = \frac{1}{\lambda+r}$ is a fixed point. This corroborates that this natural rate of duration is granted without any exclusions. For any higher initial value of d, the sequence of durations generated upon additional arrivals is strictly increasing. Note also that any durations to the right of $\widehat{d}$ require $p < 1$. The dynamics for optimal duration is thus the increasing sequence generated by the upper line until a value $\widehat{d} \le d \le \frac{1}{r}$ is reached, where p is chosen so that next period duration is $\frac{1}{r}$.

This dynamics assumes the initial level of d is above the natural rate $\frac{1}{\lambda+r}$, which is proved in Lemma 5.

Lemma 5. $d_2 > 1/(r + \lambda)$.

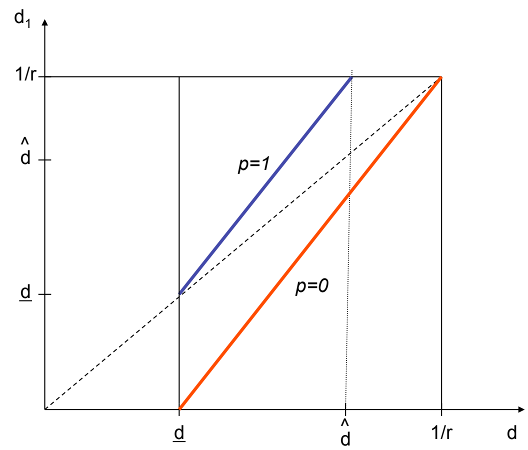

Figure 2: Evolution of duration

The Lemma shows that new incumbents are promised some future exclusion of their competitor. 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 restricted to set $d_1(d) = d$, so that backloading is disallowed, the planner's first order condition is just satisfied at $d_2 = 1/(r+\lambda)$: since duration is smoothed equally across the incumbent's arrivals when $d_1(d) = d$, the planner then smooths across incumbent and outsider by granting all innovations $1/(r+\lambda)$. Removing the restriction of $d_1(d) = d$ increases the marginal benefit of d_2, since that duration can generate more incumbent innovations by being backloaded, which generates a first order benefit to increasing d_2 to a level greater than $1/(r+\lambda)$.

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 lost innovations, 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.

The following Proposition summarizes the results developed above.

Proposition 6. ***(Optimal patent policy)*** *There exist an integer N and $p_{N-1}\epsilon\,[0,1]$ such that an arrival of the outsider is implemented for sure iff it occurs before $N-1$ consecutive arrivals of the incumbent; it is implemented with probability zero if it comes after N consecutive arrivals of the incumbent and with probability p_{N-1} if it comes after $N-1$ consecutive arrivals.*

Backloading of rewards is similar to the quantitative result in [1].[12] 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. The driving force is not the promise of a large prize after N successes, as arrivals are exogenous. In our model it is the cummulative nature of innovations that plays a central role. To clarify the difference, consider a restricted set of patent mechanisms that allows the outsider to use all but the last step of the incumbent, so that after any sequence of consecutive innovations, the incumbent will derive profits only from the last step. With this restriction, the optimal policy will have no backloading and the outsider's ideas will never be excluded. For backloading to arise, it is central that exclusions that occur after several arrivals reward the incumbent for all interim innovation steps.

As shown in the Appendix the derivative of the value function has a simple expression:

$$V'\left(d\right)=\sum_{n=1}^{N}R'\left(d_{n}\right)-r\left[R\left(d_{2}\right)+V\left(d_{2}\right)-V\left(\frac{r}{r}\right)\right], \tag{3}$$

where N corresponds to the exclusion period and $d_n$is the sequence generated using the promiss-keeping constraint starting from d. The two components in this formula highlight the benefits and costs of higher initial duration.

[12] 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.

The benefit is that higher current duration allows to provide more duration for future arrivals and consequently additional value. On the other hand, higher duration increases the probability of exclusion and this carries a cost as measured by the term in brackets.

Equation (3) can be used to show the value of exclusion. In absence of exclusions, at the natural rate $\underline{d}$ where N is infinite, the formula shows that $V'(\underline{d}) = \infty$, as all terms in the series are identical and positive. Finally note that strict concavity of the value function follows immediately from equation (3). [13]

3.3.2 Implementation

Next we turn to the question of how this allocation could be implemented by the planner when arrivals are private to the firms. Our implementation also requires minimal information on the side of the firms. To simplify our exposition we focus on the special case where $p_{N-1} = 0$ so there is no exclusion until the N^{th} arrival from the incumbent and full exclusion afterwards. This simplifies the necessary tools to implement the optimal allocation, and shows the general principal behind the implementation. The planner employs two fees $\{f, t\}$. The first is a typical *patent fee* paid by an outsider who wants to become the new incumbent, 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. The second one 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 implementation can be interpreted as a screening device. The way in which screening of the outsider is obtained is standard: the value of 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. As for the incumbent, screening on the number of consecutive

[13] In addition, it also follows that $V$is differentiable at all points including those where N changes, as the terms dropped or added have $d_N = \frac{1}{r}$ and thus $R'(d_N) = 0$.

innovations works because the greater is the number of steps that an innovator is profiting from, the higher is the value of exclusivity. Therefore the fee needs to be set sufficiently high that only an innovator with at least N arrivals will be willing to pay the fee.

The complication is that, if an innovator planned to foreclose the market early, they could potentially benefit from a double-deviation, where they also invest a greater amount in anticipation of the foreclosure. The proposition below shows that the fees can be set such that this strategy is never profitable.

Proposition 7. *There exists f and t that decentralizes the optimal allocation. That is, an outsider pays f only upon receiving an idea and the incumbent pays t after exactly N arrivals of ideas.*

One can interpret t as the "price of exclusivity," a fee paid to make the firm never face competition from additional innovations. An alternative way to interpet the price of exclusivity is as a handicap rule in an instant-by-instant auction of the right to sell, or more precisely the right to exclude the other firm from using any of the last consecutive step increases. In the absence of such a handicap, an outsider would never have the incentive to innovate, since the profitability of a new innovation that competes with an old one is always the lower than for any other innovation, so the incumbent would always win the subsequent auction. Such an auction would succeed at maximizing joint profits for the two firms, by focusing all innovation at one firm, but at the expense of consumer surplus. While exclusivity is inevitable in the optimal allocation, the firms' joint interests would be served if exclusivity arrived immediately. The handicap or price of exclusivity trades off the incentive benefit of backloaded rewards that come with eventual exclusivity against the fact that, left to decide on their own, the firms would choose exclusivity too soon in order to limit competition.

3.3.3 Different arrival rates or more than two firms

In this section we consider a situation in which the incumbent has an arrival rate that is different from the outsider. The same logic and results apply to the case where there are more than 2 innovators, as the outsiders' total arrival rate is the sum over individual arrivals. To focus on heterogeneity and not on the effect of differing aggregate innovation opportunities, our analysis maintains the sum of arrivals constant and equal to λ, where $\alpha\lambda$ corresponds to the incumbent's recurrence rate and $(1-\alpha)\,\lambda$ to the sum of outsiders' arrivals. As an example, if the innovation ladder comprises M firms with equal arrival rates, then $\alpha = 1/M$.

The dynamic program is

$$\begin{aligned} rV(d;\alpha) &= \max_{d_1,d_2,p} \left\{ \begin{array}{c} \lambda\alpha\left(R(d_1)+V(d_1;\alpha)-V(d;\alpha)\right)+ \\ \lambda\left(1-\alpha\right)p\left(R(d_2)+V_M(d_2)-V_M(d)\right) \end{array} \right\} \\ & \quad s.t. \\ rd &= 1+\lambda\alpha(d_1-d)-\lambda\left(1-\alpha\right)pd \end{aligned} \tag{4}$$

This problem is very similar to the one described above; indeed, the characterization of section 3.3.1 applies directly, so that p is maximized until $d_1(d) = 1/r$. In particular, in the application to many firms, no competitors are excluded until all are, and a decentralization like the one provided above can still be used. We prove two facts regarding comparative statics in α. First:

Proposition 8. *$V(d;\alpha)$ is increasing in α.*

The intuition is straightforward and can be easily explained for the case $\alpha = \frac{1}{M}$, where concentration of ideas (lower number of firms M) benefits the planner by allowing better dynamic contracting opportunities. We describe the intuition more precisely after we characterize the comparative statics on d_2. To derive these we first show that the marginal cost of allocating duration decreases with α (increases with the number of firms M)

Proposition 9. *Take $d' > d$. $V(d';\alpha) - V(d;\alpha)$ is increasing in α.*

This is intuitive: the marginal cost of allocating duration is $\partial V(d;\alpha)/\partial d$, which decreases when the current promise holder is more likely to be the next innovator. This implies that duration d and recurrence rate α are complements and given that the optimal initial duration d_2 maximizes the sum $R(d_2) + V(d_2;\alpha)$, it follows immediately that d_2 increases in α.

Corollary 10. *d_2 is increasing in α.*

While this corollary states that itinital duration increases with α, this not necessarily means that there will be more exclusions. This follows from the fact that for a fixed exclusionary policy (N,p) initial duration increases with α as it becomes more likely and faster for the incumbent to reach the exclusion region. However, we can establish that exclusion must increase for very low levels of recurrence. As α approaches zero (M approaches infinity), there is no recurrence so every innovator is unique. In that case there are no gains from backloading and the optimal policy approaches one that offers every idea the natural rate $d = 1/(r + \lambda)$, with no exclusion of others' ideas. But for higher α (smaller M) we have proved that exclusions are optimal.

Exclusions arise as the planner can do better in two ways. First, when the same innovator has consecutive innovations the planner can allow the innovator to profit from both, so the allocation of duration to the second innovation does not compete away duration of the first. Moreover, there are benefits from backloading and those benefits are only achievable to the extent that innovators recur. The more firms are competing, the less is offered to a new incumbent, and therefore the less likely is monopolization by that incumbent. Monopolization is still inevitable, but slower since the cost of exclusion is increasing in M.

Tables 1 and 2 provide numerical results for a set of arrival rates $\lambda \epsilon \{0.1, 1, 12\}$ with expected time of arrivals of 10 years, 1 year and 1 month, respectively and incumbent recurrence rate $\alpha \epsilon \left\{0, \frac{1}{8}, \frac{1}{2}, \frac{9}{10}\right\}$. In addition to these parameters, the interest rate $r = 0.05$ consistent with a yearly time unit and the cost function quadratic. For value of λ, the first column in Table 1 gives our measure of duration, the second column $N.p$, the number of periods prior to exclusion of entrants plus the probability of exclusion thereafter and the third column the probability that entrants are excluded at the beginning of incumbency. For fixed λ, duration increases with α (as predicted by the theory), though it takes more arrivals to exclude the rival. The intuition behind this result is that when the rate of recurrence is high, the probability of being displaced is lower conferring the incumbent a higher natural barrier and protection. The combined effect is much higher probabilities of exclusion. An increase in λ reduces the natural rate $1/(r + \lambda(1 - \alpha)$ for every α. To mitigate this reduction, the number of periods to exclusion are decreased leading at the end to a higher total probability of exclusion.

Table 2 gives two measures of relative values for each cell: the ratio of value to the first best (e.g. one that would be achieved with perfect licensing) and the ratio of the value obtained with no exclusions (the natural rate) to our constrained optimal value. In all cases higher recurrence rates α reduce the gap between the constrained optimal and the first best. This confirms the intuition highlighted before, that recurrence in innovation enhances incentives for investment. We also find that the gap is widened when the frequency of innovation λ increases. Our results show that the gain from

Table 1: Duration and Exclusion

α	$\lambda = 0.1$			$\lambda = 1$			$\lambda = 12$		
	d	$N.p$	$P\%$	d	$N.p$	$P\%$	d	$N.p$	$P\%$
0	6.7			1.0			0.1		
$\frac{1}{8}$	8.0	1.17	5.6	3.2	1.02	11.2	2.4	1.01	11.7
$\frac{1}{2}$	11.6	1.73	28.9	8.6	1.27	39.5	8.2	1.23	40.7
$\frac{9}{10}$	17.0	4.73	60.7	14.7	3.25	70.9	14.4	3.14	71.8

Table 2: Values

α	$\lambda = 0.1$		$\lambda = 1$		$\lambda = 12$	
	V/opt	V_nat/V	V/opt	V_nat/V	V/opt	V_nat/V
0	0.56	1.0	0.09	1.0	0.08	1.0
$\frac{1}{8}$	0.62	0.99	0.21	0.52	0.13	0.07
$\frac{1}{2}$	0.79	0.96	0.55	0.35	0.50	0.03
$\frac{9}{10}$	0.98	0.99	0.92	0.61	0.90	0.09

exclusion -as opposed to just providing the natural rate of protection- are very moderate when the frequency of innovation λ is low but are very large for when it is high. Interestingly, this suggests that a patent system that eventually excludes competiting innovators is more desirable in areas with faster rates of innovation.

The numerical results also highlight the importance of repeated innovation. The first row can be interpreted as a situation when innovators don't recur, while the $\alpha = 1/2$ case can be identified with a symmetric duopoly of innovators. Since for both cases the unconstrained optimal is the same, the values in Table 2 can be used to compare their performance. The range goes from a 40% increase in value for $\lambda = 0.1$ to over a 5 fold increase for $\lambda = 1$ or $\lambda = 12$. This shows that the concentration of innovators is particularly important in areas where there is a fast rate of innovation.

4 Complete Exclusion Rights and Competition "For the Market"

In this section we relax the requirement of forward exclusion, giving the rights over the whole ladder to one firm at a time, which we call the leading firm. As argued earlier, on the class of patent mechanisms this maximizes ex-ante incentives for innovation as it rewards an innovator for *all* contributions at any time where that innovator is profiting, and not just the most recent sequence of consecutive innovations as is the case with forward exclusion. In the language of [16], this corresponds to a patent with infinite lagging breadth. Competition for the market is the motivation for innovation.

Given there are no static distortions, this patent mechanism maximizes total discounted welfare from innovations, as it maximizes the scarce market time delivered to the innovator for all of his innovations. As social value and the sum of private value for the two firms are identical, one can interpret this section also as a private contract, signed at time zero, between the duopolists to maximize their joint payoff. For conciseness we develop these results for the case where all ideas are identical.

4.1 Dynamic Program

In the appendix, we show that scarcity of time requires that all instants be allocated to one innovator or the other. As a result the planner's problem can be written as a function of the outstanding duration d promised to innovator 1. As the other innovator will be profiting from all innovations in the complementary time, its entitled duration is $1/r - d$. This provides a simple recursive representation with state variable d.

The outstanding duration d of firm one evolves in an optimal fashion as a function of contingencies. If innovator 1 receives the next innovation we denote the new promised duration by d_1; if innovator 2 is the first to innovate, the new duration for innovator 1 will be denoted by d_2. In the first case, innovator two receives durations $\frac{1}{r} - d_1$ and $\frac{1}{r} - d_2$, respectively. While $0 < d_1 < \frac{1}{r}$, all ideas will be implemented, notwithstanding the fact that incentives for investment will be low for one player when the promised duration is close to one of these extremes.

The planner must also decide whom to allocate rights to sell in the interim period before the next arrival; in the prior section it was always to the most recent innovator. Here $x \in [0,1]$ determines the share of instants until the next arrival where innovator 1 has exclusive rights. So when $x = 1$, innovator one is the seller while for $x = 0$ it is innovator two. The dynamic program that solves the optimal assignments is given by the following functional equation:

$$
\begin{aligned}
rV(d) &= \max_{d_1,d_2,x} \left\{ \begin{array}{c} \lambda\left((R(d_1) + V(d_1) - V(d)\right) + \\ \lambda\left(R(\frac{1}{r} - d_2) + V(d_2(\theta)) - V(d)\right) \end{array} \right\} \qquad (5)\\
& \quad s.t. \\
rd &= x + \lambda(d_1 - d) + \lambda(d_2 - d)
\end{aligned}
$$

Since the problem is symmetric, the value function must be symmetric around the midpoint $\frac{1}{2r}$,so we restrict our discussion on the shape of V to 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 examing the above problem it is clear that x could not be interior unless the promise keeping constraint were not binding.

As in the case of forward exclusion, we can identify a natural rate corresponding to the case where selling rights are assigned to the most recent innovator rights until the next arrival. This corresponds to the duration $d = 1/(r + \lambda)$ from the last section, when there is no exclusion and therefore all innovations are treated identically. Denote by $\hat{d}$ the duration this offers to the most recent arrival. By symmetry, when the other firm has an arrival duration 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 to reward the other innovator in the future. We study this trade off using the dynamic program in the remainder of this section.

4.2 Characterization

The appendix shows that V is once again concave; since it is symmetric, it is maximized at $1/2r$. This is intuitive: when duration promise is identical for the two agents, they will be treated 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

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

An identical allocation to future innovations is feasible if $d \in [1/r - \hat{d}, \hat{d}]$, by adjusting x between zero and one. 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 11. *$V(d)$ 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.[14] 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.[15]

Proposition 12. $d_1(d) > d$

Since $d_1(d)$ cannot depend on d for d in the range of $[1/r - \hat{d}, \hat{d}]$, the proposition implies that $d_1(d) > \hat{d}$ for d in that range. 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, just as in the prior case. An increasing sequence on an interval must 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 $1/r$. 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 during this period. Suppose that, due to past success by innovator 1, $d > \hat{d}$. Then $x = 1$: the intervening period until the next innovation is entirely allocated to firm 1. If d is high enough, it may be the case that $d_2 > \hat{d}$, meaning that even after an arrival

[14]It is easy to show that this is a self generating property of the value function, and therefore must be true of $V(d)$ which is a fixed point of the Bellman operator.

[15]Proofs of results from this point forward are contained in the Appendix.

by firm 2, firm one will still maintain rights to all innovations, including the new one that firm 2 has just produced. Firm 2 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 contract can equally be viewed 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.

To illustrate the potential gains and the degrees of assymmetry that can arise, we provide results of numerical computations. The benchmark values used are $\lambda = 1/2$, $r = 0.05$ and a quadratic cost function. Table 3 give the ratio of values to the unconstrained optimal achieved through forward

Table 3: Values: Total exclusion vs. Forward exclusion

2λ	Forward (V/opt)	Total (V/opt)	Ratio
0.1	0.79	0.90	1.13
1	0.55	0.80	1.46
12	0.50	0.76	1.51

exclusion and total exclusion. Since the first best value is the same in both scenarios, the ratio given in the third colum shows the relative advantage of total exclusion, ranging from about 10 to 50%. Gains are larger again in industries with higher pace of innovation, due mostly to the shortcomings of protection under forward exclusion.

Figure 3 gives the distribution of duration generated by a sequence of 500,0000 draws and iterative application of the policy function. The range of values should go from zero to 20=1/.05, corresponding to the absorbing extreme points of duration. In spite of the fact that these two points are absorbing, convergence is so slow and unlikely that for a finite sample of as large a size as the one we have generated, there appears to be a non-degenerate distribution. This distribution shows the prevalence of assymetry (meaning durations that exceed the midpoint 10), obviously in favor of both players. The average duration to the right side of the midpoint is slightly over 12 (where player one is assigned 50% more duration than player two) and symmetrically below 10 is approximately 8 (where player two is assigned 50% more duration than player one.) Durations of 15 are clearly in the support (here player one is assigned 3 times the duration of player two.). As a reference point, it is worth noting that the natural duration $\hat{d}$ (the one that provides *equal treatment* to all innovations) is roughly 10.5, considerably away from the area where the mass of the distribution is concentrated. *Unequal treatment* seems more the norm resulting from the optimal incentive mechanism.

5 Static Distortions

Our assumption that the promise of duration for a given innovation is sufficient for computing social benefit from that innovation is important, and has strong implications. It implies that the social return to delivering duration for a given innovation does not depend on the way the firm's other

Figure 3: Distribution of duration

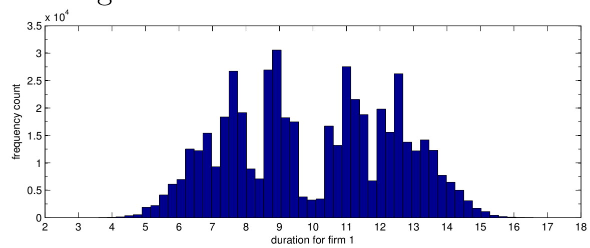

innovations are being treated: the firm's incentive to innovate is determined entirely by the duration promise for the given innovation, and not what the firm is granted for other innovations. This does not imply that the firm can only profit while it gets preference, it simply implies that all units 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 future 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

$d^* = \arg\max_{0 \le d \le 1/r} R(d)$. 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 in those papers, that innovation is granted duration d^* in order to maximize the planner's value in the space of rewards by product market treatment.

Consider first the no exclusion case with natural rate $\underline{d}$. If $d^* \le \underline{d}$, it is immediate that the planner can implement every innovation at the Arrow-Nordhaus duration d^*. So the model only has a dynamic tradeoff if $d^* > \underline{d}$, to ensure that the planner faces scarcity in market time. 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 has no selling rights are not zero, but just less than when it hast them. This might be due to imperfect licensing or services it provides for the leading edge firm, 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. The marginal value of duration is also likely to decrease and so promised duration and the probability of exclusion. For instance, in the case of quadratic cost function it is easy to show that $R'(d) = \left(\frac{1}{r} - d\right)(1-\gamma)^2$ and thus the marginal value of providing duration decreases with γ.

6 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 near-permanent 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 polices where rights are is always granted immediately for any innovation that is implemented (forward exclusions). 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 A: Proofs for Forward Exclusion Rights

Proof of Lemma 1

Lemma 1. The function

$$F\left(\overline{x}\right) \equiv \max_{a \leq x(z) \leq b, y} \int_0^1 f\left(x\left(z\right), z\right) dz + \int_0^1 g(y(z), z) dz$$
$$s.t. \int_0^1 x\left(z\right) dz + \int_{z:y(z)>0} dz = \overline{x}$$

is weakly concave.

Proof. Take any two $\bar{x}_1 < \bar{x}_2$ with optimal actions $x_1(z), y_1$ and $x_2(z), y_2$. We have □

Claim 13. There exists a set A such that

$$\begin{aligned}
\int_A x_1(z)dz &= \phi\bar{x}_1 \\
\int_A y_1(z)dz &= \phi\bar{x}_1 \\
\int_A x_2(z)dz &= \phi\bar{x}_2 \\
\int_A y_2(z)dz &= \phi\bar{x}_2 \\
\int_A f(x_1(z),z)dz + \int_A g(y_2(z),z) &= \phi F(\bar{x}_1) \\
\int_A f(x_2(z),z)dz + \int_A g(y_2(z),z) &= \phi F(\bar{x}_2)
\end{aligned}$$

Proof. Define four measures

$$\begin{aligned}
\mu_1(E) &= \int_E x_1(z) + \int_{z\in E:y_1(z)>0} dz \\
\mu_2(E) &= \int_E x_2(z) + \int_{z\in E:y_2(z)>0} dz \\
\mu_3(E) &= \int_E f(x_1(z),z) + \int_E g(y_1(z),z) \\
\mu_4(E) &= \int_E f(x_2(z),z) + \int_E g(y_2(z),z)
\end{aligned}$$

and the mapping from measurable subsets of [0,1] into R^4

$$\mu(E) = \{\mu_1(E), \mu_3(E), \mu_3(E), \mu_4(E)\}$$

Note that the range of μ includes $\{0,0,0,0\}$ and $\{\bar{x}_1, \bar{x}_2, F(\bar{x}_1), F(\bar{x}_2)\}$, moreover, by Theorem 5.5 in Rudin (1991), the range of μ is convex, so there must be some set A with $\mu(A) = \{\phi\bar{x}_1, \phi\bar{x}_2, \phi F(\bar{x}_1), \phi F(\bar{x}_2)\}$

Let $\bar{x} = \phi\bar{x}_1 + (1-\phi)\bar{x}_2$ and define

$$\hat{x}(z) = \begin{matrix} x_1(z) \text{ if } z \in A \\ x_2(z) \text{ if } z \in A^c \end{matrix}$$

$$\hat{y}(z) = \begin{matrix} y_1(z) \text{ if } z \in A \\ y_2(z) \text{ if } z \in A^c \end{matrix}$$

where A is as described in the claim. Then $\int_0^1 \hat{x}(z)\,dz + \int_{z:\hat{y}(z)>0} dz = \phi\bar{x}_1 + (1-\phi)\bar{x}_2$ (i.e. it is is feasible at $\bar{x}$), so $F(\bar{x}) \geq \int f(\hat{x}(z), z)dz + \int_0^1 g(y(z), z)dz$. But $\int f(\hat{x}(z), z)dz + \int g(y(z), z)dz = \phi F(\bar{x}_1) + (1-\phi)F(\bar{x}_2)$, so F is concave, since we have constructed a feasible policy that does at least as well as the convex combination of the payoffs. □

Note that the problem in (1) fits into this setting directly, with the interpretation of d_1 as x and d_2 as y, with $d_2 = 0$ interpreted as firm one continuing as the incumbent.

To show that V is differentiable, for any d, define $\underline{V}(\epsilon)$ by

$$r\underline{V}(\epsilon) = \left\{ \begin{array}{c} \lambda\left(\int (R(d_1(\theta),\theta) + V(d_1(\theta)) - V(d)\right) dF(\theta) + \\ \lambda \int_{\bar{\theta}_\epsilon}^{\infty} (R(d_2(\theta),\theta) + V(d_2(\theta)) - V(d))\, dF(\theta) \end{array} \right\}$$

Where, for duration d, $\bar{\theta}_\epsilon$ solves

$$r(d+\epsilon) = 1 + \lambda(\int d_1(\theta)dF(\theta) - (d+\epsilon)) - \lambda(1 - F(\bar{\theta}_\epsilon))(d+\epsilon)$$

so that promise keeping is satisfied for fixed $d_1(\theta)$ and $d_2(\theta)$, which are taken to be their optimum values at duration d.[16] The function $\underline{V}$ is differentiable, $\underline{V}(0) = V(d)$, and, since it is computed for a feasible policy, $\underline{V}(\epsilon) \leq V(d+\epsilon)$, so V is differentiable.

Proof of Proposition 4:

Proof. Starting at d, let $\{d_n, p_n\}$ denote the sequence giving the continued duration and probability of replacement after n consecutive arrivals of he incumbent. Naturally, promise keeping satisfies:

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

We now show that in the optimal mechanism, $p_{n+1} > 0$ implies $p_n = 1$. Suppose towards a contradiction that $p_{n+1} > 0$ and $p_n < 1$. Consider an alternative path identical to the original one with the exception that $\tilde{p}_n = p_n + \varepsilon$ and $\tilde{p}_{n+1} = p_{n+1} - \delta$ where δ is calculated so that $\tilde{d_n} = d_n$. We show

[16] If necessary,$d_2(\theta)$is extended in any continuous way below $\bar{\theta}_0 = \bar{\theta}$

now that there exist such $\delta < p_{n+1}$ for ε small. Using the functional equation for d_n on the original and alternative path:

$$\begin{aligned} rd_n &= 1 + \lambda \left(d_{n+1} - d_n\right) - \lambda p_n d_n \\ &= 1 + \lambda \left(\tilde{d}_{n+1} - d_n\right) - \lambda \tilde{p}_n d_n \end{aligned}$$

giving the necessary and sufficient condition:

$$\tilde{d}_{n+1} - d_{n+1} = \left(\tilde{p}_n - p_n\right) d_n \equiv \varepsilon d_{n.}$$

Now observe that since $p_{n+1} > 0$ it follows that $\frac{1}{r+2\lambda} \leq d_{n+1} < \frac{1}{r}$ so it can be increased by reducing p_{n+1}. This proves that the alternative path is feasible and gives $d_n = \tilde{d}_n$ for $\varepsilon > 0$ small enough. □

Finally, we show that this alternative path is strictly better. Both sequences of durations are identical except for $\tilde{d}_{n+1} > d_{n+1}$. From the perspective of the incumbent, it induces the same sequence of investments except for $\tilde{\Delta}_{n+1} > \Delta_{n+1}$, so it is strictly better. Moreover, the variation is neutral with respect to the expected value generated when replacing the incumbent since since initial duration (the expected discounted time of this arrival) is unchanged.

Having established that $p_{n+1} > 0$ implies $p_n = 1$, it follows that the sequence $\{p_n\}$ is weakly decreasing. Moreover, there can be at most one period where $0 < p_n < 1$ for otherwise the condition that we proved would be violated. This proves that the optimal sequence has $p_n = 1$ for $n = 1, .., N-1$, $0 \leq p_N < 0$ and $p_n = 0$ for $n \geq N+1$. This also implies that the sequence d_n is increasing to the point where $d_{N+1} = \frac{1}{r}$ and the outsider is excluded forever.

Proof of Lemma 5:

Proof. Starting at time $t = 0$ consider the stopping time T_1 defined by player one gets its first arrival before player two. This stopping time has an associated density $f(t) = \lambda \int e^{-2\lambda t} dt$ and associated expected discount factor $Ee^{-rT_1} = \lambda \int e^{-(r+2\lambda)t} dt = \frac{\lambda}{r+2\lambda}$. Denote this expected discount factor by β. By independence, the stopping time T_{1n} defined by player one gets its n^{th} arrival prior to player two getting an arrival has expected discount factor β^n. This is also the expected discount factor for the event: "player two gets its first arrival after $n-1$ arrivals of player one."

Start with the baseline policy where all arrivals are implemented granting the incumbent player (say player one) duration $d^* = \frac{1}{r+\lambda}$. Consider the following deviation: if after $n-1$ consecutive arrivals for player one there is an arrival for player two, do not implement that arrival but then return to the baseline plan. If instead player one gets its n^{th} arrival, return to the baseline plan. This is a plan that delivers a sequence $\{d_0, d_1, d_2, ..., d_{n-1}, d_{n+1,...}\}$ with $d_i > d^*$ for all $i < n$ and equal to d^* for $i \geq n$. The increased durations for player one have the benefit of larger innovations but there is the cost of missing one potential innovation and the corresponding value $R^* = R(d^*)$. As of time zero, the expected discounted cost is $\beta^n R^*$. The calculation of the benefits is slightly more complex and follows here. We start with $n-1$ and using the recursive definition of duration, it easily follows that:

$$(r + 2\lambda)(d_{n-1} - d^*) = \lambda d^*$$

Incidentally, note that

$$\frac{d_{n-1} - d^*}{d^*} = \frac{\lambda}{r + 2\lambda} = \beta.$$

More generally,

$$(r + 2\lambda)(d_i - d^*) = \lambda (d_{i+1} - d^*)$$

which dividing through by d^* gives

$$\frac{d_i - d^*}{d^*} = \beta \frac{d_{i+1} - d^*}{d^*} = \beta^{n-i}.$$

In consequence, $d_i = (1 + \beta^{n-i}) d^*$ and in particular $d_0 = (1 + \beta^n) d^*$. The value of the given alternative plan:

$$\begin{aligned} W_n - W^* &= \sum_{i=0}^{n-1} \beta^i (R(d_i) - R(d^*)) - \beta^n R(d^*) \\ &\geq \sum_{i=0}^{n-1} \beta^i R'(d_{n-1})(d_i - d^*) - \beta^n R(d^*) \\ &= \sum_{i=0}^{n-1} \beta^i R'((1+\beta) d^*) \beta^{n-i} d^* - \beta^n R(d^*) \\ &= \beta^n [nR'((1+\beta) d^*) d^* - R(d^*)] \end{aligned}$$

It is straightforward to see that since $(1+\beta) d^* < \frac{1}{r}$, $R'((1+\beta) d^*) > 0$ and so for sufficiently large n, $W_n > W^*$. $\square$

Derivation of formula for $V'(d)$

In region with $p = 1$

Using the functional equations:

$$\begin{aligned} (r+\lambda) V(d) &= \lambda\alpha \left[R(d_1) + V(d_1)\right] + \lambda(1-\alpha)\left[R(d_2) + V(d_2)\right] \\ (r+\lambda) d &= 1 + \lambda\alpha d_1 \end{aligned}$$

Using the second equation, $\lambda\alpha\partial d_1/\partial d = r + \lambda$ and now totally differentiating the first equation with respect to d after substitution yields:

$$V'(d;\alpha) = R'(d_1) + V'(d_1;\alpha). \tag{6}$$

In region with $0 < p < 1$

The promiss-keeping constrant can be used to rewrite the functional equation in a more convenient way:

Multiplying both sides of the PK constraint by r and adding to both sides $\lambda(1-\alpha)p$, gives:

$$[r + \lambda(\alpha + (1-\alpha)p)]\, rd + \lambda(1-\alpha)p = r + \lambda(\alpha + (1-\alpha)p)$$

so

$$\frac{\lambda(1-\alpha)p}{r + \lambda(\alpha + (1-\alpha)p)} = 1 - rd. \tag{7}$$

In addition, solving for rd and simplifying gives

$$rd = \frac{r + \lambda\alpha}{r + \lambda(\alpha + (1-\alpha)p)}$$

implying

$$\frac{\lambda\alpha}{r + \lambda(\alpha + (1-\alpha)p)} = \frac{\lambda\alpha}{r + \lambda\alpha} rd. \tag{8}$$

Now substituting (7) and (8) in the dynamic programing equation gives:

$$V(d) = \frac{rd}{r + \lambda\alpha}\lambda\alpha\left(R\left(\frac{1}{r}\right) + V\left(\frac{1}{r}\right)\right) + (1 - rd)\left(R(d_2) + V(d_2)\right) \tag{9}$$

Using $V\left(\frac{1}{r}\right) = \frac{\lambda\alpha}{r} R\left(\frac{1}{r}\right)$ gives:

$$R\left(\frac{1}{r}\right) + V\left(\frac{1}{r}\right) = \frac{r + \lambda\alpha}{r} R\left(\frac{1}{r}\right) = \frac{r + \lambda\alpha}{\lambda\alpha} V\left(\frac{1}{r}\right)$$

and substituting this in (9) gives the simplified dynamic programming equation:

$$V(d) = rdV\left(\frac{1}{r}\right) + (1 - rd)\left(R(d_2) + V(d_2)\right) \tag{10}$$

Solving forwards

Equations (6) and (10) can be used to forward substitute $V'(d_n)$ and this gives:

$$V'(d;\alpha) = \sum_{n=1}^{N} R'(d_n) - r\left[R(d_2) + V(d_2) - V\left(\frac{1}{r}\right)\right]. \tag{11}$$

Proof of Proposition 7

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 an outsider 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 the outsider without an idea does not find it worthwhile to pay f as one who does is indifferent between paying f or not.

Consider the deviation of paying t (the foreclosure fee), n arrivals after the intial 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

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

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 14. 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 $v_{n-1} - \beta^{n-1} t \geq v_{n-2} - \beta^{n-2} t$. □

Claim 15. $v_n - \beta^n t \geq v_{n+1} - \beta^{n+1} t$ 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 size 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 $v_{n+1} - \beta^{n+1} t \geq v_{n+2} - \beta^{n+2} t$. □

Proof of Propositions 8 and 9

Consider the problem:

$$\begin{aligned} rV(d;\alpha) &= max_{d_1,d_2,p}\alpha\lambda\left[R(d_1)+V(d_1;\alpha)\right]+(1-\alpha)\lambda p\left[R(d_2)+V(d_2;\alpha)\right] \\ s.t\ d &= 1+\alpha\lambda(d_1-d)-(1-\alpha)pd \end{aligned}$$

We show that $V(d;\alpha)$ is increasing in α for $d>d_2(\alpha)$ and small change in α.For induction, assume right hand side V is increasing in α. Take $\alpha'>\alpha$. Consider the following (suboptimal policy). For an arrival of the outsider, choose the same p and d_2 that are optimal for α. For an arrival of the incumbent grant the incumbent the same value d_2 with probability $p\left(\frac{\alpha'-\alpha}{\alpha'}\right)$ and d_1' with probability α/α' and the same d with probability $\frac{1-p}{\alpha'}(\alpha'-\alpha)$. Upon substitution, the promiss-keeping constraint reads

$$rd = 1+\lambda\alpha(d_1'-d)+p\lambda(\alpha'-\alpha)d_2-\lambda(1-\alpha)pd,$$

identical to the promiss-keeping for α except for the third term. It follows that $d_1'<d$. Choose α' sufficiently close to α so that $d_1'>d_2$. Denoting by $\tilde{V}$ the value of this policy, it satisfies the functional equation:

$$\begin{aligned} r\tilde{V} &= \lambda\alpha\left(R(d_1')+V(d_1',\alpha')-\tilde{V}\right)+\lambda(1-p)(\alpha'-\alpha)\left(R(d)+V(d;\alpha)-\tilde{V}\right) \\ &\quad +\lambda(1-\alpha)p\left(R(d_2)+V(d_2;\alpha')-\tilde{V}\right) \\ &\geq \lambda\alpha\left(R(d_1')+V(d_1',\alpha)-\tilde{V}\right)+\lambda(1-\alpha)p\left(R(d_2)+V(d_2;\alpha)-\tilde{V}\right) \\ &> \lambda\alpha\left(R(d_1)+V(d_1;\alpha)-\tilde{V}\right)+\lambda(1-\alpha)p\left(R(d_2)+V(d_2;\alpha)-\tilde{V}\right) \end{aligned}$$

where the first inequality follows from the induction hypothesis and the fact that since this policy is suboptimal $V(d;\alpha)\geq\tilde{V}$ and the second inequality from the fact that strict concavity$R(\cdot)+V(\cdot;\alpha)$ implies it is strictly decreasing for durations above its maximizer d_2 and that $d<d_1'<d_1$. This proves that $\tilde{V}>V(d,\alpha)$. Since the policy considered is feasible (and actually suboptimal), it follows that $V(d,\alpha')>V(d,\alpha)$.

We consider now the proof of Proposition 9. Since the value functions are concave, they are almost everywhere differentiable. We show that $\partial V(d;\alpha)/\partial d$ is increasing in α. Using the promiss-keeping constraint and maintaining fixed p (an envelope condition argument)

$$\partial d_1/\partial d = \frac{r+\lambda(\alpha+(1-\alpha)p)}{\lambda\alpha}.$$

Differentiating the functional equation it follows that:

$$(r+\lambda(\alpha+(1-\alpha p)))(\partial V(d;\alpha)/\partial d)=\lambda\alpha\frac{\partial}{\partial d_1}(R(d_1)+V(d_1;\alpha))\frac{\partial d_1}{\partial d}$$

that after substitution implies $\partial V(d;\alpha)/\partial d=\frac{\partial}{\partial d_1}(R(d_1)+V(d_1;\alpha))$. Assume inductively that $\partial V(d_1;\alpha)$ is increasing in α. Using the promiss keeping condition, d_1 decreases in α and by the strict concavity of R and the induction assumption it follows that $\partial V(d;\alpha)/\partial d$ is increasing in α.

Your old proof

Rewrite (4) as

$$\begin{aligned} V_M(d) &= \max_{d_1,p}\left\{\ \tfrac{1}{r+\lambda}\left(R(d_1)+V_M(d_1)\right)+p\kappa\ \right\} \\ s.t.& \\ d &= 1+\frac{\lambda}{M}(d_1-d)-\frac{\lambda(M-1)}{M}pd \end{aligned}$$

Where κ is a positive constant. Define the operator T, mapping continuous, bounded functions on $[1/(r+\lambda),1/r]\times\mathbb{N}$ by

$$\begin{aligned} TV_M(d) &= \max_{d_1,p}\left\{\ \tfrac{1}{r+\lambda}\left(R(d_1)+V_M(d_1)\right)+p\kappa\ \right\} \\ s.t.& \\ d &= 1+\frac{\lambda}{M}(d_1-d)-\frac{\lambda(M-1)}{M}pd \end{aligned}$$

First we show that V must decrease in M. To see this, fix M and the associated optimal $d_1(d)$ policy; if the planner chooses the same $d_1(d)$ for $M-1$, he can choose a strictly greater implementation probability p and still satisfy promise keeping. Therefore T maps functions decreasing in M to functions decreasing in M, and the fixed point must have this property.

We follow a similar approach for showing that differences are decreasing in M. Let $d_1^M(d)$ be the optimal policies for M. We add to the definition of T the constraint that $d_1(d)$ must be weakly increasing. Note that, since this is a feature of the optimal policy for the true value function, the fixed point that imposes this constraint must by the fixed point of the problem without the constraint.

Fix $d' > d$ and M. Now, for $M-1$, consider the optimal policy $d_1^{M-1}(d)$ and the policy for d' $x = d_1^{M-1}(d) + (d_1^M(d') - d_1^M(d))$ Since $d_1^{M-1}(d) < d_1^M(d)$, $x \leq 1/r$ But $R(x) - R(d_1^{M-1}(d)) > R(d_1^M(d') - R(d_1^{M-1}(d))$ since the difference between he policies is the same, $d_1^{M-1}(d) < d_1^M(d)$, and $R$is concave; likewise $V_{M-1}(x) - V_{M-1}(d_1^{M-1}(d)) > V_M(d_1^M(d')) - V_M(d_1^M(d))$ if V is concave and has decreasing differences. Finally, denoting $p_M(d_1)$ the p which solves the promise keeping constraint for M and d_1 , $p_{M-1}(x) - p_{M-1}(d_1^{M-1}(d)) > p_M(d_1^M(d') - p_M(d_1^M(d)$. Therefore T maps concave functions with decreasing differences into ones with decreasing differences, and therefore the fixed point must have this property.

Appendix B: Optimality of Complete Exclusion Rights

We begin with the most general problem, where both firms have promises that may not add up to all available time, and prove that the optimal policy can be solved by the dynamic program (5).

We introduce the following notation. 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 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. If nothing arrives, the planner may change the duration promise by

$\dot{d}$. The dynamic program is

$$rV(d,\underline{d}) = \max_{\substack{d_1^n,d_1^c,d_2,\dot{d},x \\ \underline{d}_1^c,\underline{d}_1^n,\underline{d}_2,\dot{\underline{d}},\underline{x}}} \left\{ \begin{array}{c} \lambda\left(R(d_1^n) + V(\max\{d_1^n, d_1^c\}, \underline{d}_2) - V(d,\underline{d})\right) + \\ \lambda\left(R(\underline{d}_1^n) + V(d_2, \max\{\underline{d}_1^n, \underline{d}_1^c\}) - V(d,\underline{d})\right) + \\ V_1(d,\underline{d})\dot{d} + V_2(d,\underline{d})\dot{\underline{d}} \end{array} \right\} \quad (12)$$

$$s.t. \quad (13)$$

$$rd = x + \lambda(d_1^c - d) + \lambda(d_2 - d) + \dot{d} \quad (14)$$

$$r\underline{d} = \underline{x} + \lambda(\underline{d}_1^c - \underline{d}) + \lambda(\underline{d}_2 - \underline{d}) + \dot{\underline{d}} \quad (15)$$

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:

$$\begin{aligned} 0 &\leq \max\{d_1^n, d_1^c\} + \underline{d}_2 \leq 1/r \\ 0 &\leq d_2 + \max\{\underline{d}_1^n, \underline{d}_1^c\} \leq 1/r \\ 0 &\leq x + \underline{x} \leq 1 \end{aligned}$$

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 $\dot{d}$ 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 is[17]

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

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

[17] Subscripts denote derivatives.

We now verify that V is in fact concave. If it is, then imposing the earlier results we have a simplified problem[18]

$$\begin{aligned} rV(d,\underline{d}) &= \max_{\substack{d_1,d_2,x \\ \underline{d}_1,\underline{d}_2,\underline{x}}} \left\{ \begin{array}{l} \lambda\left(R(d_1)+V(d_1,\underline{d}_2)-V(d,\underline{d})\right)+ \\ \lambda\left(R(\underline{d}_1)+V(d_2,\underline{d}_1)-V(d,\underline{d})\right)+ \end{array} \right\} \quad (16) \\ & s.t. \\ rd &= x+\lambda(d_1-d)+\lambda(d_2-d) \\ r\underline{d} &= \underline{x}+\lambda(\underline{d}_1-\underline{d})+\lambda(\underline{d}_2-\underline{d}) \quad (17) \end{aligned}$$

Lemma 17. *V is concave*

Proof. The Bellman equation can be rewritten as

$$V(d,\bar{d}) = \frac{1}{r}\frac{\lambda}{r+2\lambda}\max(R(d_1)+R(\underline{d}_1)+V(d_1,\underline{d}_2)+V(d_2,\underline{d}_1))$$

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. □

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

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, $d+\bar{d}=1/r$ □

[18] In the language of the dynamic program (16), the forward exclusion case adds the restriction that it $\underline{d}_1>0$, then $d_2=0$, and if $\underline{d}_2>0$ then $d_1=0$. One can use this structure to derive the problem studied in that section.

Proof of Proposition 12:

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

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, Massachussets, 1969.

[16] Ted O'Donoghue, Suzanne Scotchmer, and Jacques-FranÃ§ois 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.