

Optimal Patent Policy with Recurrent Innovators

Hugo Hopenhayn* and Matthew Mitchell†

Abstract

We characterize the solution to the problem facing a planner who must allocate preferential treatment across two firms who can use preferential treatment to make profits, and in turn are encouraged to innovate by the provision of preferential treatment. The planner, because he can allocate preference to a single firm for multiple innovations at any point in time, backloads rewards, giving the firm with the preponderance of the promise a preference that corresponds to a patent offering complete exclusion. We show that the promised duration of preference evolves in a region that is determined by the optimal promises in a static version of the model, which corresponds to a classic static patent problem of Arrow and Nordhaus. When there are no static distortions, so that the optimal static patent lasts forever, the optimal policy we study leads to monopoly, in the sense that one firm is excluded even though it is getting useful ideas. We show that these basic results hold even if the planner is forced to use a restricted set of policies where preference is always granted immediately for any innovation that is implemented.

1 Introduction

In this paper we analyze how a planner rewards two innovators who contribute to a common research agenda, for instance improving a given product. Moral hazard precludes rewarding with a cash prize; instead, rewards

*Department of Economics, UCLA.

†Rotman School of Management, University of Toronto.

must be earned through the allocation of preferential treatment, such as a patent in the product market. 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. We model this by allowing the planner to only prefer one firm at a time. As a result, time to allocate preferential treatment may be scarce. We ask how the planner can best allocate preferential treatment in the product market in order to encourage innovation.

We focus on two features of the optimal policy. The first is the way that the patent right evolves with history. Recent literature has stressed that heterogeneity may lead to different rewards for different innovators. Here there is no heterogeneity built into the structure, but the rewards are history dependent, leading to ex post heterogeneity in the reward for different innovations, and as a result, heterogeneity in the degree of innovation over time. We discuss how one can interpret this heterogeneity as a patent menu in the spirit of Scotchmer (1999), Cornelli and Schankerman (1999), and Hopenhayn and Mitchell (2001). We also discuss how heterogeneous rewards can be decentralized in a manner familiar from the buyouts used in Hopenhayn, et al. (2006).

Because the model has repeat innovators, we can use it to study the evolution of market power under the optimal policy. We show that the scarcity of market preference can lead to monopolization, in the sense that one of the innovators is eventually promised preference at nearly all instants under the optimal contract. Counterintuitively, monopolization occurs even when the marginal benefit of preferential treatment for the monopolizing firm is zero at monopoly. This arises because of the planner's incentive to backload rewards. The motive for backloading, however, is different from the standard backloading intuition from Becker and Stigler (1974) and Lazear (1981). In those papers, backloading is beneficial because not only does it generate strong incentives late, but also because it generates strong incentives early, as the agent works hard to reach the point where the backloaded incentives kick in. Here we introduce a different motive for backloading. Since the planner can prefer an agent at a given time for multiple innovations, it is useful to push preference to the later period where it is applicable to more than one innovation. We show more generally that the length of preference increases with successes, eventually reaching the static optimum patent length familiar from Arrow (1962) and Nordhaus (1969).

Our model links the recent literature on optimal patent menus with a

much more established literature on the role of competition between firms in patent policy. In papers in the former category, the problem is either static (for instance in Scotchmer (1999) or has a sequence of agents that never recur (Hopenhayn, et al (2006)). Here we consider a case where there is both dynamics and recurring innovators. We consider a general structure that allows for many possible market structures. The inclusion of recurrent innovators is interesting because it allows us to think about the dynamic evolution of monopoly power under the optimal policy. Does the planner use the policy to avoid foreclosing innovators, or does it rely on backloaded incentives to efficiently reward innovators for a sequence of successes?

In this sense our paper has similarities with a recent paper by Acemoglu and Akgigit (2010). In that paper, the authors use a growth theory structure similar to Aghion, et al (2001) and consider policies that change depending on the quality differential between the firm's most recent innovations. Acemoglu and Akgigit 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 takes those ideas to consider a more general class of policies in a more abstract environment. Although our structure does not nest theirs, the intuition about backloading that we develop applies to their environment, and therefore helps develop further intuition for their numerical results. Moreover, it shows a general class of problems where what they term "state dependency" of rewards is definitely optimal.

Our abstract model can be described concisely. Two innovating firms randomly receive opportunities to generate social value (innovations) at independent rates. The planner rewards the innovators by offering, at each instant, a given innovator the opportunity to profit from some subset of past innovations; we term this grant preference, but one can think of it as a patent on some set of innovations. We take the planner's payoff from delivering preference to a given innovator for an innovation he generates as a reduced form function, but section 2 describes an explicit quality ladder model that delivers the general structure we study.

The planner, then, needs to decide how to determine the allocation of preference at each state and history. This is potentially a very complicated problem, but the nature of our setup allows us to show that optimal policies always involve complete exclusion: preference for all past innovations at any time that there is preference for any past innovations granted to a given innovator. Further, the optimal policy always eventually reaches a state

where all future instants are promised to at least one innovator, a situation we term duopoly.

With these results in hand, we turn to characterize the optimal policy under duopoly more fully. We show that duration moves within an interval bounded by each firm's statically optimal promise of preference, nearing the endpoints with positive probability. As a corollary, when the static optimal level of preference is infinite (as it is in the quality ladder model we begin with), the optimal policy eventually enters a state where nearly the entire future has been promised to one firm; the other firm is (nearly) completely foreclosed, its ideas unused.

More generally, the optimal policy trades off the planner's desire to take advantage of both innovator's ideas against the benefits of taking full advantage of a given arrival by promising sufficiently long preference that some of the other firm's ideas must be used at a lower rate. In the optimal policy, the history of arrivals determines the current distribution of promises, which in turn determines how much preference will optimally be promised to each firm. This varying state generates the heterogeneity in otherwise homogenous ideas and innovators.

The preference itself evolves in a stark way. The firm with the greater duration promise gets preference. As a result, when the future promise is skewed sufficiently toward one firm, even an arrival of an idea by the competitor leaves the highly promised firm with preference. That preference includes preference for the new innovation, invented by the rival. The rival's payoff to generating the innovation is that the promise becomes less skewed, moving the state closer to its favor, where it gains preference, which one can interpret as taking control of the patent right.

Our results use a structure that is amenable to recursive methods. Although this is critical in our ability to make progress in solving the problem, it also leads to the extreme form of preference (exclusive rights) that arises. This, in turn, leads to situations where firm innovate, rewarded by preference not immediately, but in expectation in the future. To the extent that these policies seem unusual, and are driven by the underlying recursive nature of the setup, we consider an alternative structure that avoids these results. In particular, we consider a second regime where complete exclusion is not available to the planner.. We call this alternative regime "exclusive rights" (following Hopenhayn, et al. (2006)); the planner can grant a firm preferential treatment at a single instant for any innovations from a sequence of innovations by that firm that have arrived consecutively, without an interven-

ing innovation by the other firm. In other words, whenever the competitor is granted some preferential treatment, the incumbent leading firm loses all rights to preferential treatment. This is comparable to a standard notion of patents used in models of cumulative innovation, where a leading firm can maintain market position by filing for new patents, until a competing improvement makes a new firm the market leader.

We show that under the exclusive rights regime, we still get backloading, and the preference promises still evolve in the same interval bounded by the static patent lengths. Moreover, the chance of a competitors innovation being implemented (and the incumbent being ousted) is declining with successes by the incumbent. This is familiar from the state dependent policies that Acemoglu and Akcigit (2010) compute as optimal policies in the step-by-step model they study. We show the sense in which this structure is driven by a backloading incentive present in their paper as well.

2 Preferential Treatment and Patents in Quality Ladders

In the next section we introduce an abstract model of a planner who encourages effort through the allocation of "preference." For a concrete example of patents-as preference, consider the quality ladder structure example explored in the patent literature in papers such as O'Donoghue et al. (1998) and Hopenhayn, et al (2006). In particular, suppose that two firms periodically get ideas that can be turned into innovations of size Δ in exchange for cost $c(\Delta)$. Here Δ (and therefore $c(\Delta)$) is the non-verifiable feature that necessitates patents. It is well known in the literature (and very intuitive) that, if the degree of innovation cannot be verified by the planner, the innovators cannot be rewarded with transfers, since they could claim the transfer and not pay the costs of innovation. The planner can use the promise of rights to sell in a market to successfully induce effort.¹ There are no production costs, only the costs of innovation $c(\Delta)$.

Innovation is cumulative in the sense that an innovation reflects an improvement to the prior state-of-the art; a firm's innovation results in a product that is Δ units of quality q better than the last innovation. A single

¹See Hopenhayn, et al. (2006) for more on this issue in the cumulative context, in a quality ladder model like this one.

consumer purchases one physical unit of quality q , choosing which variety in order to maximize $q - p$, where p is the price paid for the quality q variety.² We take competition to be Bertrand, so that the leading edge product is always the one sold in equilibrium, and the social surplus at any point in time is the quality q either in the form of profits for the firm selling the leading edge product, or as consumer surplus if $p < q$. As a result, if r is the discount rate, every innovation yields $\Delta/r - c(\Delta)$ additional units of present discounted surplus for the planner. The leading edge product (i.e. the most recent innovation) sells for a price equal to its quality differential over the other firm's best product, since that trailing product will be provided at cost (zero) in the pricing game, and the leader charges the quality differential.

We define *preference* for an innovation as the ability to generate profits from it. In this case, then, we will say that a firm is given preference for a given innovation if no other firm can offer a (weakly) better product, that is, one that embodies the preferred innovation or one that is better. The firm, therefore, makes profits equal to the sum of the quality level of its preferred innovations.³

As in Hopenhayn, et al. (2006) we study policies that are described by contingent rights for a given innovation. In particular, the planner promises, at any arrival, an expected discounted length of time $d \geq 0$ during which the innovator will be given preferential treatment for an innovation made under that idea. Since the value of d , which we term the *duration* of preference, is in present discounted terms, it might come in many ways, for instance, a T period patent (where preference is guaranteed for all T periods) would have $d = (1 - e^{-rT})/r$. We use the language of duration to describe recursively how the optimal policy proceeds, considering arbitrary duration policies, which may be contingent on future arrivals as well as the passage of time. A patent that offered T periods of protection for sure, followed by T' units of additional protection with probability $1/2$ would have $d = (1 - e^{-rT})/r + \frac{1}{2}e^{-rT}(1 - e^{-rT'})/r$. Since the planner can choose a preference policy at every instant, this duration can be delivered in any contingent way, evolving over time or with later arrivals of innovators, and with the identity of the innovator that arrives with an idea. Of course, since (discounted) time is not unbounded,

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

³Note that here it is impossible to give preference for a discontinuous set of innovations (by definition), but that is without loss since it is never optimal to offer such preference in the policies we study.

the maximum possible promise of sure preferential treatment forever is $1/r$. This dynamic budget constraint of the planner's incentive tool is the key feature of the model.

Since the firm's profits increase by Δ with an innovation of size Δ , when the firm makes an innovation it solves

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

The planner's benefit from the innovation, since the innovation contributes forever to either profit or consumer surplus, is $R(d) = \Delta(d)/r - c(\Delta(d))$. Note that

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

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

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

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

We then have that

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

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

One can generalize this example so that profits when the firm is not preferred are not zero, but just less than when the firm is preferred. This might be due to services it provides for the leading edge provider, in order to make the innovations work efficiently. Such an environment would mean that

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

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

While these examples have the feature that there are no static distortions, so that $R(d)$ is maximized at $1/r$, we will not limit ourselves to such examples in what we study below. Focusing on the case with no static distortions, however, is interesting for at least two reasons. First, it highlights the role of the dynamic force that we study, namely the scarcity of preference relative to the optimal level, without any other source of inefficiency. Further, the work of Gilbert and Shapiro (1990) suggests that a planner who allocates patent rights together with the ability to regulate the strength of preference per period (for instance through patent breadth or direct price controls) will choose, in many circumstances, a long, narrow patent in the single innovation context.

3 Model

We now introduce the abstract model of allocation of preference that we study for the remainder of the paper, and which nests the example of the previous section. There is continuous time and an infinite horizon. There are two agents (which we sometimes call firms or innovators) and a principal (or planner). Each firm receives an opportunity to generate value for the planner with independent Poisson arrival rate λ .⁴ The planner cannot simply pay for the effort. One might imagine that cash for transfer is not available, or that, as in the patent example described above, the planner cannot observe inputs or outputs and therefore cannot effectively use cash as a reward. The planner can, however, commit to future “preferential treatment” for the agent, which will allow the holder to generate value from the innovation.

We call the use of these opportunities by the firms *innovations*, but it is not the only one that can be mapped into our structure. For instance, the planner could be a firm selling a product, where customers can be induced to buy with promises of future good treatment. An airline can reward today’s purchase with promises of future good seat assignments available only to frequent flyers. These rewards are scarce; only a limited number of customers can be allocated the good seat assignment. The firm trades off using current promises to encourage sales against the fact that current promises restrict the possibility of later promises to other potential customers. Perhaps an organization has a limited ability to assign a single “boss” who enjoys a share

⁴Although 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.

of the organization's output while in place, but multiple agents contribute to the improvement of the organization's ability to produce output. We use exclusively the language of the preference-as-patents example in what follows, but later discuss an alternative interpretation in the innovation case, where preference is viewed as preference by a competition authority.

At any given instant preferential treatment can be granted to a given firm for any list of past innovations.⁵ In other words, for a list of arrivals of ideas $I = \{1, 2, \dots, N\}$, the preference is the identity of the innovator who receives preference $i \in \{1, 2\}$ and a set $P \subseteq I$ of innovations for which the agent receives preference. Since P could be empty, it is possible that no agent receives preference at a given instant.

When an idea arrives, the contract prescribes future instants during which the innovator will be assigned preference for that idea. We summarize the contract by the *duration* of preference it promises the innovator for the idea: the expected discounted amount of time when the innovator will receive preference for the idea that he just received under the contract. We take the planner's lifetime payoff from a promise of d units of time of future preference to an arriving idea as $R(d)$. Interpreted as a patent right, the function $R(d)$ describes the present discounted marginal benefit from the improvement generated from the innovator promised d , net of the innovator's cost and any product market distortions generated by the preferential treatment. In the example of the last section, for instance, the more duration of exclusive production is allocated to a given arrival, the more profits can be made by the innovator with the idea, and in turn the higher will be the level of innovation and the planner's payoff. We assume R is a differentiable and concave function, and we normalize $R(0) = 0$.⁶

⁵Without loss of generality we let preference for a given innovation be allocated to a single innovator at any instant; this is without loss because changing preference over time can effectively "split" preference for a given innovation across innovators. Such a split will, in addition, not be optimal.

⁶To add the notion of fixed costs, simply assume that planner's reward can be written as

$$R(d) = \begin{cases} 0, & d < F \\ \tilde{R}(d), & d \geq F \end{cases}$$

where \tilde{R} is concave and differentiable as before. The role of F is to model the idea that innovations may have fixed costs; nearly zero duration may not be enough for the project to be implemented at all, but above some point the fixed costs are covered and the innovation occurs (at some positive size). One can interpret $\tilde{R}(d)$ as the social return were the fixed costs not present. Note that in terms of the underlying innovation technology, one could

We interpret the shape of R as making statements about the product market where patent rights are granted. If $R'(1/r) = 0$, we say that there are 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. More generally, let $d^* = \arg \max_{0 \leq d \leq 1/r} R(d)$. The value of d^* is analogous to the (discounted) optimal patent length in a static model like Arrow (1962) or Nordhaus (1969); since the market is driven by only one innovation, that innovation is granted duration d^* in order to maximize the planner's value in the space of rewards by product market treatment.

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

assume either a fixed cost to be paid before any improvement could be generated, or an indivisibility in the spirit of Boldrin and Levine (****) where only ideas of at least a threshold size are innovations at all. Either could give rise to returns for the planner as described by $R(d)$. The analysis below still applies.

Since R does not depend on the preference offered to other innovations, in particular there is no impact of offering the current firm preference for the *other* firm's past innovations; such an offer could transfer resources to the current firm, so long as the transfer is assumed to be welfare neutral, but it has no impact on the profitability of the firm's current innovation, and therefore no impact on the amount of effort they exert, and in turn on the planner's payoff from the current innovation. We therefore track only the promise of preference to a given firm for that firm's innovations.

First, suppose that the planner offers the most recent innovator preference for all past innovations until the next arrival, and continues to follow this strategy. Denote by \hat{d} the duration this offers to the new duration. When the other firm has an arrival, duration therefore drops to $1/r - \hat{d}$. Therefore \hat{d} solves

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

or $\hat{d} = \frac{r+\lambda}{r(r+2\lambda)} > 1/2r$. If $d^* < \frac{r+\lambda}{r(r+2\lambda)}$, it is immediate that the planner can implement every innovation at the Arrow-Nordhaus duration d^* , since the planner can always provide less than what is delivered under the plan that delivers \hat{d} . So we take $d^* > \hat{d}$ to ensure that the planner faces a trade-off across innovations. A simple description of this assumption is that if the planner offers the current innovator preferential treatment until the next arrival of the competitor, the protection is still insufficient relative to the Arrow-Nordhaus patent.

3.1 Dynamic Program

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

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} ; below we discuss situations where it is (at least eventually) sufficient to track only a single duration promise. In the interim, we speak generically about duration as d ; everything is symmetric across the innovators, so all statements apply equally to \underline{d} . In order to make everything completely symmetric, we refer to the promise to firm two in the event that firm one arrives by \underline{d}_2 , and so on.

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

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

$$\begin{aligned}
rV(d, \underline{d}) &= \max_{\substack{d_1^n, d_1^c, d_2, \dot{d}, x \\ \underline{d}_1^n, \underline{d}_1^c, \underline{d}_2, \dot{\underline{d}}, \underline{x}}} \left\{ \begin{array}{l} \lambda (R(d_1^n) + V(\max\{d_1^n, d_1^c\}, \underline{d}_2) - V(d, \underline{d})) + \\ \lambda (R(\underline{d}_1^n) + V(d_2, \max\{\underline{d}_1^n, \underline{d}_1^c\}) - V(d, \underline{d})) + \\ V_1(d, \underline{d})\dot{d} + V_2(d, \underline{d})\dot{\underline{d}} \end{array} \right\} (1) \\
rd &= x + \lambda(d_1^c - d) + \lambda(d_2 - d) + \dot{d} \\
r\underline{d} &= \underline{x} + \lambda(\underline{d}_1^c - \underline{d}) + \lambda(\underline{d}_2 - \underline{d}) + \dot{\underline{d}} \quad (2)
\end{aligned}$$

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:

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

The lower end of these constraints never binds, since duration is always at least \hat{d} .

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

Since greater d only makes the feasible set of possible choices of d_1 and d_2 smaller, it is immediate that $V(d, \underline{d})$ is weakly decreasing in each argument. This in turn implies that d_1^c can always be taken to be at least as big as d_1^n ; if d_1^c were less, raising it and offsetting the increase by lowering d to maintain promise keeping always does at least as well, and strictly better if V is strictly decreasing. Similarly, for $d_1^c > d_1^n$, reducing d_1^c at the margin is identical to increasing \dot{d} , and therefore we can let $d_1^c = d_1^n \equiv d_1$. However, in the modified program where $d_1^c = d_1^n \equiv d_1$ the envelope condition is⁷

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

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 summarize this in the following lemma.

Lemma 1 $d_1^c = d_1^n$ and $\dot{d} = 0$

⁷Subscripts denote derivatives.

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

This logic implies that the planner is always preferring all of an innovator's innovations if he is preferring any of the innovator's innovations. Since protecting innovations of the other innovator has no impact, one can take this to mean that any point of preference for an innovator is preference to the entire history of innovations. We therefore call this form of preference *complete exclusion rights*. In the patent context for a quality ladder, this will naturally map to an exclusive right to the entire quality ladder. Because this form of preference is strong, we also study below a weaker form, when protection ends whenever a firm's competitor arrives with an innovation.

In interpreting preference as a patent, one can imagine a patent right allowed the leading edge product to also exclude all competing products (except perhaps an outside good with quality normalized to zero). One might imagine that older product infringe on newer products, but not vice-versa. In O'Donoghue, et al. (1998) this is referred to as "lagging breadth." Then at any time that the firm is the market leader, it profits from all of its past improvements (and all of the competitors, but this is a transfer and does not impact incentives).

The planner's simplified problem is

$$rV(d, \underline{d}) = \max_{\substack{d_1, d_2, x \\ \underline{d}_1, \underline{d}_2, \underline{x}}} \left\{ \begin{array}{l} \lambda(R(d_1) + V(d_1, \underline{d}_2) - V(d, \underline{d})) + \\ \lambda(R(\underline{d}_1) + V(d_2, \underline{d}_1) - V(d, \underline{d})) + \end{array} \right\} \quad (3)$$

$$\begin{aligned} rd &= x + \lambda(d_1 - d) + \lambda(d_2 - d) \\ r\underline{d} &= \underline{x} + \lambda(\underline{d}_1 - \underline{d}) + \lambda(\underline{d}_2 - \underline{d}) \end{aligned} \quad (4)$$

Note that we have not imposed the domain constraints such as $d_1 + \underline{d}_2 \leq 1/r$. When that constraint binds, we enter a region where the planner can keep track of just one duration promise, since he has promised every instant to either one firm or another. In the remainder of this section we study (3) in order to argue that the constraint must bind at some future point in time with positive probability; as a result, duration promises always enter the region where the two firms are promised all of the instants. This allows us to study that problem (with only one state variable) in the next section since it will always eventually be the problem that the planner finds itself solving, once the domain constraint has bound.

Lemma 2 *V is concave*

Proof. The Bellman equation can be rewritten as

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

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 argue that for any value of the state (d, \bar{d}) , there must be a positive probability that eventually $d + \bar{d} = 1/r$. Let $B(d, \underline{d})$ be the (countable) set of durations visited in an optimal policy starting from (d, \underline{d}) . Let D denote the set of all points (d, \underline{d}) such that $d + \underline{d} = 1/r$, and let $P(X|d, \underline{d})$ be the probability of visiting X starting from (d, \underline{d}) , i.e. $P(B(d, \underline{d})|d, \underline{d}) = 1$

Lemma 3 *For all (d, \underline{d}) , $P(D|d, \underline{d}) > 0$*

Proof. Suppose the probability of D starting from (d, \underline{d}) is 0. In other words, starting from (d, \underline{d}) , the probability of the domain constraints

$$\begin{aligned} d_1 + \underline{d}_2 &\leq 1/r \\ d_2 + \underline{d}_1 &\leq 1/r \end{aligned}$$

ever binding at a future state is zero.

Consider a point in $B(d, \underline{d})$. First, we claim that for any such point, the constraint

$$x + \underline{x} = 1$$

can never bind. The reason is that, since the two firms' duration is always less than $1/r$, there must be some future points where the constraint does not bind, so one of the two x 's can always be lowered now and increased at some future time when the constraint does not bind, delivering the same duration. In turn, the other firm's x can be lowered.

But then, for a point in $B(d, \underline{d})$, we can study the problem without domain constraints, since the value is identical to the problem that does not impose them. In that problem it is immediate that V is additive across its two arguments, since the Bellman operator on the countable set $B(d, \underline{d})$ maps additive functions into additive functions. However, then we can use the first order conditions. The FOC for d_2 and the envelope condition imply

$$V_1(d_2, \underline{d}_1) = \mu(d, \underline{d}) = V_1(d, \underline{d})$$

Since V is additive, this implies $d_2 = d$. The FOC for d_1 is

$$R'(d_1) + V_1(d_1, \underline{d}_2) = \mu(d, \underline{d}) = V_1(d, \underline{d})$$

Again, since this is independent of \underline{d}_2 and \underline{d} and V is concave we see that $d_1 \geq d$ anytime $R'(d_1) > 0$. Thus duration weakly climbs for $d < d^*$; therefore it must converge. It cannot converge to some level less than d^* when the PK constraint does not bind, since there is no cost to giving more duration to the current innovation, and a strict benefit. Therefore d must be increasing to d^* . This argument applies to both firms; however, both d and \underline{d} can not rise for $d < d^*$, since $d^* > \hat{d} > 1/2r$, and eventually the domain constraint must be violated. Therefore $P(D|d, \underline{d}) > 0$. ■

The planner may want to save duration for the next arrival; since the planner cannot predict which firm will have the next innovation, saving duration allows flexibility to allocate saved duration to whoever shows up, rather than

having already promised to one firm or another. The intuition for the total duration constraint eventually binding, though, is that saving duration only makes sense if there is some chance of eventually awarding the saved duration. Otherwise the savings is wasteful, and the planner should allocate a little more duration right away. Since for any point there is a positive probability of reaching D , and D is an absorbing set, duration must eventually settle in D .

Corollary 4 *The ergodic distribution for duration is contained in D*

This result sets the stage for the next section, which details of the optimal policy once D has been achieved. We call this duopoly, in the sense that one of the two firms is receiving preference at any instant.

4 Complete Exclusion Rights with Duopoly

4.1 Dynamic Program

In this section we study the case where the planner has promised $1/r$ to the two innovators. In this case, a sufficient state variable is d , since $\underline{d} = 1/r - d$. Moreover, since every instant is allocated, $\underline{x} = 1 - x$. The planner's problem can therefore be written as

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

s.t.

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

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

4.2 Characterization

Since V is globally concave and symmetric, it is maximized at $1/2r$. This is intuitive: when duration promise is identical to the two agents, you can treat

the agents identically upon the next arrival, setting $d_1 = 1/r - d_2$, which is best since R is concave. Note that having the agents treated identically requires

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

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

Lemma 5 $V(d)$ is constant in the range of $[1/r - \hat{d}, \hat{d}]$

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

Lemma 6 $d_1(\hat{d}) > \hat{d}$

Proof. 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, for $d > \hat{d}$, the planner delivers the extra ε units of duration 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). The planner then receives.

$$\tilde{V}(\hat{d} + \varepsilon) = \begin{cases} \frac{\lambda}{r}R(\hat{d} + \delta) + \frac{\lambda}{r}R(\hat{d} - \delta) & \text{if } \varepsilon > 0 \\ \frac{2\lambda}{r}R(\hat{d}) & \text{otherwise} \end{cases}$$

This is clearly a differentiable function, and since it is feasible for the planner, must be less than the payoff V from the optimal policy. But therefore V is differentiable, and $d_1(\hat{d})$ must exceed \hat{d} . ■

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

Proposition 7 *For all $d < d^*$ and $\pi < 1$ there exists a T such that duration is greater than d with at least probability π .*

Note that when $d^* = 1/r$, this state is preserved for an arbitrarily long time as d is taken arbitrarily close to $1/r$; there is near monopolization.

Duration rises and falls with arrivals by the two firms; the two firms engage in a sort of "tug of war" for duration. An interesting feature is the evolution of x . Starting from duration in the middle region where V is maximized, ideas by innovator 1 move duration up, and the promise keeping constraint binds. As a result, $x = 1$. Note that the intervening period between innovations is never split; for any duration $d > \hat{d}$, a sequence of innovations by the "trailing" innovator promised $1/r - d$ falls with every innovation by the trailing firm, but if d is high enough, innovations by the trailing firm may at first not change x , so long as the duration promise falls but remains above \hat{d} . Preference for the trailing firm only kicks in when a sufficient number of innovations by it have moved duration below $1/r - \hat{d}$, i.e. to the flat portion plus one more innovation. 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 preference.

In the interpretation of preference as a patent right, the policy has the very unusual feature that for some histories for instance where $d_2(d) > \hat{d}$, even when firm 2 comes up with an innovation, firm one remains the preferred firm – gaining preference for all innovations, including the one of the competitor! It is as if firm one, the incumbent holder of the patent, gets the new patent at a license rate that extracts all of the surplus from firm two. In Hopenhayn, et al (2006), optimal policies were decentralized through buy-outs. Here the buyout is costless; firm one has such high duration promise d that it includes the right to buy the next innovation for free. The innovation by the "laggard" firm is entirely motivated by the eventual possibility that they might eventually enter the region where $d < \hat{d}$, become the patent holder for the frontier product, and earn the rents from the patent, including the past contributions that were "bought out".for free.

An alternative interpretation of preference is not as patent policy, 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 environment we study introduces a natural desire by the planner to backload rewards. If the planner waits to provide a given firm preference, it can provide that firm preference for more innovations, since innovations are constantly arriving, and a given firm can be allocated preference for multiple innovations at a point in time. As a result, the planner waits to award the preference until the firm has a preponderance of the duration promise.

Backloading of rewards is similar to the quantitative result in Acemoglu and Akcigit (2010). 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. Our model generates backloading for a different reason, and although our model does not nest the one used by Acemoglu and Akcigit (2010), it is similar enough that the same force is likely at work in their numerical results.

One might be concerned that complete exclusion rights are overly broad, in a way that might naturally lead to excessive monopoly power. For this

reason, as well as the fact that competition from prior preferred products invented by the other firm may be inevitable, we consider a restricted class of policies where complete exclusion is not available to the planner, and show that many of the same results are preserved.

5 Exclusive Preference Promises

5.1 Definition

In Hopenhayn et al. (2006), a patent system is defined to be exclusive if, at any given point in time, only one innovator has a duration promise. In other words, whenever a new idea is awarded some preference, all prior claims to market leadership by innovators other than the current one are set to zero. In the non-recurring context, that corresponds to the idea that any new innovation eliminates the ability of all prior firms to profit. Here, we take exclusivity to mean that arrival of a new idea supersedes the duration promise of prior firms. Since we assume that firms cannot profit from future promises made after this promise expires, we are assuming that innovations must be rewarded immediately, and cannot be rewarded after the next allocation of preference to the competitor. Exclusivity is natural, in that it mirrors the sort of market structures that are generated by many dynamic models of patents such as O'Donoghue, et al (1998). It maps into a patent right that is non-infringing on past rights, but does not encompass the prior rights. That is, every patent that is granted is non-infringing on every other patent, both before and after the arrival of a given improvement.

In particular, in the language of the quality ladder model of section 2, suppose that firms have the exclusive right to produce quality levels they "invent," but no right to exclude previous products invented by others. That is, new innovations get a patent right that neither infringes nor is infringed upon by the prior innovations. Until the other firm has a patented innovation, then, the leading innovator will be able to make profits equal to the difference between the leading edge quality and the last quality invented by the other firm. Let the expected discounted duration of this time be d . For this time period until a patented innovation arrives by the other firm, we say that the leading firm is given preferential treatment for that set of inventions, since his patent right allows profits to accrue. This preferential treatment ends (and preferential treatment for the other firm begins) when a patent

is granted to the other firm. Note that the other firm's patent *permanently* ends preferential treatment, since later patent rights will still not allow the firm to profit from innovations that took place before the competitors most recent innovation.

Formally, in terms of the program (3)

Definition 8 *Preference is exclusive if $d_2 = 0$ if $\underline{d}_1 > 0$, and $\underline{d}_2 = 0$ if $\underline{d}_1 > 0$.*

In the exclusive case, the planner has less duration to allocate, since preference can be given to a smaller number of innovations (by assumption) at any instant: only the most recent string of innovation by a given firm. Therefore there is, in this environment, a natural interpretation of $d = 1/(r + \lambda)$ that is similar to \hat{d} : all future innovations are implemented, meaning that duration for the current incumbent is defined by "until the next idea of the other firm arrives," which in discounted terms is $1/(r + \lambda)$. As a result, duration $d \leq 1/(r + \lambda)$ can be delivered without excluding anything, and there is scarcity in duration if $d^* > 1/(r + \lambda)$, which is implied by $d^* > \hat{d}$.

5.2 Dynamic Program

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

The dynamic program is therefore

$$\begin{aligned}
 rV(d) &= \max_{d_1, d_2, p} \left\{ \begin{array}{l} \lambda(R(d_1) + V(d_1) - V(d)) + \\ \lambda p(R(d_2) + V(d_2) - V(d)) \end{array} \right\} & (5) \\
 & \text{s.t.} \\
 rd &= 1 + \lambda(d_1 - d) - \lambda pd
 \end{aligned}$$

where d , d_1 , and d_2 can be taken to lie in $[1/(r+\lambda), 1/r]$, since if $d < 1/(r+\lambda)$, the PK constraint does not bind, and you can set $x < 1$. Therefore the value function is independent of d in this range, and we can restrict attention to the domain where all durations are in the range $[1/(r+\lambda), 1/r]$.⁸

The first order conditions for d_1 , d_2 , and p , respectively, are

$$\begin{aligned} R'(d_1) &= -V'(d_1) + \mu(d) \\ R'(d_2) &= -V'(d_2) \\ R(d_2) + V(d_2) - V(d) &= -\mu(d)d \end{aligned}$$

Since the FOC for p does not depend on p , we can take p to be always on the boundary, i.e. either zero or one. Next we use these facts to completely characterize the optimal policy.

5.3 Characterization

First, we show that if the current promise involves any exclusion ($d > 1/(r+\lambda)$), then the current offer to the incumbent firm if he arrives with an idea, $d_1(d)$, either implements the Arrow/Nordhaus static optimum d^* , or involves as little current exclusion as is consistent with promise keeping; that is, $p(d) = 1$, so that all arrivals by the non-incumbent firm are implemented at d .

Lemma 9 *Suppose $d > 1/(r+\lambda)$. Then either $d_1(d) \geq d^*$, or $p(d) = 1$.*

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

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

Denote by τ the smallest positive integer such that $p^\tau(d) > 0$. Since $d < d^* \leq 1/r$, it must be the case that τ is finite. In words, τ is the number of consecutive arrivals by the incumbent before $p > 0$. Note that following

⁸It 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.

the promise keeping constraint, duration is weakly falling with these arrivals, so $d_1^t(d) < d^*$ for $t \leq \tau$.

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

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

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

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

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

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

The characterization Lemma (9) shows the sense in which duration is backloaded: it is backloaded maximally, subject to not exceeding Arrow/Nordhaus.

The proof makes clear the reason for the backloading, which differs from standard theories of backloaded incentives. Since the planner is committed to d , he is committed to a fixed amount of exclusions of the non-incumbent firm. When those exclusions occur is welfare neutral, in the sense that all of the exclusions cost the planner missing out on a new incumbent starting with d_2 . When the planner implements duration d by excluding arrivals of the outside firm later, he raises the duration promise for all intervening arrivals by the incumbent, though, which raises the incumbents level of innovation. Intuitively, for a given d , every implemented non-incumbent is "bad luck" for the incumbent. The planner resolves this bad luck as soon as possible by making every "unfortunate" (for the incumbent) idea of the non-incumbent end the incumbents duration early on. Whenever the incumbent gets an arrival, then, the "good luck" for the incumbent is large: he has avoided a state where all his competitors ideas are implemented, and therefore gets the maximal duration increase $d_1(d)$ that the planner could have offered and maintained the promise of d . This makes the incumbent respond to an arrival of an idea with the maximal effort. Of course the planner does not drive $d_1(d)$ beyond d^* , since although doing so might be feasible, it would not be productive.

We know that duration starts at a point (d_2) that can not exceed d^* , since lowering d_2 is unambiguously better (in terms of current R and future payoff $V(d_2)$). The following shows that d_1 never goes above d^* , so duration promises must always lie in the range $[1/(r + \lambda), d^*]$. This is natural since $V(d)$ is decreasing; there is no incentive to promise any idea more than the Arrow/Nordhaus duration.

Lemma 10 *Suppose $d \leq d^*$. Then $d_1(d) \leq d^*$.*

Proof. Suppose $d_1(d) > d^*$. If you lower p and d_1 along (PK), (same exclusions, so gains from that cancel, and then you do better from tomorrow on) ■

These results together imply that $d_1(d)$ moves in the interval $[1/(r+\lambda), d^*]$ following

$$rd = 1 + \lambda d_1 - 2\lambda d \tag{6}$$

whenever such d_1 is less than d^* , and $d_1(d) = d^*$ otherwise. Following this evolution immediately gives us the following two implications of the dynamic evolution of duration.

Corollary 11 *Suppose $d > 1/(r + \lambda)$. Then $d_1(d) \geq d$, with d reaching d^* for a finite number of consecutive ideas for one firm*

This implies the following dynamics, starting from $d > 1/(r + \lambda)$: rising duration to d^* , and then, for $d^* < 1/r$, a constant probability of returning to d_2 . With positive probability the system reaches the Arrow/Nordhaus solution. If that static solution involves permanent preference for one firm, this implies that the system reaches monopolization in finite time:

Corollary 12 *Suppose $d_2 > 1/(r + \lambda)$. Then for any probability $\pi < 1$ there exists T such that d^* has been achieved in no more than T periods with probability π .*

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

Lemma 13 $d_2 > 1/(r + \lambda)$.

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

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

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

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

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

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

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

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

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

and therefore there is an improvement if

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

■

The Lemma shows that new incumbents are promised some exclusions of their competitors. The earlier results show that these exclusions are maximally backloaded, which generates duration promises that climb, with positive probability, to d^* . If $d^* = 1/r$, the planner monopolizes the market, in the sense of excluding all of one firm's ideas from preferential treatment, in finite time.

The intuition behind offering some exclusions to new incumbents is related to backloading. If the planner were forced to grant exclusions that generate d_2 immediately (i.e. for ideas that arrive immediately after the change in incumbency), then exclusions would not be beneficial and d_2 would be exactly $1/(r+\lambda)$. The reason is concavity of R : immediate exclusions cost $R(d_2)$ in un-implemented projects, but generate $R'(d_2)$. The latter is always smaller when R is strictly concave, but would be identical if R were linear. However, backloading leaves the cost of exclusions the same, but increases

their benefit: exclusions far in the future generate benefits for all of the incumbents ideas that arrive in sequence in the meantime. For linear R , this extra benefit of exclusions shows immediately that backloaded exclusions are beneficial; by continuity they must be beneficial for R that are nearly linear. The proof extends this logic to show that for any concave R , a small amount of exclusions, sufficiently backloaded, is beneficial to the planner.

6 Conclusions

We have characterized the solution to the problem facing a planner who must allocate preferential treatment across two firms who can use preferential treatment to make profits, and in turn are encouraged to innovate by the provision of preferential treatment. The planner, because he can allocate preference to a single firm for multiple innovations at any point in time, backloads rewards, giving the firm with the preponderance of the promise a preference that corresponds to a patent offering complete exclusion. We show that the promised duration of preference evolves in a region that is determined by the optimal promises in a static version of the model, which corresponds to a classic static patent problem of Arrow and Nordhaus. When there are no static distortions, so that the optimal static patent lasts forever, the optimal policy we study leads to monopoly, in the sense that one firm is excluded even though it is getting useful ideas. We show that these basic results hold even if the planner is forced to use a restricted set of policies where preference is always granted immediately for any innovation that is implemented.

References

- [1] Acemoglu, Daron, and Ufuk Akcigit (2010). "State Dependent Intellectual Property Rights Policy," NBER Working Paper W12775.
- [2] Aghion P, Harris C, Howitt P and Vickers J (2001) "Competition , Imitation and Growth with Step-by-Step Innovation," *Review of Economic Studies* **68**(3): 467-492.

- [3] K. J. Arrow, "Economic Welfare and the Allocation of Resources for Invention," in *The Rate and Direction of Inventive Activity*, Nat. Bur. Econ. Res. conference report, Princeton 1962, 619-25.
- [4] Becker, Gary S. and Stigler, George J. (1974) "Law Enforcement, Malfeasance, and Compensation of Enforcers," *The Journal of Legal Studies* **3**(1), pp. 1-18
- [5] Cornelli, Francesca and Mark Schankerman (1999). 'Patent Renewals and R&D Incentives,' *Rand Journal of Economics*, **30**: 197-213.
- [6] Hopenhayn, Hugo, Gerard Llobet, and Matthew Mitchell (2006), "Rewarding Sequential Innovators: Patents, Prizes, and Buyouts," *Journal of Political Economy* **114**(6), pp. 1041-1068
- [7] Hopenhayn, Hugo and Matthew Mitchell (2001) 'Innovation Variety and Patent Breadth,' *Rand Journal of Economics*, **32**: 152-166.
- [8] Lazear, E. (1981) "Agency, Earnings Profiles, Productivity, and Hours Restrictions," *The American Economic Review* **71**(4), pp. 606-620
- [9] W. D. Nordhaus, *Invention, Growth, and Welfare; A Theoretical Treatment of Technological Change*, Cambridge, Mass. 1969, ch. 5.
- [10] O'Donoghue, Ted, Suzanne Scotchmer and Jacques-François Thisse, "Patent Breadth, Patent Life, and the Pace of Technological Progress," *Journal of Economics & Management Strategy* **7**(1), 1998, pages 1-32.
- [11] Scotchmer, Suzanne (1999) 'On the Optimality of the Patent Renewal System,' *Rand Journal of Economics*, **30**: 181-96.