The Direction of Innovation

Kevin A. Bryan and Jorge Lemus*

April 23, 2016

Abstract

We construct a tractable general model of the direction of innovation. Competition leads firms to pursue inefficient research lines, because firms both race toward easy projects and do not fully appropriate the value of their inventions. This dual distortion will imply that any directionally efficient policy must condition on the properties of hypothetical inventions which are not discovered in equilibrium, hence common R&D policies like patents and prizes generate suboptimal innovation direction and may even generate lower welfare than laissez faire. We apply this theory to radical versus incremental innovation, patent pools, and the effect of trade on R&D.

It has long been conjectured that laissez faire markets will not produce the optimal quantity of innovation due to indivisibilities, where the fixed cost of R&D is only fully paid by the initial inventor, and inappropriability, where research generates spillovers on subsequent inventions (Arrow (1962)). Mechanisms like patents and R&D subsidies attempt to restore efficient research

*Bryan: Rotman School of Management, kevin.bryan@rotman.utoronto.ca. Lemus: University of Illinois at Urbana-Champaign, jalemus@illinois.edu. We appreciate the helpful comments of Jeffrey Ely, Benjamin Jones, William Rogerson, David Besanko, Mark Satterthwaite, Joel Mokyr, Bruno Strulovici, Weifeng Zhong, Gonzalo Cisternas, Dan Bernhardt, Ashish Arora, Erik Hovenkamp, Joshua Gans and Michael Whinston. Thanks to seminar audiences at Carnegie Mellon, Northwestern, Northwestern Kellogg, Universidad de Chile (CEA), MIT Sloan, Central Florida, Bilkent, Duke Fuqua, Cornell Johnson, Toronto Rotman, Oregon, RPI and conference participants at London Business School, the 2014 NBER Summer Institute, TOI (Chile) and Corvinus University of Budapest. Some results presented here previously circulated within a paper entitled "R&D Policy and the Direction of Innovation."

effort, and to do so without requiring the planner to know the research production functions of firms or the ex-ante expected value of their inventions. However, firms do not simply choose how much R&D to do, but also how to allocate their scientists across different research projects. For example, early semiconductor researchers could have worked with silicon or germanium, nuclear plants could have been developed using either water or deuterium as a moderator, and early automobile designers could have focused their effort on gasoline-powered, steam-powered or electric-powered vehicles. These potential inventions may differ in how hard they are to invent, in how valuable they are, and in which future research opportunities they make possible. A natural question therefore arises: how do policies intended to optimize the *quantity* of research affect the *direction* of that research?

We construct a theoretical model of innovation direction similar in spirit to existing workhorse models of innovation effort. Our model is tractable even though we permit firms to work on an arbitrary set of inventions at any time, with arbitrary links between inventions today and the nature of inventive opportunities in the future. We generate three primary theoretical results. First, even if the total quantity of research is optimal and even if firms receive the full social value of their inventions, two distinct classes of directional distortion remain, which we call "racing" and "underappropriation" distortions. Second, moving from laissez faire to a system with patents or subsidies can make these distortions strictly worse. Third, directional inefficiency is a property of *every* innovation policy which both rewards inventors and does not condition on the properties of inventions which are not invented in equilibrium. That is, the possibility of directional inefficiency places fundamental limits on the efficacy of decentralized "autopilot" innovation policy.

Intuition for the two main classes of directional distortion generated by our model can be seen in two simple examples. In Figure 1, there are two potential inventions, A and B. Two firms have one indivisible unit of research that can be costlessly allocated to either invention. Assume that inventors appropriate the full social value of their invention, and that once either invention is discovered, the marginal value of the other invention immediately falls to zero. That is, only the first invention has any value. Let A be relatively easy, such that if one firm researches A while the other researches B, A is discovered first with probability $\frac{3}{4}$. If firms both work on the same invention, they are equally likely to discover it first.

If both firms work on A, with probability one A is discovered and \$12 of

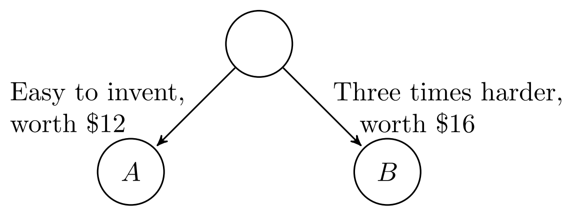

Figure 1: The racing distortion

value is created, and if both work on B, \$16 is created. If one firm works on A and the other on B, with probability $\frac{3}{4}$ A is invented first, and with probability $\frac{1}{4}$ B is invented first, creating a total value of $\frac{3}{4} \cdot \$12 + \frac{1}{4} \cdot \$16 = \$13$. The efficient solution involves both firms working on B, creating \$16 of value. However, working on B is not an equilibrium. A firm earns \$8 in expectation when both work on B, but it earns $\frac{3}{4} \cdot \$12 = \$9$ from deviating and working on A. The firm that deviates does not properly account for the fact that when it makes a discovery, the rival firm can no longer earn any surplus by discovering the now-worthless alternative invention. Notice that this is precisely the intuition of the *racing distortion* of classic patent race models like Loury (1979) in a directional context, substituting the extensive margin of which project to work on for the intensive margin of how hard to work, and the opportunity cost of foregone inventions for the cost of research effort.

Racing behavior is not the only way direction choice can induce inefficiency. In Figure 2, again let there be two firms allocating one indivisible unit of research each, and let there initially be two equally easy inventions A and B. Since they are equally easy, the probability a given firm invents first is $\frac{1}{2}$ regardless of what the other firm works on, hence there is no racing distortion. In addition, assume that once A is invented, it becomes possible for each firm to work on a third invention, C. Further, assume that once A is invented, the marginal value of B falls to zero, and once B is invented, the marginal value of both inventions A and C fall to zero.

The most social value, \$12, is created when both firms work on A and then on C. Each firm expects to earn \$6 under this research plan, but this is not an equilibrium. A firm that deviates by working on B instead of A will finish first with probability $\frac{1}{2}$, earning \$10. If A is invented before B, the deviating firm can still try to invent C at that point, earning $\frac{1}{2} \cdot \$8 = \$4$ in expectation. The expected payoff of the deviation is $\frac{1}{2} \cdot \$10 + \frac{1}{2} \cdot \$4 = \$7$, hence deviating is profitable. This example illustrates that

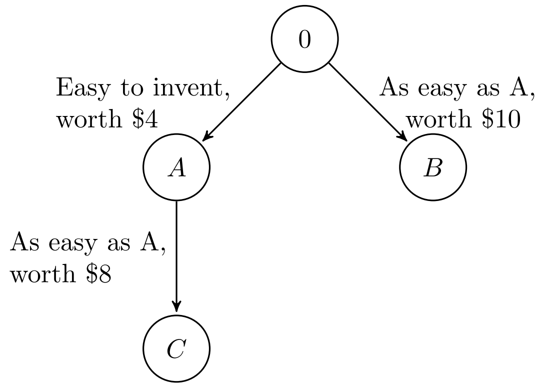

Figure 2: The underappropriation distortion

inventors do not properly account for how their inventive effort today affects the nature and availability of socially valuable projects other firms might invent in the future. This is precisely the intuition of the *underappropriation distortion* of sequential innovation models like Green and Scotchmer (1995) in a directional context, substituting distortion toward research lines where sequential inventions are relatively unimportant for inefficient effort along the intensive margin in a single sequential research line.

Our formal model will show that if we abstract away from known distortions in the market for R&D — for instance, if the total quantity of research is fixed at the socially optimal optimal level, if the research sector captures the full social value of their inventions, if researchers can be induced to work without agency problems, and if there is perfect knowledge among all firms about research opportunities today and in the future — then a competitive research sector will nonetheless be inefficient because of a combination of racing and underappropriation distortions. This dual distortion means that policies like patents and prizes will not necessarily improve efficiency, unlike in models where laissez faire research effort is inefficient.

Consider first patents, which let an inventor today capture the value of projects tomorrow which build on her invention. Patents cause firms to internalize the fact that their inventions today make future inventions easier or possible in the first place, but they do not fix, and may make worse, the racing distortion. Firms will be induced to race toward any invention which garners an industry-pivotal patent, regardless of whether that particular technology lies on a research line which is easy to productively extend. Prize contests are likewise problematic, as prizes for inventions exceeding a technological threshold exacerbate racing behavior toward lower-value projects which are just sufficient to garner the prize. Prizes given only to difficult technological achievements will push firms toward those types of projects, but if the optimal projects are easy yet avoided because of underappropriation, such a policy may simply make equilibrium directional distortion worse.

Note that patents and prize contests both condition inventor rewards solely on the properties of realized inventions, and not on properties of unrealized alternative research projects in the same technological area. Indeed, this is an important virtue of these types of policies: they can be run "automatically" by a planner who is ignorant of anything except ex-post observable features of inventions. However, when directional distortion is important, whether firms are deviating toward easy though potentially low-value inventions because of the racing distortion, or toward immediately lucrative yet potentially difficult inventions because of the underappropriation distortion, depends on the properties of all inventions including those which are not actually invented in equilibrium. Therefore, innovation policy which does not condition on the properties of those unrealized potential inventions, properties which may be very hard for the planner to observe, will be unable to restore directional efficiency across industries.

In the remainder of the paper, we develop the above intuition formally. In Section 1, we show how to construct planner-optimal and equilibrium dynamic research allocation for an arbitrary set of inventions with unrestricted linkages in how earlier inventions affect the value or difficulty of future inventions. In Section 2, we show that, in this general model, equilibrium directional inefficiency is driven by a combination of racing and underappropriation distortions qualitatively similar to those in the examples above. In Section 3, we show that laissez faire, patents of various strengths, and prizes cannot be ranked in terms of welfare, and that *every* policy which rewards inventors more than non-inventors and does not condition on the properties of off-equilibrium-path inventions cannot guarantee directional efficiency. We discuss four applications of the theory in Section 4. All proofs, and a number of generalizations, are left to the appendices.

Our results differ from the existing literature in restricting attention to the *distortion between the planner and the firms* when there are multiple projects available at any time, and success on a project changes the nature of research targets available in the future.[1] That is, we study inefficiency

[1]That firms may lack correct directional incentives in R&D is a longstanding worry,

in *research lines*. The distortions generated, and the relative advantages of different mitigating policies, do not depend in any way on information externalities (as in bandit models like Keller and Oldale (2003) and Chatterjee and Evans (2004)), changing preferences (Acemoglu (2011)), changing factor prices (Kennedy (1964), Samuelson (1965), Acemoglu (2002)), simultaneous discovery (Dasgupta and Maskin (1987)), heterogeneity across firms in size or internal organization (Holmstrom (1989), Aghion and Tirole (1994)) or differences in researcher desire for autonomy (Aghion, Dewatripont and Stein (2008)).[2] Our distortions arise even if the total quantity of research is optimal, and even if there is no gap between the social and private return to individual inventions.[3]

Directional inefficiency may be of particular importance due to limits on the ability of policy to affect the rate of innovation. Even when basic research has a high marginal return, both privately and socially, the return to government-sponsored R&D is often much more limited (e.g., David, Hall and Toole (2000), Lerner (1999)). This result is partially due to crowding out: the supply of trained scientists is essentially fixed in the short run (Goolsbee (1998)). If crowding out limits how planners can affect the *rate* of inventive activity, affecting R&D direction may be a first-order effect of government policy. Our results suggest fundamental limits on existing policy levers in

however. Nelson (1983) argues that "[i]t is not so much that private expenditures will be too little in the absence of government assistance. The difficulties lie rather in the fact that the market, left to itself, is unlikely to spawn an appropriate portfolio of projects."

[2]A discrete version of Acemoglu's 2011 result can be seen in our framework. He allows consumer preferences over technology lines to change. Patents are of finite length, so with some probability work on a line consumers do not value may not be worth anything until after the patent expires, and hence firms exert too much effort on lines where rents can be accrued in the near future. In our model, the possibility of preferences changing just feeds into the continuation value following some invention; firms undervalue social payoffs that accrue through the continuation value rather than immediately, since part of that continuation is captured by competitors.

[3]The most similar result we are aware of, though in the context of a structural endogenous growth model, is Akcigit, Hanley and Serrano-Velarde (2013), which suggests like our model that "neutral subsidies" like R&D tax credits operate by increasing the total amount of R&D without correcting the distortion toward applied research. A mathematically similar model to ours of direction choice applied to the context of competing patent races appears in a new working paper by Hopenhayn and Squintani (2016). In their model, firms in the analogue of our laissez faire equilibrium work first on relatively easy inventions even when working in this "hot" area is socially undesirable. The fundamental reason is very similar to the racing distortion in our model.

restoring directional efficiency.

1 The Model

Consider a finite set of states Ω, where $s \in \Omega$ represents a level of technology, or a collection of existing inventions. Transitions between states are associated with two parameters $\{\lambda, \pi\}$, where we define $\lambda : \Omega \times \Omega \to \mathbb{R}_+$ as the *simplicity* to transition between any two states, and $\pi : \Omega \times \Omega \to \mathbb{R}_+$ as the incremental *immediate social payoff* from a transition. When $\lambda(s, s') = 0$, state s' cannot be reached with one invention from state s. For each state $s \in \Omega$, the set $S(s) \subseteq \Omega$ represents all the states such that $\lambda(s, s') > 0$. We refer to the initial state of technology as s_0.

Definition 1. *An* invention graph *is represented by the triplet* $\{\Omega, \lambda, \pi\}$ *describing all potential states* Ω *of knowledge in a technological area, the simplicity* $\lambda : \Omega \times \Omega \to \mathbb{R}_+$ *of transitioning between them, and the immediate social payoff* $\pi : \Omega \times \Omega \to \mathbb{R}_+$ *from such a transition.*

The invention graph is common knowledge and all discoveries and inventive effort are publicly observed.[4]

As an example, consider the case of three inventions represented by the invention graph in Figure 3. In this example, the possible states of technology are given by $\Omega = \{s_0, \{1\}, \{2\}, \{1, 2\}, \{1, 3\}, \{1, 2, 3\}\}$. The transition and payoffs are given by the parameters $\{(\lambda_k, \pi_k)\}_{k=1}^{9}$. There is an arrow between states s and s' if and only if $\lambda(s, s') > 0$. In state s_0, only inventions 1 or 2 can be discovered in one step, and once either 1 or 2 have been discovered, research on invention 3 can begin.

[4]Relaxing the assumption that all future parameter values are commonly known to an assumption that only the distribution of parameter values in future states is known will not change our qualitative results. With perfect information, this assumption would merely change the expected continuation value following any invention, which will appear as an arbitrary parameter in Proposition 2. Indeed, if inventions today publicly resolve uncertainty about future parameter values, then inventions potentially create information useful to all parties, generating a positive externality. Thus, inventions which create a lot of useful information will be undersupplied by firms in equilibrium. We will not focus on these types of information transmission externalities in the remainder of this paper, as they are well-known from the multi-armed bandit literature (e.g., Keller, Rady and Cripps (2005)).

Our model allows state contingent payoffs and simplicities. In Figure 3, λ_2 and λ_4 are unrestricted, meaning that the discovery of invention 1 may increase ($\lambda_2 > \lambda_4$), decrease ($\lambda_2 < \lambda_4$), or keep constant ($\lambda_2 = \lambda_4$) the difficulty of discovering invention 2. Similarly, π_2 and π_4 can differ, capturing substitutability or complementarity between inventions 1 and 2.

In the remainder of the paper, we refer abstractly to states without reference to the exact bundle of inventions a particular state embodies, calling states $s' \in S(s)$ projects, research targets, or inventions.

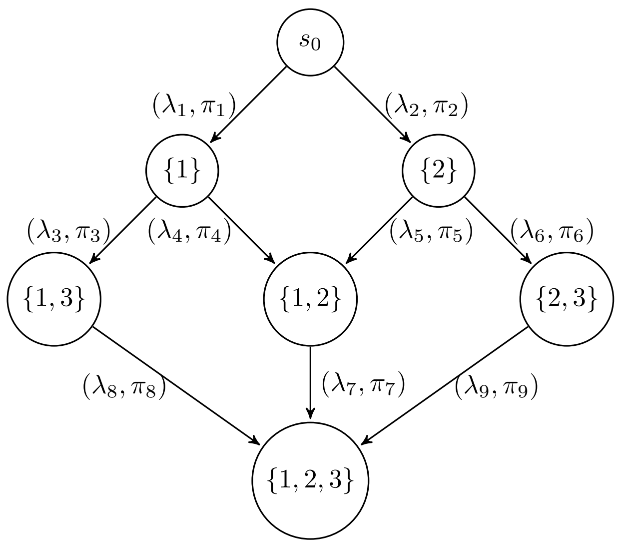

Figure 3: A simple invention graph.

1.1 Production Technology

There are N risk-neutral firms each endowed with $\frac{M}{N}$ units of research, where M represents the total measure of researchers in society. Let $x_i(s, s') \leq \frac{M}{N}$ be the flow amount of research allocated by firm i toward state $s' \in S(s)$ when the current state is s, and let $x(s, s') = \sum_i x_i(s, s')$ be the aggregate flow amount of research toward state s'.[5] Research is costless, so the problem

[5] Although time is continuous in our model, optimal and equilibrium strategies will be constant between state transitions, hence we omit time subscripts. Intuitively, no information is revealed and no changes in the strategy set or payoffs occur between state transitions.

is one of pure allocation of research resources.

As in patent race models, the probability of discovering s' given $x_i(s,s')$ in a given interval of time is determined by the exponential distribution, with hazard rates $\lambda(s,s')x_i(s,s')$ linear in effort, independent across firms, and independent across research lines within any firm.[6] Therefore, the unconditional probability of a transition from s to s' in an interval of time τ is given by $1-\exp(-\lambda(s,s')x(s,s')\tau)$.

In the remainder of the paper, we will omit some indexes for ease of notation, denoting $x_i(s,s')$ simply as $x_{is'}$, and likewise for similar variables, when it is clear that $s' \in S(s)$.

1.2 Planner Problem

Since the invention hazard rates are linear and independent across firms, a risk-neutral social planner needs only decide how to allocate all M units of research across projects. The expected discounted value of the invention graph for the planner at state $s \in \Omega$ is defined recursively as

$$V_{ps} = \max_{\substack{\sum_{s'\in S(s)} x_{s'}=M, \\ x_{s'}\geq 0,\ \forall s'\in S(s)}} \int_0^\infty e^{-rt} e^{-\sum\limits_{s'\in S(s)} \lambda_{s'}x_{s'}\cdot t} \sum_{s'\in S(s)} \lambda_{s'}x_{s'}\cdot(\pi_{s'}+V_{ps'})dt$$

That is, after reaching state s, the planner chooses the allocation $x = (x_{s'})_{\{s'\in S(s)\}}$ to maximize the future discounted payoff: the integral with respect to time of the probability that no invention has occurred ($e^{-\sum_{s'\in S(s)}\lambda_{s'}x_{s'}t}$), times the immediate hazard rate of each research line ($\lambda_{s'}x_{s'}$), times the discounted (e^{-rt}) payoff summed over all possible inventions inclusive of continuation value from a discovery along that line ($\pi_{s'}+V_{ps'}$). Simplifying the expression above, the social planner problem is to solve the recursive maximization problem:

$$V_{ps} = \max_{\substack{\sum_{s'\in S(s)} x_{s'}=1, \\ x_{s'}\geq 0,\ \forall s'\in S(s)}} \frac{\sum\limits_{s'\in S(s)} M\lambda_{s'}x_{s'}[\pi_{s'}+V_{ps'}]}{r+\sum\limits_{s'\in S(s)} M\lambda_{s'}x_{s'}}$$

[6] In the Online Appendix, we generalize to a hazard rate that is concave or convex in $x_i(s,s')$. There are technical difficulties with this objective function (in particular, non-pseudoconcavity) which do not appear, for example, in one-shot models like Reinganum (1981), but our main qualitative results do not change.

1.3 Firm Problem

Let policy $\mathcal{P}$ determine the transfers received by firms following an invention. In particular, $\mathcal{P}$ specifies transfer $w(s,s')$ received by inventors and transfer $z(s,s')$ received by noninventors following any invention. At each state $s \in \Omega$, the equilibrium continuation value for firm i following a transition to $s' \in S(s)$ given policy $\mathcal{P}$ is denoted $V_{i\mathcal{P}s'}$.[7] The firm problem is to allocate its $\frac{M}{N}$ units of research among the projects $s' \in S(s)$, conditional on other firms' allocations and the policy rule. Once any firm discovers some invention $s' \in S$, all firms reallocate effort across the new set of potential research targets $S(s')$.

Denote $a_{-is'} = \sum_{j \neq i} x_{js'}$ as the total research allocated towards invention s' by firms other than i. Given the strategies of rivals $a_{-i} = (a_{-is'})_{s' \in S(s)}$ and the policy rule $\mathcal{P}$, the expected discounted value of firm i at state s can be written recursively as

$$V_{i\mathcal{P}s|a} = \max_{\substack{\sum_{s' \in S(s)} x_{is'} = \frac{M}{N}, \\ x_{is'} \geq 0,\ \forall s' \in S(s)}} \int_0^\infty e^{-rt - \sum_{s' \in S(s)} (a_{-is'} + x_{is'})\lambda_{s'} t} \sum_{s' \in S(s)} \lambda_{s'} [x_{is'}(w_{s'} + V_{i\mathcal{P}s'}) + a_{-is'}(z_{s'} + V_{i\mathcal{P}s'})] dt$$

The strategy of each player depends on the current state, the equilibrium continuation value, and the current allocation of effort of rival firms. Simplifying this expression, firms solve:

$$V_{i\mathcal{P}s|a} = \max_{\substack{\sum_{s'} x_{is'} = \frac{M}{N}, \\ x_{is'} \geq 0,\ \forall s' \in S(s)}} \frac{\sum_{s' \in S(s)} \lambda_{s'} [x_{is'}(w_{s'} + V_{i\mathcal{P}s'}) + a_{-is'}(z_{s'} + V_{i\mathcal{P}s'})]}{r + \sum_{s' \in S(s)} \lambda_{s'} (x_{is'} + a_{-is'})}$$

1.4 Common Transfer Policies

We permit general transfer policies $\mathcal{P}$, but four classes will be of special relevance: laissez faire, patents, neutral prizes, and information-constrained policies. Laissez faire gives inventors the immediate social payoff of their invention, but permits all firms to equally build on that invention.

Definition 2. *The transfer policy $\mathcal{P}_{LF}$ is* laissez faire *if inventing firms receive the full immediate social payoff of their invention, i.e.* $w(s,s') =$

[7] Forcing the continuation value to be identical for inventing and noninventing firms is without loss of generality since w and z are unrestricted.

$\pi(s, s')$, and if noninventing firms receive no immediate transfer following an invention but are equally able to build on today's invention ($z(s, s') = 0$).[8]

We model patents as a tractable reduced form of a licensing game. Let parameter $\gamma \in [0, 1]$ indicates what fraction of the total continuation value following any invention non-inventors have to cumulatively pay to the inventor.[9] If $\gamma = 1$, patents are so strong that the inventor of s' is immediately granted the entire discounted surplus generated by their invention including surplus from any invention which builds on it in the future. If $\gamma = 0$, patents are equivalent to laissez faire.

Definition 3. *The transfer policy $\mathcal{P}_\gamma$ involves* patents *if inventing firms receive transfers $w(s, s') = \pi(s, s') + (N - 1)\gamma V_{i\mathcal{P}_\gamma s'}$ and noninventors pay (receive a negative transfer) $z(s, s') = -\gamma V_{i\mathcal{P}_\gamma s'}$, for $\gamma \in [0, 1]$.*

Prizes and contests are another common innovation inducement scheme. Let a neutral prize q in state s be a lump sum awarded to the inventor of any project $s' \in S(s)$, in addition to the immediate payoff $\pi_{s'}$ and continuation value $V_{i\mathcal{P}s'}$. Many real-world prizes share this structure, where any invention achieving a given technological threshold is rewarded with a constant prize. For example, the Netflix contest awarded \$1M to the firm that "substantially improves the accuracy of predictions about how much someone is going to enjoy a movie based on their movie preferences." If q is interpreted more broadly as a form of utility for an inventor, then it may also represent credit in the Mertonian sense; merely passing a technological threshold, regardless of economic significance, is the cutoff upon which scientific credit is distributed.

Definition 4. *The transfer policy $\mathcal{P}_q$ involves* neutral prizes *in state s if the first firm to successfully invent any invention in state s receives transfers $w(s, s') = \pi(s, s') + q$ and noninventors receive transfer $z(s, s') = 0$.*

Patents, neutral prizes and laissez faire are all policies which do not condition transfers $w(s, s')$ and $z(s, s')$ on off-equilibrium-path parameters: the

[8] The assumption that firms earn the full social surplus under laissez faire from each invention is not critical for our results. Under the laissez faire policy, since π enters the firm value function linearly, a straightforward induction argument shows that if all immediate firm payoffs π are scaled by $\sigma > 0$, the firm problem is unchanged. Only the relative values of immediate payoffs affect the choice of direction.

[9] Specifying patent payments in this way allows us to retain the earlier assumption that, following these side payments, all firms receive equal continuation values.

transfer following the invention s' does not depend on the parameters of projects $\ell \neq s' \in S(s)$. In this sense, these policies all lie in a class we call *information-constrained*, meaning any policy where transfers following an invention do not condition on the parameters of inventions which are not invented in equilibrium. An example of a policy which is not information-constrained is an NIH funding panel, which explicitly takes into account the value and challenge of alternative projects when choosing which projects should receive funding.

Definition 5. *A transfer policy $\mathcal{P}$ is* information-constrained *if transfers $w(s, s')$ and $z(s, s')$ do not condition on the parameters of inventions $\ell \neq s'$.*

2 Planner Optimum and Firm Equilibrium

The planner and firm problems both involve choosing vectors of effort across an arbitrarily large number of projects. In principle, then, comparing the efficiency of the optimal and equilibrium allocations involves comparing the welfare induced by two arbitrary vectors describing aggregate effort in every state. Worse, neither the planner maximand nor firm equilibrium have analytically tractable first order conditions, hence a mathematical workaround is required to make non-numerical statements about efficient and inefficient directional policy.

Although the objective function for the planner and the firms is non-linear in the research allocation, both problems are linear fractionals with linear inequality constraints. The Charnes-Cooper transformation (Charnes and Cooper (1962)) converts programs of this type into analytically tractable linear programs. These linear programs have corner solutions, where full effort is optimally exerted on a single project at a single time, and the corner solutions can be stated in terms of the maxima of simple indices. The problem of directional efficiency is therefore tractable. In any state, the planner will direct all researchers to work on a single project, and knowing which project is optimal simply involves checking which invention maximizes a particular index derived from the linear program solution. A firm equilibrium may involve firms all working on the same project or on different projects, but critically the existence of a socially optimal firm equilibrium requires only knowing whether there exists an equilibrium where all firms exert full effort on the planner's preferred project. This can again be checked using a similar simple index.

Define $\tilde{r} = Nr + M(N-1)\lambda_{s'}$, which corresponds to a virtual discount rate for competing firms, and denote by Θ the expected discounted continuation value for firm i when firm i does not exert any effort until the next invention is completed by some firm:

$$\Theta = \frac{\sum_{s' \in S} a_{-is'} \lambda_{s'} (z_{s'} + V_{i\mathcal{P}s'})}{r + \sum_{s' \in S(s)} a_{-is'} \lambda_{s'}}.$$

Denote the immediate payoff plus the continuation value as $P_{i\mathcal{P}s'} = w_{s'} + V_{i\mathcal{P}s'}$ for an inventing firm i under policy $\mathcal{P}$ and $P_{ps'} = \pi_{s'} + V_{ps'}$ for the planner. We refer to $\lambda_{s'}\pi_{s'}$ as the flow immediate social payoff, $\lambda_{s'}V_{ps'}$ as the flow social continuation value, $\lambda_{s'}P_{ps'}$ as the flow total social payoff, and the time between any two inventions as a "period."

Proposition 1. *In state $s \in \Omega$:*

1. *The planner optimum puts all research effort toward states $s' \in S(s)$ which maximize the index*

$$\frac{M\lambda_{s'}}{r + M\lambda_{s'}} P_{ps'}$$

2. *The best response of firm i given rival effort a_{-i} and policy $\mathcal{P}$ is to distribute all of its effort among states $s' \in S(s)$ which maximize the index*

$$\frac{M\lambda_{s'}}{\tilde{r} + M\lambda_{s'}} (P_{i\mathcal{P}s'} - \Theta)$$

The firm's best response index differs from the planner's in three ways: transfers to firms inclusive of continuation value may differ from the total social payoff ($P_{ps'} \neq P_{i\mathcal{P}s'}$); firms maximize payoffs only marginal to the Θ which is earned from doing nothing in the current state; and firms effectively discount at a different rate from the planner since their research decision today only has a partial effect on the eventual time the next invention in society is completed, and hence the time at which all firms can begin work on new projects ($r \neq \tilde{r}$).

The planner index generically has a unique maximum. It may seem surprising that the planner does not mix across projects, but generating an underlying reason for R&D diversity is a tricky modeling problem, frequently misunderstood in informal discussion. The intuition that a diverse research agenda provides "more lottery tickets" is false in the absence of decreasing returns to scale in the research production function, since simultaneous diversified research with any constant returns to scale can be replicated by sequentially exploring projects, with the benefit that sequential exploration allows a high level of effort to be exerted first on projects which are believed to be more valuable. It can be shown in a variant of our main model that uncertainty about parameter values, switching costs, asymmetry across firms, or even low levels of decreasing returns to scale do not necessarily lead a planner to work on multiple projects at a time.[10]

2.1 Efficiency of the Firm Equilibrium

The planner optimum generically involves full effort in each state on a single invention s'. From the best response characterization in Proposition 1, we can construct firm equilibria, calling them efficient if there exists any equilibrium where firms allocate their research in the planner optimal way. We restrict to Markov perfect stationary equilibria throughout to rule out equilibria where firms collude and punish each other across states.[11]

Definition 6. *The firm equilibria in state s are* efficient under policy $\mathcal{P}$ *if there exists a Markov perfect stationary equilibrium involving full effort from all firms toward planner optimal $s' \in S(s)$.*

In the next proposition, we characterize the trade-off faced by firm when deviating from the efficient research path. First, define

$$\Delta(s', \ell) = \frac{M(\lambda_\ell - \lambda_{s'})}{r + M\lambda_{s'}} = \frac{\dfrac{1}{r + M\lambda_{s'}} - \dfrac{1}{r + M\lambda_\ell}}{\dfrac{1}{r + M\lambda_\ell}}$$

which measures the relative difference of the value of \$1 forever with discount rates indexed by $\lambda_{s'}$ and λ_ℓ.

Proposition 2. *Let the current state be s.*

[10] See Online Appendices B and C for details.

[11] The Online Appendix discusses the existence, multiplicity, and potential asymmetry of firm equilibria.

1. *Project s' is planner optimal if and only if, for all $\ell \in S(s)$*

$$\lambda_{s'} P_{s'} \geq \lambda_\ell P_\ell - \lambda_{s'} P_{s'} \Delta(s', \ell)$$

2. *Project s' is a firm equilibrium if and only if, for all $\ell \in S(s)$*

$$\lambda_{s'} P_{s'} \geq \lambda_\ell P_\ell - \lambda_{s'} P_{s'} \Delta(s', \ell) + D_1(s', \ell) + D_2(s', \ell) + D_3(s', \ell),$$

where

$$\begin{aligned} D_1(s', \ell) &= \lambda_{s'}(P_{ps'} - (w_{s'} + V_{i\mathcal{P}s'})) - \lambda_\ell(P_{p\ell} - (w_\ell + V_{i\mathcal{P}\ell})) \\ D_2(s', \ell) &= \left(\frac{N-1}{N}\right) \Delta(s', \ell) \lambda_{s'} P_{ps'} \\ D_3(s', \ell) &= \frac{1}{N} \Delta(s', \ell) \lambda_{s'} (P_{ps'} - (w_{s'} + (N-1) z_{s'} + N V_{i\mathcal{P}s'})) \end{aligned}$$

The first part of Proposition 2 simply restates the planner optimum in Proposition 1 in terms of discounted flow payoffs. The second part of Proposition 2 fully decomposes the source of inefficiency in the firm equilibrium into three parts.

$D_1(s', \ell)$, the *underappropriation distortion*, is positive when a firm deviating from project s' to project ℓ receives a higher portion of the total social value of the invention. $D_2(s', \ell)$, the *racing distortion*, captures the incentive to deviate toward easier projects because firms do not account for how their effort affects the probability other firms succeed with alternative projects in a given period of time. $D_3(s', \ell)$, the *industry payoff distortion*, is zero when the total payoff to all firms under policy $\mathcal{P}$ is equal to the total social payoff of each invention. When that condition does not hold, the racing externality is either minimized or exacerbated.[12] If s' is the planner optimum, any policy $\mathcal{P}$ such that $D_1(s', \ell) + D_2(s', \ell) + D_3(s', \ell) \leq 0$, for all $\ell \in S(s)$, implements the efficient direction.

The decomposition in Proposition 2 is perhaps surprising. Research by firms affects what projects other firms can work on tomorrow, when these future projects become available, the probability a given firm actually invents the project it is currently working on, and so on, and yet *any* innovation policy can generate inefficiency in only three ways: either firms are overincentivized to race toward projects easier than the planner preferred ones; or

[12] Recall that $\Delta(s', \ell) > 0$ when invention ℓ is easier than s'.

inventing firms do not appropriate a sufficiently large share of the surplus their inventions generate; or researching firms overall receive a different share of the social surplus of invention depending on which research lines are pursued. These are the fundamental ways competition in the resaerch sector can generate directional efficiency, as they can exist even when the total aggregate amount of research is fixed at the (unmodeled) socially optimal level, and even though our model deliberately shuts down any distortions which lead to inefficiency with a single private firm.[13]

Consider, for example, the laissez faire policy, which generates transfers to the inventor $w(s, s') = \pi(s, s')$ and to noninventors $z(s, s') = 0$. If the firm equilibrium in all future states is efficient, then by induction the firm continuation value under laissez faire is $V_{i\mathcal{P}s'} = \frac{V_{ps'}}{N}$; each firm collects, in expectation, an equal share of the social continuation value. In this case, $D_3 = 0$ because total industry transfers are exactly the total social payoff. Since inventors only receive a $\frac{1}{N}$ share of the social continuation value, and firms are overincentivized to work on relatively easy projects, the underappropriation distortion D_1 and racing distortion D_2 distort behavior.

Corollary 1. *When firms research efficiently in all future states, under the laissez faire innovation policy $\mathcal{P}_{LF}$:*

$$D_1(s', \ell) = \left(\frac{N-1}{N}\right)(\lambda_{s'} V_{s'} - \lambda_\ell V_\ell),\ D_2(s', \ell) = \left(\frac{N-1}{N}\right)\lambda_{s'}\Delta(s', \ell)P_{ps'},\ D_3(s', \ell) = 0$$

The firm equilibrium condition in Proposition 2 collapses to:

$$\lambda_{s'} P_{s'} \geq \lambda_\ell P_\ell - \lambda_{s'} P_{s'} \Delta(s', \ell) + (N-1)(\lambda_\ell \pi_\ell - \lambda_{s'} \pi_{s'}).$$

Corollary 1 says that under laissez faire, firms are incentivized to deviate toward projects with high immediate flow payoffs $\lambda\pi$. These projects may be easier than the planner optimum ($\lambda_\ell > \lambda_{s'}$), or have a higher immediate payoff ($\pi_\ell > \pi_{s'}$), or both. The magnitude of the distortion is increasing in N, and hence Proposition 3 shows that sufficient fragmentation of the research sector guarantees inefficiency unless the planner optimal project in a given state has higher flow immediate payoff than any potential deviation.

Proposition 3. *Let the firm equilibrium in future states be efficient.*

[13] It is trivial to note that when $N = 1$, invention is always directionally efficient, as a single private sector firm in our model will behave identically to a social planner.

1. *If planner-optimal s' is not a laissez faire equilibrium when there are $\bar{N}$ firms, s' is still not an equilibrium for any $N \geq \bar{N}$.*
2. *If the planner optimal invention does not maximize $\lambda_{\bar{s}}\pi_{\bar{s}}, \forall \bar{s} \in S(s)$, then there exists a level N^* of fragmentation in the research sector such that if $N \geq N^*$, the laissez faire firm equilibrium is inefficient.*

Not every type of invention graph leads to laissez faire inefficiency. In the Online Appendix, we prove that directional inefficiency requires both the existence of multiple research targets in some states, and some form of state dependence linking invention today to inventive opportunities tomorrow. If there are multiple research targets, but invention today does not change the social value or simplicity of the remaining targets, the laissez faire equilibrium is efficient. Since there is no benefit in continuation value from avoiding projects with high flow immediate payoff, and the future is discounted, the planner will work first on projects with maximal $\lambda\pi$, hence by Corollary 1 the firms will not deviate. However, once a single element of state dependence is introduced - for example, an invention which requires a precursor, or an invention which is made easier by complementary inventions, or an invention whose value is reduced once a substitute exists - the laissez faire firm equilibrium is no longer efficient in general.

3 General Policy Solutions to Directional Inefficiency

In the previous section, we saw that laissez faire is not directionally efficient. In this section, we show that prize contests, and patents of various strengths, also do not induce efficient equilibria, that patents, prizes and laissez faire can each be more efficient than the others, and that any efficient invention graph *must* either condition on off-equilibrium-path parameters of the invention graph or reward noninventors as highly as inventors. That is, information-constrained policies, whose transfers are simple functions of on-equilibrium-path observables, are insufficient when it comes to directional efficiency.

3.1 Neutral Prizes

Recall that the neutral prize policy $\mathcal{P}_q$ awards a lump sum q to the inventor of any project $s' \in S(s)$

Corollary 2. *In state s, under neutral prize policy $\mathcal{P}_q$, total distortions are*

$$D^*_{NP}(s',\ell) = D^*_{LF}(s',\ell) + \frac{\Delta(s',\ell)}{N} q\tilde{r}$$

where $D^*_{LF}(s',\ell)$ is the equilibrium distortion in Proposition 2 under laissez faire.

Neutral prizes of any size do not guarantee efficiency, and indeed can generate an inefficient outcome even when laissez faire is efficient. Note that neutral prizes still generate the underappropriation distortion of laissez faire, since firms only collect a portion of the social continuation value of their inventions. In addition, since $\Delta(s',\ell) > 0$ only for projects ℓ that are easier than the planner optimal project s', neutral prizes exacerbate the racing distortion toward projects that are easier than the planner optimum. Intuitively, a fixed prize q increases the payoff, in percentage terms, of low-value projects more than high-value projects. Therefore, prizes can only make firms more likely to work on lower-value yet easier projects than the laissez faire equilibrium, as they race to finish inventions which are easy yet just sufficient to garner the prize.

In practice, then, large prizes will cause firms to race toward inventions which can be completed more quickly because the incentive from winning the prize overwhelms the incentive of developing a potentially more difficult technology that is easier for the inventor to build on in the future. If q represents the value to an inventor of Mertonian credit for a breakthrough, the exact same distortion arises. However, if the prize designer believes that firms are working on inefficiently hard projects because the inventor of those projects can capture a large share of the total value of a research line—e.g., projects that can be kept secret—then prizes may be effective in reducing that type of directional inefficiency. These directional distortions must be traded off against whatever increase in effort, unmodeled here, a prize designer hopes to generate toward research in a given technological area.

3.2 Patents

Patents are thought to play an important role when sequential innovation is critical, since patent rights limit double marginalization when inventions build on each other (Green and Scotchmer (1995)). When multiple research lines can be pursued, however, patents can distort *ex-ante* incentives even when they ameliorate ex-post double marginalization problems. The grant of a broad patent which covers substitutes and downstream inventions can cause a race among upstream inventors to develop relatively easy yet socially inefficient early-stage inventions in order to obtain this broad patent.

Recall that under the patent policy $\mathcal{P}_\gamma$, inventors receive transfers $w(s,s') = \pi(s,s') + (N-1)\gamma V_{i\mathcal{P}s'}$ and noninventors pay $z(s,s') = -\gamma V_{i\mathcal{P}s'}$, where $\gamma \in [0,1]$ represents the fraction of the continuation value following a patented invention which is collected by the initial inventor.

Corollary 3. *Under patent policy $\mathcal{P}_\gamma$, distortions are*

$$D^*_{P_\gamma}(s',\ell) = D^*_{LF}(s',\ell) + \gamma(N-1)(\lambda_\ell V_{\mathcal{P}_\gamma \ell} - \lambda_{s'} V_{i\mathcal{P}_\gamma s'}) + \mathcal{V}(\gamma),$$

where

$$\mathcal{V}(\gamma) = (1 + \Delta(s',\ell))\lambda_{s'}(V_{i\mathcal{P}_{LF}s'} - V_{i\mathcal{P}_\gamma s'}) - \lambda_\ell(V_{i\mathcal{P}_{LF}\ell} - V_{i\mathcal{P}_\gamma \ell})$$

Suppose that invention in all future states is efficient ($V_{i\mathcal{P}_{LF}s} = V_{i\mathcal{P}_\gamma s} = \frac{V_{ps}}{N}, \forall s$), hence $\mathcal{V}(\gamma) = 0$. In this case, patents of maximal strength $\gamma = 1$ exactly cancel out the laissez faire underappropriation distortion $D_1(s',\ell)$, as might be expected. Patents, however, do not affect the laissez faire racing distortion $D_3(s',\bar{s})$. If the underappropriation distortion under laissez faire is helping counteract the racing distortion—e.g., if firms are not deviating toward an inefficient easy project under laissez faire because they would only capture a small portion of the total social value of that invention—then increasing the strength of patents can actually make directional inefficiency *worse*. In practical terms, with strong patents, firms may avoid hard inventions with large payoffs because racing to invent something easier gives them a claim over the value of future discoveries, including some substitutes which may have been potential research targets from the start.

An immediate implication of the previous two corollaries is that there exist invention graphs for which patents of various strengths, neutral prizes, and laissez-faire each dominate the others in terms of social welfare. We give numerical examples in Online Appendix B.

3.3 Directional Efficiency With Information Constrained Policy

Patents and prizes both condition on very little information: the incentives they provide to firms depend only on the parameters of on-equilibrium-path inventions, and hence can operate "automatically." Both policies were shown to have difficulty simultaneously limiting racing behavior while still causing firms to care sufficiently much about the social continuation value their research generates. This finding raises the question: does there exist *any* policy which can generate directional inefficiency without conditioning on the parameters of off-path inventions?

Recall from Definition 5, in section 1.3, that a policy is information-constrained if transfers do not condition on the parameters of inventions which are not ever invented in equilibrium. We show in Proposition 4 that any information-constrained policy which is always efficient must reward inventing firms at least as much as non-inventing firms. That is to say, if inventors are to be rewarded more than non-inventors, the dual nature of directional distortions, coming both from racing behavior and underappropriation, cannot be wholly corrected by *any* information-constrained policy.

Proposition 4. *Let transfer policy $\mathcal{P}$ be information-constrained.*

1. *If the payoff of inventors and non-inventors can be equalized, the information-constrained policy $w(s,s') = z(s,s') = \alpha\pi(s,s'), \forall\alpha \geq 0$ implements efficiency on any invention graph.*

2. *If the payoff (inclusive of continuation value) for inventors must be strictly higher than that of non-inventors, then there exists no information-constrained policy which is efficient for all invention graphs.*

The first part of Proposition 4 is trivial: since we have shut down all nondirectional distortions in our model, if the payoff to inventors and non-inventors is identical, there is no benefit for any firm from taking any action that lowers the cumulative payoff to all firms. If the cumulative payoff to all firms is maximized along the efficient path, then direction will not be distorted in equilibrium.[14] For many reasons aside from directional efficiency, however, we may wish to rule out policies that reward inventors and non-inventors equally, the most obvious one being that getting the total *rate* of effort to the optimal level may require rewarding inventors in some way.

The second part of the Proposition 4 shows that any efficient policy which gives larger rewards to inventors than non-inventors must condition on the

[14] Other efficient policies like "paying firms only if they invent along the socially optimal research line" require the mechanism to condition on the full vector of simplicities in order to compute the optimal direction.

parameters of off-equilibrium-path inventions. Therefore, the result should be read as an impossibility result concerning uniform, non-targeted policies, which condition on ex-post public features of an invention to generate correct directional incentives, rather than as an implementation result in the formal sense of the term in mechanism design.

The intuition of Proposition 4, as can be seen in the formal proof, is as follows. Consider *any* invention graph with states where the planner is indifferent between two inventions, and where the indifference may result because an invention s' is harder or easier than an alternative ℓ. Since $w > z$ by assumption, if the relative total transfer to inventors of two inventions is equal to the relative total social payoffs, so that there is no underappropriation, then firms will race toward easier inventions. An arbitrarily small change in the simplicity or payoff of the off-equilibrium-path invention will make the planner optimum unique, and hence the proposed transfer policy will be inefficient. On the other hand, whether transfers should be biased toward s' or ℓ depends on which one is easier and inducing racing as a result, a comparative question that necessarily depends on the simplicity of the off-equilibrium-path invention. This result does not require any unusual invention parameters or interactions: all that is required is for the planner to be potentially unsure in some states whether the racing and underappropriation distortion is dominant.[15]

Empirically, many common innovation policies are information-constrained, and hence the inefficiency result in Part 2 binds. Governments appear to desire "neutral policies" since they impose less cost in terms of information gathering and less scope for political considerations to factor into the reward system for inventors. Whatever the reason, our result suggests that avoiding targeted policies has a cost in terms of efficiency. Note that our definition of "information-constrained" allows more planner information than is usually assumed when justifying policies like patents, since we permit conditioning

[15] The reader may wonder whether certain types of invention graphs are efficient under laissez faire or under a particular policy: that is, is the inefficiency result driven by "weird" invention graphs, or is directional inefficiency a generic property? It is trivial to construct classes of invention graphs where simple policies generate efficiency. For example, if in every state there is one invention that is both easier and more immediately lucrative than any other, laissez faire will be efficient. However, both the nature of the inefficiency proof, and results proved in the Online Appendix, suggest to us that directional inefficiency occurs broadly. In particular, a single element of state dependence, where the nature of inventive opportunity tomorrow depends on what is invented (or not invented) today, is sufficient to construct graphs where laissez faire policy is inefficient.

policy on the full social value, including continuation value, of all inventions when setting transfers.

4 Applications and Implications

Though our primary results are theoretical, the model is general enough to apply to a broad set of policy problems, of which we consider four. First, competition among innovators can cause firms to work on projects that are *either* more radical or more incremental than those preferred by the social planner. Second, patent pools or research joint ventures can play a role in reducing directional distortion in addition to their role in limiting holdup which is well-known in the existing literature. Third, when trade expansions enlarge the size of the market, the endogenous entry of additional firms will distort directional incentives. Fourth, the breadth of "pioneer patents", often determined ex-post in the courts, ought to account for the potential of ex-ante distortion.

4.1 Incremental Steps versus Large Steps

Proposition 5 shows that, perhaps counterintuitively, firms in competitive equilibrium may work on inventions that are either too incremental *or* too radical.

Proposition 5. *Let an invention graph contain an incremental line with two sequential inventions 1 and 3, and a radical invention with a single invention 2. Assume that the radical project is harder than either of the incremental steps ($\lambda_2 > \max\{\lambda_1, \lambda_3\}$), that the radical invention payoff exceeds the total payoff of the incremental line ($\pi_2 > \pi_1 + \pi_3$), and that once either the radical invention or the incremental line have been invented, the value of the other line falls to zero.*

1. *If the planner is indifferent between the incremental line and the radical line then the incremental line (radical line) is a laissez faire equilibrium if and only if $\lambda_1\pi_1 \geq \lambda_2\pi_2$ ($\lambda_2\pi_2 \geq \lambda_1\pi_1$).*

2. *There is an open set of parameters where the radical (incremental) line is strictly preferred by the planner yet the radical (incremental) line is not a laissez faire equilibrium.*

When $\lambda_1\pi_1 \geq \lambda_2\pi_2$, the racing distortion is stronger than the underappropriation distortion: competitive pressure to finish some project quickly pushes firms off the difficult radical invention 2, leading them to work on incremental project 1 even though they will only capture a fraction of the value of the follow-on invention 3. On the other hand, if $\lambda_2\pi_2 \geq \lambda_1\pi_1$, the underappropriation externality is stronger: firms work on the radical invention because although the incremental first step is easy, inventing firms must in expectation share the continuation value generated once 3 is eventually invented. Note also from Corollary 3 that under patents of maximal strength $\gamma = 1$, only racing behavior distorts the firm equilibrium. Therefore, contrary to intuition, innovation will be excessively *incremental* in technological areas where patents are de facto effective in allowing originating firms to accrue most of the rents from follow-on innovation.

Essentially, if rivals are trying to invent a very difficult, very valuable new invention, a firm can instead shift to a less valuable substitute research line where the initial steps are not that challenging. Since the initial steps are not that hard, it is likely the firm will get the patent before its rivals make the radical discovery, and hence even though the incremental line offers a less valuable industry, it offers the firm a high probability of holding an industry-controlling patent. The usual intuition that patents are necessary for radical invention is based on the idea that, in the absence of patents, firms will not capture enough of the social value ex-post of their invention to justify a large research investment. A directional model, on the other hand, clarifies that strong patents also encourage inefficient ex-ante racing for critical patents which may very well be incremental in nature.

4.2 Patent Pools

Patent pools and research joint venture cross-licensing agreements, made in advance of R&D investments, are common in many industries and their potential welfare-enhancing role in solving hold-up problems is well-known (Lerner and Tirole (2004), Denicolo (2002)). Our main result suggests an alternative welfare-enhancing role for patent pools: reducing directional distortion. The socially optimal inventions in an industry like semiconductors might reasonably be known better by private sector firms than the planner, or for political economy reasons the planner may not want to be seen differentially incentivizing specific inventions. Firms would prefer to commit to innovating only along the optimal research line, but as we have shown, research on these inventions is not generally a laissez faire equilibrium, hence such a commitment must be binding if the firms are to avoid deviating.

Patent pools serve this commitment role. From Proposition 4, if the payoff to inventing firms and non-inventing firms is identical, then there is no directional inefficiency. A pool that involves the R&D intensive firms in an industry committing to share revenue from the stream of inventions that follow will generate first-best directional efficiency even if there is no ability to commit to precisely which type of innovation will be funded by each firm, and even if the pool organizer does not know which inventions lie along the planner-optimal research line. We completely abstract away from deadweight loss generated either in the product market or in the technology licensing market, and so answering questions about the exact welfare implication of patent pools is beyond the bounds of our model as currently constructed. That said, the role of patent pools, research joint ventures and standard setting in generating an efficient industry research portfolio appears understudied: organizations like SEMATECH (Irwin and Klenow (1996)) rather explicitly set the goal of coordinating industry research on particular research lines, rather than just hoping to reduce licensing frictions ex-post, or free riding ex-ante.

4.3 Trade Expansion and the Direction of Innovation

Trade is often considered a net positive for innovation, both because it expands the size of the markets, and because it assists in the diffusion of knowledge (e.g., Bloom, Draca and van Reenen (2015)). However, trade can be problematic if it distorts the direction of innovation. Consider a version of our base model where the number of firms is endogenous, retaining the assumption that the total measure of science in society M is fixed.[16] Firms are assumed to pay a fixed cost F at time 0 to enter, and no entry or exit occurs after that date. If entry decisions are made simultaneously, then the number of firms is the largest integer such that $V_{i\mathcal{P}}(s_0, N) \geq F$. That is, firms enter as long as their expected discounted profits exceed the fixed cost.

Assume that an expansion of trade increases the immediate social payoff π to all inventions by a constant factor ζ. From Proposition 1 and Footnote 8, we know that holding the number of firms N constant, neither the planner

[16] For example, an expansion in North-South trade may expand the potential product market without changing the size of the research sector.

optimal projects nor the firm equilibria change. However, the value of the invention graph for each firm increases to $\zeta V_{i\mathcal{P}}(s_0, N)$, which is equivalent to a reduction in the entry cost to $\frac{F}{\zeta}$ when calculating the number of firms who enter in the long-run equilibrium. Thus, an expansion of trade implies that the equilibrium number of firms also increases.

The following proposition shows that a large expansion in the size of the market, caused by trade, eventually causes so much entry in the R&D space that it distorts the direction of invention away from the planner optimum.

Proposition 6. *Assume that prior to an expansion of trade, when $\zeta = 1$, the direction of invention is efficient. Suppose there exists $s \in \Omega$ such that the planner optimal project is $s' \in S(s)$ and $\lambda_{s'}\pi_{s'} < \lambda_{\bar{s}}\pi_{\bar{s}}$ for some $\bar{s} \in S(s)$. Then, there exists an expansion of trade $\bar{\zeta} > 1$ such that the direction of invention is inefficient.*

That an increase in the size of the product market can cause an increase in the number of producing firms under constant returns to scale is intuitive. The idea that increased competition among R&D performing firms following a decrease in trade barriers can force firms to switch their research toward projects which are either more immediately lucrative or quicker to complete is a complaint that has been made by industry participants. For instance, Zheng and Kammen (2014) show that solar R&D spending fell following the rapid entry of Chinese firms into the photovoltaic industry after 2010, and firms both inside and outside of China decreased investment particularly in more fundamental research programs.

4.4 Pioneer patents

Patents considered pioneering in a technological field are often granted wide breadth ex-post by the courts (Love (2012)).[17] Granting broad patent protection to technologically significant inventions which pioneer a valuable field, when those inventions were first but not necessarily best, inefficiently distorts ex-ante incentives in infant industries, as firms race to garner the industry-controlling patent. For example, the famous Wright Brothers '393 patent cov-

[17] As Merges and Nelson (1990) note, the scope of a patent is determined to a large extent ex-post. This subjective evaluation allows judges to condition the breadth, and hence value, of patents on information including the social value and difficulty of substitutes for the invention in question. As Proposition 4 showed, it is precisely this type of non-neutral payment to inventors that is required to align directional incentives.

ers airplanes which maintain lateral control with aileron flaps, even though the Wright invention itself uses a simpler technology called wing warping which was rarely seen thereafter. Corollary 3 shows that the distortionary effects of patents are particularly strong when there are multiple potential inventions with much higher social continuation value than immediate payoff, precisely the situation where broad pioneer patents are often seen.

When a judge is considering how broadly to interpret a patent for a pioneer technology, the question ought not be whether a particular invention made possible a valuable field, but rather whether alternative, potentially-infringing, yet easier-to-improve technologies were feasible research targets at the time the supposed pioneer technology was created. Though it may often be difficult to assess whether a potentially infringing technology could feasibly have been invented at the time the pioneer technology came about, this does not strike us as a qualitatively more difficult problem than the existing legal question, assessing whether a pioneer technology was, in fact, the technological linchpin in the industry that followed.

5 Conclusion

We provide three novel contributions in the paper.

First, we construct a tractable dynamic model of the direction of innovation, with an arbitrary number of inventions where the value and difficulty to invent each can vary arbitrarily depending on what other inventions have already been discovered. We show that it is possible to transform the planner and firm problems into linear programs, allowing us to characterize their maximand as a simple index which can be tractably analyzed. Many economic situation involve agents choosing actions that affect an exponential hazard rate, generating linear functional maximands, hence transforming these maximands to a linear program and working directly with the index functions as in Proposition 2 may prove to be a useful technique in other contexts.

Second, we show that firms allocate their scientists inefficiently in equilibrium due to the contribution in a directional context of both an underappropriation and a racing distortion; these distortions are analogous to well-known distortions in models of the *rate* of innovation. Neither patents nor prizes fully ameliorate these distortions, and hence neither class of policy can generically generate optimal direction. Indeed, patents and prizes can both make directional distortion worse than laissez faire.

Third, there is a fundamental limit to decentralized "autopilot" innovation policy when it comes to directional incentives. Any innovation policy which both rewards inventors and conditions inventor incentives only on ex-post observable parameters like the difficulty or social value of realized inventions is incapable of always generating directional efficiency. This is quite different from models of the rate of invention, where a planner need not know everything firms know about a technological area to induce efficiency. For instance, in sequential invention models, laissez faire is inefficient due to underappropriation of the value early inventors grant to those who will build on that invention. A patent causes firms to internalize those positive spillovers, but the planner need not know exactly how large the spillovers will be ex-ante for the patent system to work. On the contrary, directional inefficiency is caused by the interaction of two distortions, and precisely which distortion is dominant and hence must be counteracted by policy depends on the nature of potential inventions a firm could have worked on. The planner needs both to correct distortions and to *know which type of distortion needs correcting*. Ex post observation of the equilibrium path is not sufficient to solve the latter problem.

We make a number of assumptions to permit the cleanest possible understanding of the fundamental mechanisms of directional distortion. Loosening these assumptions, in the Online Appendix we show that our qualitative results hold when firms are no longer symmetric, when the hazard rate of invention is nonlinear, and when some of the benefits of invention spill over to rival firms. In order to maintain symmetry across firms in our main model, we do not permit firms to keep inventions secret, for firms to apply learning from unsuccessful research to future inventions, or for firms to license patents except in the most reduced form manner. These restrictions provide analytic tractability for a model powerful enough to investigate general policies while remaining stylized enough to clearly separate the unique distortions introduced by direction choice. A model of innovation direction which permits other distortions already examined in the theoretical literature, as may be required for empirical models of the severity of directional distortion and its harm on welfare, is a particularly productive extension.

References

Acemoglu, Daron. 2002. "Directed Technical Change." *Review of Economic Studies*, 69.4: 781–809.

Acemoglu, Daron. 2011. "The Rate and Direction of Inventive Activity Revisited." , ed. Scott Stern and Josh Lerner, Chapter Diversity and Technological Progress, 319–356. University of Chicago Press.

Aghion, Philippe, and Jean Tirole. 1994. "The Management of Innovation." *Quarterly Journal of Economics*, 109.4: 1185–1209.

Aghion, Philippe, Mathias Dewatripont, and Jeremy C. Stein. 2008. "Academic Freedom, Private-Sector Focus, and the Process of Innovation." *RAND Journal of Economics*, 39.3: 617–635.

Akcigit, Ufuk, Douglas Hanley, and Nicolas Serrano-Velarde. 2013. "Back to Basics: Basic Research Spillovers, Innovation Policy and Growth." *Working Paper.*

Arrow, Kenneth J. 1962. "The Economic Implications of Learning by Doing." *The Review of Economic Studies*, 29.3: 155–173.

Becker, Johannes Gerd, and Damian S. Damianov. 2006. "On the Existence of Symmetric Mixed Strategy Equilibria." *Economic Letters*, 90.1.

Bloom, Nicholas, Mirko Draca, and John van Reenen. 2015. "Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT and Productivity." *The Review of Economic Studies.*

Charnes, A., and W. W. Cooper. 1962. "Programming with Linear Fractional Functionals." *Naval Research Logistics Quarterly*, 9(3-4): 181–186.

Chatterjee, Kalyan, and Robert Evans. 2004. "Rivals Search for Buried Treasure: Competition and Duplication in R&D." *RAND Journal of Economics*, 35.1: 160–183.

Dasgupta, Partha, and Eric Maskin. 1987. "The Simple Economics of Research Portfolios." *The Economic Journal*, 97: 581–595.

David, Paul A., Bronwyn Hall, and Andrew A. Toole. 2000. "Is Public R&D a Complement or Substitute for Private R&D? A Review of the Econometric Evidence." *Research Policy*, 29(4-5): 497–529.

Denicolo, Vincenzo. 2002. "Sequential Innovation and the Patent-Antitrust Conflict." *Oxford Economic Papers*.

Goolsbee, Austan. 1998. "Does Government R&D Policy Mainly Benefit Scientists and Engineers?" *The American Economic Review Papers and Proceedings*, 88.2: 298–302.

Green, Jerry R., and Suzanne Scotchmer. 1995. "On the Division of Profit in Sequential Innovation." *RAND Journal of Economics*, 26.1: 20–33.

Holmstrom, Bengt. 1989. "Agency Costs and Innovation." *Journal of Economic Behavior and Organization*.

Hopenhayn, Hugo, and Francesco Squintani. 2016. "On the Direction of R&D."

Irwin, Douglas A., and Peter Klenow. 1996. "High-tech R&D subsidies Estimating the effects of Sematech." *Journal of International Economics*.

Keller, Godfrey, and Alison Oldale. 2003. "Branching Bandits: A Sequential Search Process with Correlated Pay-offs." *Journal of Economic Theory*, 113.2: 302–315.

Keller, Godfrey, Sven Rady, and Martin Cripps. 2005. "Strategic Experimentation with Exponential Bandits." *Econometrica*, 73.1: 39–68.

Kennedy, Charles. 1964. "Induced Bias in Innovation and the Theory of Distribution." *The Economic Journal*, 74: 541–547.

Lerner, Josh. 1999. "The Government as Venture Capitalist: The Long-Run Impact of the SBIR Program." *Journal of Business*, 72.3: 285–318.

Lerner, Josh, and Jean Tirole. 2004. "Efficient Patent Pools." *American Economic Review*.

Loury, Glenn C. 1979. "Market Structure and Innovation." *The Quarterly Journal of Economics*, 93.3: 395–410.

Love, Brian J. 2012. "Interring the Pioneer Patent Doctrine." *North Carolina Law Review*, 90.

Merges, Robert P., and Richard R. Nelson. 1990. "On the Complex Economics of Patent Scope." *Columbia Law Review*, 90.4.

Nelson, Richard R. 1983. "Government Support of Technical Progress: Lessons from History." *Journal of Policy Analysis and Management*, 2.4: 499–514.

Reinganum, Jennifer F. 1981. "On the Diffusion of New Technology: A Game Theoretic Approach." *The Review of Economic Studies.*

Samuelson, Paul A. 1965. "A Theory of Induced Innovations along Kennedy-Weisacker Lines." *Review of Economics and Statistics*, 47: 444–464.

Zheng, Cheng, and Daniel M. Kammen. 2014. "An innovation-focused roadmap for a sustainable global photovoltaic industry." *Energy Policy.*

6 Online Appendix A: Proofs

Proofs are presented here in the order the propositions and corollaries appear in the main text. Corollary 1 and Proposition 3 involve straightforward algebraic manipulation of earlier results, so they are omitted below.

6.1 Preliminaries: Charnes-Cooper transformation

A linear fractional program is defined as

$$\max \frac{c^T \cdot \mathbf{x} + \alpha}{d^T \cdot \mathbf{x} + \beta} \quad \text{subject to } \mathrm{A}\mathbf{x} \leq \mathbf{b},\ \mathbf{x} \geq 0.$$

Using the Charnes-Cooper transformation, a linear fractional program can be transformed into an equivalent linear program (Charnes and Cooper (1962)), by defining the auxiliary variables $\mathbf{y} = \frac{1}{d^T\mathbf{x}+\beta}\mathbf{x}, \quad t = \frac{1}{d^T\mathbf{x}+\beta}$. Then, the original problem is equivalent to

$$\max c^T \cdot \mathbf{y} + \alpha t \quad \text{subject to} \quad A\mathbf{y} \leq \mathbf{b}t \quad \text{and} \quad d^T\mathbf{y} + \beta t = 1,\ \mathbf{y} \geq 0,\ t \geq 0.$$

6.2 Proof of Proposition 1

6.2.1 Part 1: Planner Optimum

1. It is easy to show that there exists a symmetric solution to the planner's problem (even with a weakly concave rate hazard rate $h(x)$).[18]

2. Charnes-Cooper transformation. Let $c_{s'} = \lambda(s,s')[\pi(s,s') + V_p(s')]$, $d_{s'} = \lambda(s,s')$, $\alpha = 0$, $\beta = r$, $A = [1, ..., 1]^T$, and $b = M$. The original planner problem can be transformed into the equivalent optimization program

$$\max_{\{y(s,s')\}_{s' \in S(s)}} \sum_{s' \in S(s)} \lambda(s,s')[\pi(s,s') + V_p(s')]y(s,s')$$

subject to $\sum_{s' \in S(s)} y(s,s') \leq Mt$, $\sum_{s' \in S(s)} \lambda(s,s')y(s,s') + rt = 1$, $y(s,s') \geq 0$, and $t \geq 0$. Notice that in our case $t \geq 0$ is redundant. Solving for t, defining $u(s,s') = \dfrac{\lambda(s,s')[\pi(s,s') + V_p(s')]}{1 + \frac{M}{r}\lambda(s,s')}$ and $v(s,s') =$

[18] If $\{(x_i(s'))_{s' \in S(s)}\}_{i=1}^N$ is a solution, $x(s,s') = h^{-1}\left(\frac{1}{N}\sum_{i=1}^N h(x_i(s,s'))\right)$ is a symmetric solution.

$y(s,s')\left(1+\frac{M}{r}\lambda(s,s')\right)$ we can rewrite the above problem as

$$\max_{\{v(s,s')\}_{s'\in S(s)}} \sum_{s'\in S(s)} u(s,s')v(s,s') \quad \text{subject to} \quad \sum_{s'\in S(s)} v(s,s') \leq \frac{M}{r}.$$

Define $T(s) = \arg\max_{\tilde{s}\in S(s)} u(s,\tilde{s})$. The solution to the above maximization is given by $\sum_{s'\in T(s)} v(s,s') = \frac{M}{r}$ and $v(s,s') = 0$, for $s' \notin T(s)$. In terms of the original variables we have the solution:

$$\sum_{s'\in T(s)} x(s,s') = M, \quad x(s,s') = 0, \quad s' \notin T(s).$$

6.2.2 Part 2: Firm Best Response

Using the Charnes-Cooper transformation, we identify $c_{s'} = \lambda(s,s')[w(s,s') + V_{\mathcal{P}s'}]$, $d_{s'} = \lambda(s,s')$, $\alpha = \sum_{s'\in S(s)} x_{-i}(s,s')d_{s'}(z(s,s') + V_{\mathcal{P}s'})$, $\beta = r + \sum_{s'\in S(s)} x_{-i}(s,s')d_{s'}$, $A = [1,....,1]^T$, and $b = \frac{M}{N}$. Similar to Part 1, the problem is

$$\max \sum_{s'} (\beta c_{s'} - \alpha d_{s'})\, y(s') \quad \text{subject to} \ \sum_{s'} \left(A + \frac{b}{\beta} d_{s'}\right) y(s') \leq \frac{b}{\beta}.$$

Define $T(s) = \arg\max_{\tilde{s}\in S(s)} \dfrac{c(\tilde{s}) - \frac{\alpha}{\beta}d(\tilde{s})}{1 + \frac{b}{\beta}d(\tilde{s})}$. Analogous to the previous proposition, the solution is to allocate all the effort on states in $T(s)$. The solution of the original problem is:

$$\sum_{s'\in T(s)} x_i(s,s') = \frac{M}{N}, \qquad x_i(s,s') = 0 \text{ otherwise.}$$

6.3 Proof of Proposition 2

6.3.1 Part 1: The Planner Optimum with Δ

The condition for the planner optimum is:

$$\begin{aligned}
\frac{\lambda_{s'}}{r + M\lambda_{s'}} P_{ps'} &\geq \frac{\lambda_\ell}{r + M\lambda_\ell} P_{p\ell}, \forall \ell \in S(s) \\
\lambda_{s'} P_{ps'} \frac{r + M\lambda_\ell}{r + M\lambda_{s'}} &\geq \lambda_\ell P_{p\ell} \\
\lambda_{s'} P_{ps'} &\geq \lambda_\ell P_{p\ell} - \lambda_{s'} P_{ps'} \Delta(s',\ell)
\end{aligned}$$

6.3.2 Part 2: The Firm Equilibrium with Δ

The planner-optimal s' is an equilibrium if for all $\ell \in S(s)$

$$\frac{\lambda_{s'}}{\tilde{r} + M\lambda_{s'}} \bar{P}_{fs'} \geq \frac{\lambda_\ell}{\tilde{r} + M\lambda_\ell} \bar{P}_{f\ell}$$

where $\bar{P}_{f\ell} = w_\ell + V_{\mathcal{P}\ell} - \Theta$. Rearranging terms, as in the first part of this proposition, we obtain

$$\lambda_{s'} \bar{P}_{fs'} \geq \lambda_\ell \bar{P}_{f\ell} - \lambda_{s'} \bar{P}_{fs'} \Delta(s', \ell) + (N-1)(\lambda_\ell \bar{P}_{f\ell} - \lambda_{s'} \bar{P}_{fs'})$$

which is equivalent to

$$\lambda_{s'} P_{ps'} \geq \lambda_\ell P_{p\ell} - \lambda_{s'} P_{ps'} \Delta(s', \ell) + (N-1)(\lambda_\ell \bar{P}_{f\ell} - \lambda_{s'} \bar{P}_{fs'}) + \Lambda_2(s', \ell),$$

where $\Lambda_2(s', \ell) = \lambda_{s'}(P_{ps'} - \bar{P}_{fs'}) - \lambda_\ell(P_{p\ell} - \bar{P}_{f\ell}) - \lambda_{s'} \Delta(s', \ell)(\bar{P}_{fs'} - P_{ps'})$. This is also equivalent to

$$\lambda_{s'} P_{ps'} \geq \lambda_\ell P_{p\ell} - \lambda_{s'} P_{ps'} \Delta(s', \ell) + \Lambda_3(s', \ell),$$

where $\Lambda_3(s', \ell) = \lambda_{s'}(P_{ps'} - N\bar{P}_{fs'}) - \lambda_\ell(P_{p\ell} - N\bar{P}_{f\ell}) - \lambda_{s'} \Delta(s', \ell)(\bar{P}_{fs'} - P_{ps'})$. Using the definition of $\Delta(s', \ell)$, and Θ we can show that this is equivalent to

$$\lambda_{s'} P_{ps'} \geq \lambda_\ell P_{p\ell} - \lambda_{s'} P_{ps'} \Delta(s', \ell) + \Lambda_4(s', \ell),$$

where $\Lambda_4(s', \ell) = \lambda_{s'}(P_{ps'} - N(w_{s'} + V_{\mathcal{P}s'})) - \lambda_\ell(P_{p\ell} - N(w_\ell + V_{\mathcal{P}\ell})) - \lambda_{s'} \Delta(s', \ell)(w_{s'} + NV_{\mathcal{P}s'} - P_{ps'} + (N-1)z_{s'})$. Notice that if we had added NP_p rather than P_p and divide by N, we obtain the condition:

$$\lambda_{s'} P_{ps'} \geq \lambda_\ell P_{p\ell} - \lambda_{s'} P_{ps'} \Delta(s', \ell) + D_1(s', \ell) + D_3(s', \ell) + D_2(s', \ell),$$

This expression can be written as three terms:

$$\begin{aligned}
D_1(s', \ell) &= \lambda_{s'}(P_{ps'} - (w_{s'} + V_{\mathcal{P}s'})) - \lambda_\ell(P_{p\ell} - (w_\ell + V_{\mathcal{P}\ell})) \\
D_2(s', \ell) &= \left(\frac{N-1}{N}\right) \lambda_{s'} \Delta(s', \ell) P_{ps'} \\
D_3(s', \ell) &= \frac{1}{N} \lambda_{s'} \Delta(s', \ell)(P_{ps'} - (w_{s'} + (N-1)z_{s'} + NV_{\mathcal{P}s'}))
\end{aligned}$$

6.4 Proof of Corollary 2 and 3

Let $D^*_{LF}(s',\ell) = D_1(s',\ell) + D_2(s',\ell) + D_3(s',\ell)$, the distortion under laissez faire. Straightforward algebra applied to Proposition 2 shows that a prize q changes $D^*_{LF}(s',\ell)$ by adding the term $q(\lambda_\ell - \lambda_{s'} - \frac{\lambda_{s'}\Delta(s',\ell)}{N})$. Using the definition of Δ, this term is equal to $q\frac{\Delta(s',\ell)\tilde{r}}{N}$.

Likewise, distortions under the policy $\mathcal{P}_\gamma$ can be calculated directly by plugging the patent transfers w and z into Proposition 2.

6.5 Proof of Proposition 4: Information-Constrained Efficiency

With arbitrary transfers it is always possible to implement the efficient solution with transfers that depend only on on-path parameters. For any $\ell \in S(s)$, set $w_\ell = z_\ell = \pi_\ell$. Because our invention graph is finite, every research line eventually reaches a state $\bar{s}$ where $S(\hat{s}) = \emptyset, \forall \hat{s} \in S(\bar{s})$. For any invention in $S(\bar{s})$, then, the future is trivially efficient under any policy. If the future is efficient, then $V_{i\mathcal{P}s'} = V_{Ps'}$

Applying induction, if the future is efficient, then under transfers $w(s,s') = \pi(s,s')$ and $z(s,s') = \pi(s,s')$, $w(s,s') + V_{i\mathcal{P}s'} = z(s,s') + V_{i\mathcal{P}s'} = P_{ps'}$. From Proposition 2, we get that $D_1(s',\ell) = 0$ and $D_2(s'\ell) = -D_3(s'\ell)$. That is, there is no distortion in the firm choice, hence there is an efficient equilibrium.

To show part 2, let the planner be indifferent between two inventions s' and ℓ. By Proposition 2,

$$\lambda_{s'}P_{s'}(1 + \Delta(s',\ell)) = \lambda_\ell P_\ell \tag{1}$$

and

$$\lambda_\ell P_\ell(1 + \Delta(\ell,s')) = \lambda_{s'}P_{s'} \tag{2}$$

Let the firm choose transfers w and z without conditioning those transfers on off-equilibrium-path parameters. That is, $w_{s'}$ and $z_{s'}$ cannot condition on λ_ℓ or π_ℓ, and likewise for w_ℓ and z_ℓ. By the assumption that inventors are paid at least as much as non-inventors, let $w_{s'} - z_{s'} = \epsilon_{s'} > 0$. Let $f_s = w_s + V_{i\mathcal{P}s}$ be the total transfer to inventing firms inclusive of continuation value.

Again using Proposition 2 and rearranging terms, we have that s' is a firm equilibrium if

$$\lambda_{s'}f_{s'}(1 + \Delta(s',\ell)) \geq \lambda_\ell f_\ell + \frac{N-1}{N}\lambda_{s'}\Delta(s',\ell)\epsilon_{s'} \tag{3}$$

and ℓ is an equilibrium if

$$\lambda_\ell f_\ell(1+\Delta(\ell,s')) \geq \lambda_{s'} f_{s'} + \frac{N-1}{N}\lambda_\ell \Delta(\ell,s')\epsilon_\ell \tag{4}$$

From equations 3 and 1 we get:

$$\frac{f_{s'}}{f_\ell} \geq \frac{P_{s'}}{P_\ell} + \left(\frac{N-1}{N}\right)\frac{\lambda_{s'}P_{s'}}{\lambda_\ell P_\ell f_\ell}\Delta(s',\ell)\epsilon_{s'} \tag{5}$$

From equations 4 and 2 we get:

$$\frac{f_\ell}{f_{s'}} \geq \frac{P_\ell}{P_{s'}} + \left(\frac{N-1}{N}\right)\frac{\lambda_\ell P_\ell}{\lambda_\ell P_{s'} f_{s'}}\Delta(\ell,s')\epsilon_\ell \tag{6}$$

We have two cases:

1. Suppose $\frac{f_{s'}}{f_\ell} \leq \frac{P_{s'}}{P_\ell}$. If $\Delta(s',\ell) > 0$, then equation 5 implies $\frac{f_{s'}}{f_\ell} > \frac{P_{s'}}{P_\ell}$, and therefore s' cannot be a firm equilibrium.
2. Suppose $\frac{f_{s'}}{f_\ell} \geq \frac{P_{s'}}{P_\ell}$. If $\Delta(\ell,s') > 0$, then equation 6 implies $\frac{f_{s'}}{f_\ell} < \frac{P_{s'}}{P_\ell}$, and therefore ℓ cannot be a firm equilibrium.

Note that by construction $\Delta(s',\ell)$ and $\Delta(\ell,s')$ have opposite signs, and their signs depend on the simplicity of both ℓ and s'. In words, the planner needs to provide a higher relative transfer to the more difficult project in order to stop racing behavior when $w < z$. However, by assumption firms cannot condition transfers on the parameters of off-path inventions, and hence cannot condition on the sign of Δ.

By continuity, we can drop the assumption that the planner is indifferent and instead give the planner an arbitrarily small strict preference η for one invention or the other, and proceed with the proof as above (since $\epsilon_{s'}$ and $\Delta(s',\ell)$ are fixed, when we take $\eta \to 0$ we get the same result). Hence, for any information-constrained payoff functions such that $w_{s'} > z_{s'}$, a set of inventions can be chosen so the firm equilibrium is inefficient.

6.6 Proposition 5: Radical vs Incremental Steps

Let there be two potential research lines: A *radical* line with a single relatively difficult invention (invention 2), and an *incremental* line with two sequential inventions (inventions 1 and 3, where 3 cannot be worked on unless 1 has been invented). Invention 2 and 3 are perfect substitutes: if invention 2 is discovered before 3, $\pi_3 = 0$, and vice versa. We also assume that the difficulty of each research line is such that the planner is indifferent between working on either line. The planner indifference condition in the initial state is:

$$\left(\frac{M\lambda_1}{r + M\lambda_1}\right)\left(\pi_1 + \frac{M\lambda_3\pi_3}{r + M\lambda_3}\right) = \frac{M\lambda_2\pi_2}{r + M\lambda_2}. \tag{7}$$

When condition 7, and the given assumptions $\lambda_1 > \lambda_2$, $\lambda_3 > \lambda_2$, and $\pi_2 \geq \pi_3$ hold, it can be shown using the firm equilibrium condition that once invention 1 is invented, all firms working on invention 3 is an equilibrium. Thus, the continuation value after invention 1 is $V_{i1} = \frac{M\lambda_3}{r+M\lambda_3}$. By Corollary 1, all firms working on invention 1 is a firm equilibrium if

$$\lambda_1 P_1 \geq \lambda_2 P_2 - \lambda_1 P_1 \Delta(1,2) + (N-1)(\lambda_2\pi_2 - \lambda_1\pi_1).$$

In this case, $P_1 = \pi_1 + \frac{M\lambda_3\pi_3}{r+M\lambda_3}$ and $P_2 = \pi_2$. Using the equality from the planner's indifference condition, we can replace the value of P_1 and obtain the equivalent condition

$$(1 + \Delta(1,2))\frac{M\lambda_2\pi_2(r + \lambda_1)}{r + M\lambda_2} \geq \lambda_2\pi_2 + (N-1)(\lambda_2\pi_2 - \lambda_1\pi_1).$$

Finally, using the definition of $\Delta(1,2)$, the condition is equivalent to:

$$\lambda_1\pi_1 \geq \lambda_2\pi_2.$$

Analogously we show that all firms working on 2 is an equilibrium when $\lambda_2\pi_2 \geq \lambda_1\pi_1$. Part 2 follows immediately from that result and local continuity of the firm best response condition.

6.7 Proposition 6: Trade Expansion and Endogenous Firm Entry

By assumption, without the trade expansion there is an equilibrium number of firms $\bar{N}$ such that the firm equilibrium is efficient. A market expansion, caused by trade, is equivalent to a reduction in the entry cost, with the number of firms rising to ∞ and entry costs falling to zero. Let $s \in \Omega$ such that $\lambda_{s'}\pi_{s'} < \lambda_\ell\pi_\ell$, where s' is the efficient solution. By Proposition 3, as N increases firms will deviate from s', since D_2 and D_3 are bounded.

7 Online Appendix B: Extensions of the Baseline Model

In this appendix, we show that adding decreasing or increasing returns to scale to our model does not change the underlying source of firm inefficiency, that decreasing returns to scale make inefficiency in the firm equilibrium more likely, that there is no inefficiency when the parameters of inventive opportunities tomorrow do not depend on which inventions are discovered today, that a single element of state dependence in conjunction with multiple research lines generates inefficiency, and that permitting both short-lived and infinite-lived research firms exacerbates the racing distortion.

7.1 Planner Problem with Nonlinear Hazard Rates

First, consider alternative assumptions about returns to scale. Let the hazard rate on invention k for firm i be $\lambda_k h(x_k)$, where h is twice-differentiable, $h' > 0$, $h(0) = 0$ and, without loss of generality, $h(\frac{1}{N}) = \frac{1}{N}$. Under decreasing returns to scale, $h'' < 0$, and under increasing returns, $h'' > 0$. Note that, in the results presented in the body of this paper, constant returns to scale under the above assumptions simply means that $h(x) = x$. To simplify notation, throughout this section we assume that there is no inefficiency in future states.

In section 6.2.1 in Appendix A, we showed that independence of hazard rates across firms means the planner optimizes with symmetric effort across firms. Without loss of generality, we assume $M = 1$, so the planner solves

$$\max_{\sum\limits_{s' \in S(s)} x_{s'} \leq \frac{1}{N},\ x_{s'} \geq 0, \forall s' \in S(s)} \frac{\sum_{s'} \lambda_{s'} P_{s'} N h(x_{s'})}{r + \sum_{s'} \lambda_{s'} N h(x_{s'})}$$

The KKT necessary condition imply that exist $\mu_{s'} \geq 0$ such that $\mu_{s'} x_{s'} = 0$ and γ such that

$$\frac{\partial f(x)}{\partial x_{s'}} = \gamma - \mu_{s'}.$$

A corner solution, where all effort goes to $k \in S(s)$, that is $x_k = \frac{1}{N}$ and

$x_\ell = 0$ for $\ell \neq k$ is characterized by

$$\frac{\lambda_k P_k h'(x_k)(r_N + \sum_{s'} \lambda_{s'} h(x_{s'})) - \lambda_k h'(x_k)(\sum_{s'} \lambda_{s'} P_{s'} h(x_{s'}))}{(r_N + \sum_{s'} \lambda_{s'} h(x_{s'}))^2} \geq$$
$$\frac{\lambda_\ell P_\ell h'(x_\ell)(r_N + \sum_{s'} \lambda_{s'} h(x_{s'})) - \lambda_\ell h'(x_\ell)(\sum_{s'} \lambda_{s'} P_{s'} h(x_{s'}))}{(r_N + \sum_{s'} \lambda_{s'} h(x_{s'}))^2}$$

where $r_N = \frac{r}{N}$. Using that $h(0) = 0$, this simplifies to

$$\lambda_k P_k h'(x_k)(r_N + \lambda_k h(x_k)) - \lambda_k h'(x_k)\lambda_k P_k h(x_k) \geq$$
$$\lambda_\ell P_\ell h'(0)(r_N + \lambda_k h(x_k)) - \lambda_\ell h'(0)\lambda_k P_k h(x_k)$$

Let $C = \frac{h'(\frac{1}{N})}{h'(0)}$. Note that under decreasing returns to scale, $C \in (0, 1)$. Thus, we can write

$$\lambda_k P_k C(r_N + \lambda_k h(x_k)) - \lambda_k C \lambda_k P_k h(x_k) \geq \lambda_\ell P_\ell (r_N + \lambda_k h(x_k)) - \lambda_\ell \lambda_k P_k h(x_k)$$

Using that $h(x_k) = \frac{1}{N}$ and rearranging terms, and defining $\Delta_C(k, \ell) = \dfrac{\lambda_\ell - C\lambda_k}{r + \lambda_k}$, we get

$$\lambda_k P_k C \geq \lambda_\ell P_\ell - \Delta_C(k, \ell)\lambda_k P_k.$$

Notice that this condition is equivalent to the planner's condition in Proposition 2. Similar derivation for an arbitrary number of scientists M, defining $C(M) = \frac{h'(\frac{M}{N})}{h'(0)}$, gives the same result.

The only caveat is that KKT are only necessary and not sufficient conditions. However, we show that when $h(x) = x$ the only solution is the corner solution x_k and in that case the condition above holds ($C = 1$). Thus, if inequality holds strictly for $C = 1$, it still holds for C close to 1, in which case we have full effort toward a single invention even with nonconstant returns to scale. Thus, even with small levels of decreasing or increasing returns to scale, the planner corner solution is retained.

7.2 Firm Problem with Nonlinear Hazard Rates

Under the assumption that parameters are such that the planner works on a single invention under decreasing returns to scale, we now show that the firms deviate for almost exactly the same reason as under constant returns. Indeed, decreasing returns to scale make it more likely that firms will deviate because minor deviations to new research lines will generate a higher relative hazard rate under decreasing returns than under constant returns, hence exacerbate the racing distortion.

Suppose that all rivals are exerting efforts towards invention k. Recall the firm problem, if all other firms exert full effort towards invention k, is

$$\max_{\sum_{s' \in S(s)} x_{s'} \leq \frac{1}{N}, x_{s'} \geq 0, \forall s' \in S} \frac{\sum_{s'} \lambda_{s'} P_{fs'} h(x_{s'}) + A_k}{\tilde{r} + \sum_{s'} \lambda_{s'} h(x_{s'})}$$

where $A_k = (N-1)\lambda_k h(\frac{1}{N})V_{fk}$, and $\tilde{r} = r + (N-1)\lambda_k h(\frac{1}{N})$

As in Section 7.1, the first order necessary condition for positive effort on invention k and no effort on any other invention is

$$\frac{\lambda_k P_{fk} h'(x_k)(\tilde{r} + \lambda_k h(x_k)) - \lambda_k h'(x_k)(\lambda_k P_{fk} h(x_k) + A_k)}{(\tilde{r} + \lambda_k h(x_k))^2} \geq$$
$$\frac{\lambda_\ell P_{f\ell} h'(x_\ell)(\tilde{r} + \lambda_k h(x_k)) - \lambda_\ell h'(x_\ell)(\lambda_k P_{fk} h(x_k) + A_k)}{(\tilde{r} + \lambda_k h(x_k))^2}$$

This simplifies to

$$\lambda_k P_{fk} h'(x_k)(\tilde{r} + \lambda_k h(x_k)) - \lambda_k h'(x_k)(\lambda_k P_{fk} h(x_k) + A_k) \geq$$
$$\lambda_\ell P_{f\ell} h'(x_\ell)(\tilde{r} + \lambda_k h(x_k)) - \lambda_\ell h'(x_\ell)(\lambda_k P_{fk} h(x_k) + A_k)$$

Retaining the assumptions that $h(\frac{1}{N}) = \frac{1}{N}$ and $C = \frac{h'(\frac{1}{N})}{h'(0)}$, after simple algebra we get

$$\lambda_k P_{fk} C \geq \lambda_\ell P_{f\ell} + \frac{1}{N}\Delta_C(k,\ell)\lambda_k P_{fk} - \frac{1}{N}\Delta_C(k,\ell)(N-1)\lambda_k V_{fk}$$

Adding and subtracting terms, we get

$$\lambda_k P_k C \geq \lambda_\ell P_\ell - \Delta_C(k,\ell)\lambda_k P_k + D^*$$

where

$$D^* = \lambda_\ell(P_{f\ell} - P_\ell) - \lambda_k C(P_{fk} - P_k) + \frac{1}{N}\Delta_C(k,\ell)\lambda_k(P_k - (P_{fk} + (N-1)V_{fk})) + \frac{N-1}{N}\Delta_C(k,\ell)\lambda_k P_k$$

This distortion are analogous to the distortions in Proposition 2, with

$$\begin{aligned}
D_1^C(k,\ell) &= \lambda_\ell(P_{f\ell} - P_\ell) - \lambda_k C(P_{fk} - P_k) \\
D_2^C(k,\ell) &= \frac{N-1}{N}\Delta_C(k,\ell)\lambda_k P_k \\
D_3^C(k,\ell) &= \frac{1}{N}\Delta_C(k,\ell)\lambda_k(P_k - (P_{fk} + (N-1)V_{fk}))
\end{aligned}$$

Thus, adding small amounts of increasing or decreasing returns to scale does not change our main qualitative results.

7.3 Graphs Without State Dependence Have an Efficient Equilibrium

The decomposition in Proposition 2 allows a simple categorization of the nature of inefficiency generated by a particular policy in a particular *type* of invention graph. Inefficiency under laissez faire does not result from the simple existence of multiple projects. Rather, in order to generate laissez faire inefficiency, a necessary though not sufficient condition is that one firm's actions today must affect the existence of future research targets, or their value, or the difficulty of inventing them. This can be seen with the following simple cases.

First, let there be a set of research targets which are *technologically independent*.

Definition 7. *An invention graph involves* technologically independent inventions *if, in every state, the set of research targets* $S(s)$ *includes every invention in* $S(s_0)$ *which has yet to be invented, and the payoff* π *and simplicity* λ *of each undiscovered invention never change.*

With technological independence, no matter what is invented today, the options available to inventors tomorrow, and the simplicity and payoff of those inventions, does not change; there is nothing resembling a set of research lines, where invention today affects the nature of inventive opportunity tomorrow. As a result, Proposition 7 shows that on the technologically independent graph, laissez faire firm activity is efficient.

Proposition 7. *In an invention graph with technologically independent inventions, the planner optimally works on inventions in decreasing order of their immediate flow social payoff* $\lambda_{s'}\pi_{s'}$*. Further, there exists an* efficient *laissez faire firm equilibrium.*

Proof. We prove by induction. Let there be two remaining inventions. If invention i is discovered first, the expected discounted continuation value for the planner is $V_p(i) = \frac{\lambda_{-i}}{r+\lambda_{-i}}\pi_{-i}$. By Proposition 1, the planner works on invention i that node maximizes the index

$$\frac{\lambda_i}{r + M\lambda_i}[\pi_i + V_i]$$

Define $p_i = \frac{\lambda_i}{r+M\lambda_i}$. The planner discovers 1 first and 2 second if and only if

$$\left(\frac{p_1}{1-p_1}\right)\pi_1 \geq \left(\frac{p_2}{1-p_2}\right)\pi_2.$$

Using the definition of p_i, that inequality simplifies to

$$\lambda_1\pi_1 \geq \lambda_2\pi_2.$$

Now we prove the inductive step. Without loss of generality let $\lambda_1\pi_1 \geq \lambda_2\pi_2 \geq ... \geq \lambda_K\pi_K$. Define $p_i = \frac{\lambda_i}{r+\lambda_i}$ and notice that $\frac{p_i}{1-p_i} = \frac{\lambda_i}{r}$.

We know the result holds for $K = 2$. Assume the result is true for any set of $K-1$ inventions (Induction Hypothesis). Let's prove the result for K inventions. We need to show that starting from 1 is better than starting from any other invention k. By the characterization result, we start from 1 instead of k iff:

$$p_1(\pi_1 + V_p(1)) \geq p_k(\pi_k + V_p(k)), \quad \text{for all } k.$$

Since after one invention there are $K-1$ left, using the induction hypothesis we know that the planner discovers in decreasing order of $\lambda\pi$. Hence,

$$V_p(1) = \sum_{m=2}^{K}\left(\prod_{j=2}^{m} p_j\right)\pi_m \quad \text{and } V_p(k) = \sum_{m=1}^{k-1}\left(\prod_{j=1}^{m} p_j\right)\pi_m + \sum_{m=k+1}^{K}\left(\prod_{j=1,\ j\neq k}^{m} p_j\right)\pi_m.$$

Thus, the condition is equivalent to

$$\sum_{m=1}^{K}\left(\prod_{j=1}^{m} p_j\right)\pi_m \geq p_k\pi_k + p_k\sum_{m=1}^{k-1}\left(\prod_{j=1}^{m} p_j\right)\pi_m + \sum_{m=k+1}^{K}\left(\prod_{j=1}^{m} p_j\right)\pi_m, \quad \text{for all } k.$$

Notice that the terms from $k+1$ to K cancel out. This is because the expected time at which we reach invention $k+1$ is the same if we start from 1 or from k. Thus, we get

$$\sum_{m=1}^{k}\left(\prod_{j=1}^{m} p_j\right)\pi_m \geq p_k\pi_k + p_k\sum_{m=1}^{k-1}\left(\prod_{j=1}^{m} p_j\right)\pi_m.$$

which is equivalent to

$$\sum_{m=1}^{k-1}\left(\prod_{j=1}^{m} p_j\right)\pi_m(1-p_k) \geq p_k\pi_k\left(1-\left(\prod_{j=1}^{k-1} p_j\right)\right).$$

Thus, the planner start from invention 1 if and only if

$$\sum_{m=1}^{k-1} \lambda_m \pi_m \frac{\left(\prod_{j=1}^{m-1} p_j\right)(1-p_m)}{\left(1-\left(\prod_{j=1}^{k-1} p_j\right)\right)} \geq \lambda_k \pi_k, \quad \text{for all } k.$$

This always holds when the inventions are ordered by $\lambda\pi$, since the left hand side of the inequality is a convex combination of $\{\lambda_m \pi_m\}_{m=1}^{k-1}$, since the coefficients

$$a_m = \frac{\left(\prod_{j=1}^{m-1} p_j\right)(1-p_m)}{\left(1-\left(\prod_{j=1}^{k-1} p_j\right)\right)}$$

satisfy that $a_m \geq 0$ and $\sum_{m=1}^{k-1} a_m = 1$. The firm equilibrium then follows immediately: since the future is by induction efficient, by Proposition 2 the firms never deviate when the planner is working on the project with highest flow immediate payoff. □

7.4 State Dependent Invention Graphs Generate Inefficiency

Adding an element of state dependence, where invention today affects what can be worked on tomorrow, to the mere existence of multiple projects is enough to induce inefficiency under laissez faire. Consider a case where all inventions are available in the initial state, but there is no continuation value: once anything has been invented, the immediate social payoff of every other potential invention falls to zero.

Definition 8. *An invention graph involves* perfect substitutes *if all inventions are available in s_0 and any discovery reduces the immediate social payoff of all other inventions to $\pi = 0$.*[19]

With the social continuation value equal to zero, and inventing firms paid exactly the immediate social payoff of their invention, laissez faire on the perfect substitutes invention graph generates distortions $D_1(s', \ell) = D_3(s', \ell) =$

[19]Our model takes the immediate social payoff of an invention as the reduced form value from an unmodeled demand system. As such, we are in a sense abusing the term "perfect substitutes," but the manner in which the term is used here - two inventions are perfect substitutes if the marginal value of each is zero once the other has been invented - should nonetheless be clear.

0, leaving only the racing distortion D_3. Therefore, under perfect substitutes, firms only deviate toward projects which are easier than the planner optimum.

Proposition 8. *Under the laissez faire policy $\mathcal{P}_{LF}$ on the perfect substitutes invention graph, s' is planner optimal if $\forall \ell \in S(s)$*

$$\lambda_{s'}\pi_{s'} \geq \lambda_\ell \pi_\ell - \lambda_{s'}\pi_{s'}\Delta(s', \ell),$$

and s' is a firm equilibrium under laissez faire if and only if

$$\lambda_{s'}\pi_{s'} \geq \lambda_\ell \pi_\ell - \lambda_{s'}\pi_{s'}\Delta(s', \ell) + \underbrace{\left(\frac{N-1}{N}\right)\lambda_{s'}\pi_{s'}\Delta(s', \ell)}_{D_2(s',\ell)}.$$

The proof of Proposition 8 is straightforward algebra, hence is omitted.

The technologically independent inventions example shows that equilibrium direction choice is efficient, when all inventions are available from the beginning and there is not state dependency. The perfect substitutes example shows that simple forms of state contingency can generate inefficiency in the equilibrium direction. This case is a particular form of state dependency in parameter values, changing the immediate payoff π. We now show that another type of state contingency, availability of inventions only after other inventions, can also generate laissez faire directional inefficiencies.

Consider three inventions. Inventions 1 and 2 are available from the beginning. However, invention 3 becomes available only after 1 is invented. Figure 4a shows the inventions and Figure 4b the states representation.

Proposition 9. *Consider the invention graph in Figure 4. Then:*

1. *If $\lambda_3\pi_3 \leq \max\{\lambda_1\pi_1, \lambda_2\pi_2\}$, then the planner always works on the available invention with largest flow payoff $\lambda\pi$. By Proposition 2 this can be implemented as a firm equilibrium.*

2. *If $\lambda_3\pi_3 > \max\{\lambda_1\pi_1, \lambda_2\pi_2\}$, then the planner* opens a path *(works on invention 1 first) iff*

$$\lambda_1\pi_1 \geq \lambda_2\pi_2 + \left(\frac{\lambda_1}{r+\lambda_3}\right)[\lambda_2\pi_2 - \lambda_3\pi_3].$$

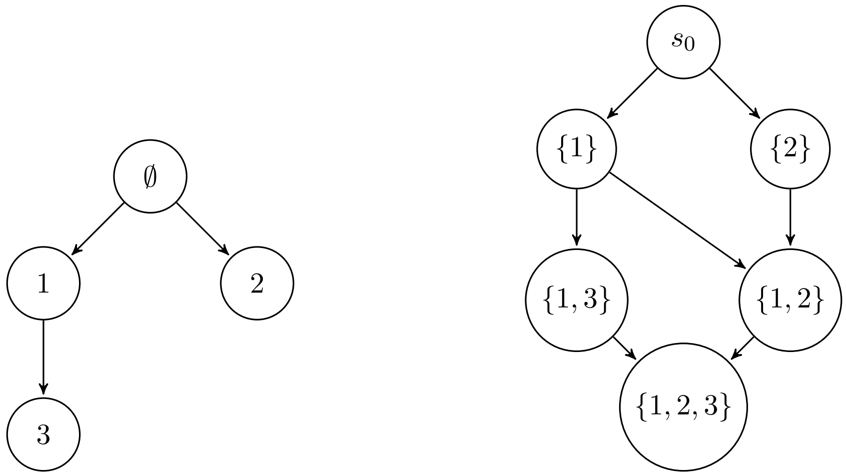

(a) Possible Inventions

(b) Graph of states.

Figure 4: Invention 3 is only available once 1 is invented.

Applying the firm equilibrium condition, the planner solution can be implemented as an equilibrium iff

$$\lambda_1\pi_1 \geq \lambda_2\pi_2 + \frac{r}{r+(N-1)(r+\lambda_1)}\left(\frac{\lambda_1}{r+\lambda_3}\right)[\lambda_2\pi_2 - \lambda_3\pi_3].$$

Proposition 9 says that the planner may work on invention 1 even when $\lambda_1\pi_1 < \lambda_2\pi_2$ as long as doing so makes available a third invention with even higher expected flow payoff and the future is not discounted too heavily. We prove this by examining all six permutations of flow immediate payoff across the inventions.

Proof. In the cases:

$$\lambda_1\pi_1 \geq \lambda_2\pi_2 \geq \lambda_3\pi_3, \quad \lambda_1\pi_1 \geq \lambda_3\pi_3 \geq \lambda_2\pi_2, \quad \lambda_2\pi_2 \geq \lambda_1\pi_1 \geq \lambda_3\pi_3,$$

the solution (for both planner and firms) is to discover in decreasing order of $\lambda_i\pi_i$, since the graph does not impose any binding constraints. This can be shown directly with Proposition 2.

Consider the following cases:

- Case (a): $\lambda_3\pi_3 \geq \lambda_1\pi_1 \geq \lambda_2\pi_2$.

- Case (b): $\lambda_3\pi_3 \geq \lambda_2\pi_2 \geq \lambda_1\pi_1$.
- Case (c): $\lambda_2\pi_2 \geq \lambda_3\pi_3 \geq \lambda_1\pi_1$.

In these cases, the planner optimum may involve working on 1 first in order to "open up" valuable invention 3. In case c, by Proposition 4, we know the planner works on 2 before 3 conditional on inventing 1. The planner would invent 1 before 2 if and only if

$$p_1\pi_1 + p_1p_2\pi_2 + p_1p_2p_3\pi_3 \geq p_2\pi_2 + p_1p_2\pi_1 + p_1p_2p_3\pi_3,$$

Algebraic manipulation shows this condition is equivalent to $\lambda_1\pi_1 \geq \lambda_2\pi_2$. Therefore, the planner will always work on the project with the highest available flow profit and therefore we can implement the planner solution as a laissez faire equilibrium.

Consider now cases (a) and (b). By Proposition 4, we know the planner will work on $3 \to 2$ after discovering 1. Therefore, the planner will first invent 1 if and only if

$$p_1\pi_1 + p_1p_3\pi_3 + p_1p_2p_3\pi_2 \geq p_2\pi_2 + p_1p_2\pi_1 + p_1p_2p_3\pi_3,$$

Moving terms around and multiplying the expression by $\frac{r}{(1-p_1)(1-p_2)} = \frac{(r+\lambda_1)(r+\lambda_2)}{r}$ we get

$$\lambda_1\pi_1 \geq \lambda_2\pi_2 + \left(\frac{\lambda_1}{r+\lambda_3}\right)[\lambda_2\pi_2 - \lambda_3\pi_3]$$

Now, using the result about equilibrium implementation of the planner solution we get the statement in the proposition. □

7.5 Spillovers

In the main results, under laissez faire, inventing firms collect the entire immediate social payoff of their invention, and non-inventing firms collect zero. Consider a policy where only a fraction α of the immediate social payoff is collected by inventors, with the remaining surplus accruing to all other firms, shared equally.

Definition 9. *Let a* spillover policy $\mathcal{P}_\alpha$ *provide inventors transfers* $w(s,s') = \pi(s,s')(1-(N-1)\alpha)$ *and noninventors* $z(s,s') = \alpha\pi(s,s')$*, Assume that* $\alpha \leq \frac{1}{N}$*, meaning inventors receiveweakly more than non-inventors.*

From proposition 2, it is easy to see that the distortions can be written as

$$D_\alpha(s,s') = D_{LF}(s,s') - (N-1)\alpha(\lambda_\ell \pi_\ell - \lambda_{s'}\pi_{s'}) + \mathcal{V}(\alpha) = (1-\alpha)D_{LF}(s,s') + \mathcal{V}(\alpha)$$

where $\mathcal{V}(\alpha)$ is the distortion from the difference between the social continuation value under laissez faire policy and spillover policy $\mathcal{P}_\alpha$. Thus, letting non-inventors get a share of the immediate payoff weakens the directional distortion caused by the laissez faire policy.

7.6 Short Run vs Long Run Firm Equilibrium

In the main results, we look only at homogenous, infinitely-lived firms with perfect information about parameter values. Much of the intuition in those results can be generalized. In this subsection, let there be one long run innovator who plays until everything is discovered, and a sequence of short run innovators who play only one period each. Short run players may be R&D firms who only have the technological ability to work on exactly the present set of invention opportunities; they hence put no weight on the social value created when their inventions open up future opportunities for other firms.

Consider an invention graph with two technologically independent inventions. Let the total number of scientists $M = 1$, with the long run and the short run firm both having $\frac{1}{2}$ scientist. Since the number of scientists is constant, just as in the case of technologically independent inventions the planner works first on 1 rather than 2 if and only if $\lambda_1\pi_1 \geq \lambda_2\pi_2$.

The long run firm has the same best response as in the technologically independent inventions case since the identity of the rivals is irrelevant. The short run innovator at any stage has the best response:

$$s' \in \arg\max_{\tilde{s}\in S(s)} \frac{\lambda_{\tilde{s}}\pi_{\tilde{s}}}{N(r + \sum_{z\in S(s)} a_{-iz}\lambda_z) + \lambda_{\tilde{s}}}$$

The continuation values for the long run player are

$$V(1) = \frac{\lambda_2\pi_2}{2r + 2\lambda_2}$$

and

$$V(2) = \frac{\lambda_1\pi_1}{2r + 2\lambda_1}$$

Suppose the long run firm initially works on invention 1. The short run firm, when both inventions are available, works on invention 1 if and only if:

$$\lambda_1\pi_1 \geq \lambda_2\pi_2 + \frac{\lambda_2\pi_2(\lambda_1 - \lambda_2)}{2r + \lambda_1 + \lambda_2} \Leftrightarrow \frac{\lambda_1\pi_1}{\lambda_2\pi_2} \geq 1 + \Delta_1.$$

Suppose the long run innovator initially works on 2. The short run innovator when both inventions are available works on 1 if and only if:

$$\lambda_1\pi_1 \geq \lambda_2\pi_2 + \frac{\lambda_2\pi_2(\lambda_1 - \lambda_2)}{2r + 2\lambda_2} \Leftrightarrow \frac{\lambda_1\pi_1}{\lambda_2\pi_2} \geq 1 + \Delta_2,$$

where $\Delta_2 > \Delta_1$ as long as $\lambda_1 \neq \lambda_2$.

Therefore, when $\lambda_1 = \lambda_2$, there is no inefficiency. When $\lambda_1 < \lambda_2$,

- If $\frac{\lambda_1\pi_1}{\lambda_2\pi_2} \geq 1$: Both long and short run firms working on 1 is an equilibrium (and it is efficient).
- If $\frac{\lambda_1\pi_1}{\lambda_2\pi_2} \leq 1+\Delta_1$: Both working on 2 is an equilibrium (and it is efficient)
- When $1 + \Delta_1 \leq \frac{\lambda_1\pi_1}{\lambda_2\pi_2} \leq 1$, the short run and long run firm working on 1 is not an equilibrium. In this case, the equilibrium is asymmetric, hence inefficient.

Analogous conditions hold if $\lambda_1 > \lambda_2$.

The equilibrium is depicted in the following figure, where Δ_2 is negative and Δ_1 is positive.

Figure 5: Equilibrium project choice with sequence of short run firms and a long run firm

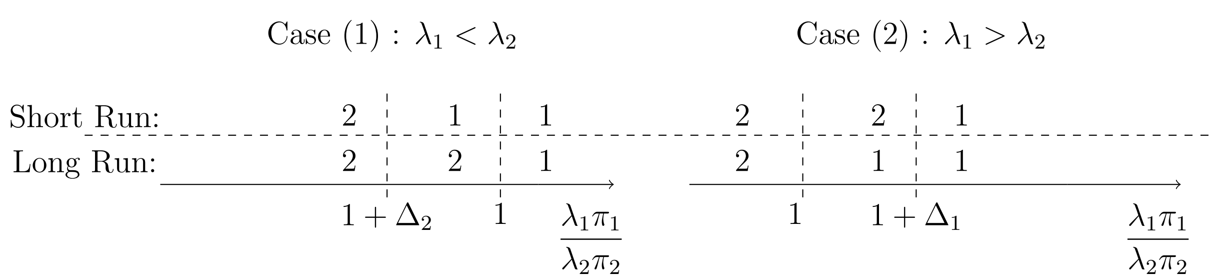

It may seem counterintuitive that short run players deviate to the harder project. The short run player puts no value on being able to work on a second project after the first invention is completed. When the long run player works on the easy project first, a deviation by a short run player to the hard project delays the total expected time until both projects are completed. Since the short run player receives no continuation value, he completely ignores the harm of delaying the completion of both projects. Note how extreme this effect is: short run firms can work on a project in equilibrium even when it has a strictly lower flow immediate payoff than the social optimum.

7.7 Laissez faire, patents and neutral prizes cannot be ranked

Laissez faire $\mathcal{P}_{LF}$, patents of various strengths $\mathcal{P}_\gamma$, and neutral prizes of various sizes $\mathcal{P}_q$ can each be preferred (in terms of efficiency) to the others depending on the nature of the parameter space. In other words, if the planner has only $\mathcal{P}_{LF}, \mathcal{P}_\gamma$ and $\mathcal{P}_q$ available as policy tools, and the planner's prior about the value of parameters is a correct point estimate, then the following cases show that there exist inventions graphs where each policy is preferred to the others.

Consider the following invention graph. Let the number of scientists $M = 1$, the number of firms $N = 2$ and the discount rate $r = 1$. In each of the following examples, the planner optimally works first on 1 then $A1$.

Figure 6: Ranking laissez faire, patents and prizes

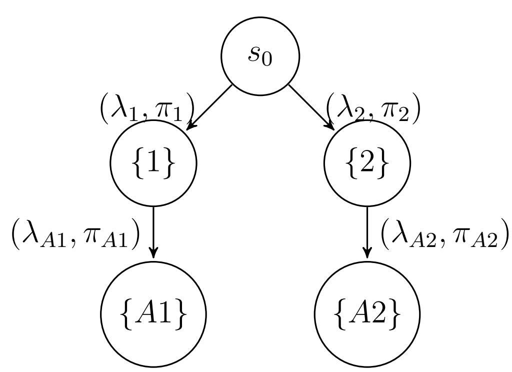

Case 1: $(\lambda_1, \pi_1) = (1, 1), (\lambda_2, \pi_2) = (2, 2), (\lambda_{A1}, \pi_{A1}) = (1, 16), (\lambda_{A2}, \pi_{A2}) = (1, 9)$. In this case, the laissez faire equilibrium is inefficient, and the equilibrium remains inefficient under patents of any strength or neutral prizes of any size.

Case 2: $(\lambda_1, \pi_1) = (1, 1), (\lambda_2, \pi_2) = (2, 2), (\lambda_{A1}, \pi_{A1}) = (1, 6), (\lambda_{A2}, \pi_{A2}) = (1, 2)$. In this case, the laissez faire equilibrium is inefficient, as is the equilibrium with any size prize, but the equilibrium with maximal patents is efficient.

Case 3: $(\lambda_1, \pi_1) = (3, 1), (\lambda_2, \pi_2) = (1, 8), (\lambda_{A1}, \pi_{A1}) = (1, 12), (\lambda_{A2}, \pi_{A2}) = (1, 10)$. In this case, the laissez faire equilibrium is inefficient, but efficiency is generated under sufficiently strong prizes or patents.

Case 4: $(\lambda_1, \pi_1) = (1, 2), (\lambda_2, \pi_2) = (2, 1), (\lambda_{A1}, \pi_{A1}) = (1, 14), (\lambda_{A2}, \pi_{A2}) = (1, 10)$. In this case, the laissez faire equilibrium is efficient, as is the equilibrium when prizes are small, but the equilibrium is inefficient under maximal patents or sufficiently large prizes.

Case 5: $(\lambda_1, \pi_1) = (1, 2), (\lambda_2, \pi_2) = (2, 1), (\lambda_{A1}, \pi_{A1}) = (1, 16), (\lambda_{A2}, \pi_{A2}) = (1, 10)$. In this case, the laissez faire equilibrium is efficient, and the equilibrium remains efficient under patents of any strength, but inefficiency arises as prizes grow sufficiently large.

Therefore, without knowing ex-ante what the parameter space will look like in a technological area, it is impossible to rank laissez faire, patents and neutral prizes in terms of their effectiveness at reducing directional inefficiency.

8 Online Appendix C: Existence of Equilibria, Mixed Equilibria and Multiplicity of Equilibria

In this appendix, we prove the existence of an equilibrium in our main model, and show the possibility of open sets of parameters with mixing equilibria, asymmetric equilibria and multiple equilibria. Note that since the planner optimum generically involves full effort on a unique invention, the existence of these alternative equilibria do not in any way change our efficiency results. For simplicity, we show all examples using laissez faire firm transfers.

8.1 Equilibrium Existence

Consider first the problem of existence. Since the invention graph is finite, we can use best responses to compute equilibria by backward induction. Consider the stage game and take the continuation values $V_i(s)$ as given. To prove equilibrium existence, we use the following result: a symmetric game whose strategy set S is a nonempty, convex, and compact subset of some Euclidean space, and whose utility functions $u(s_i, s1, \ldots, s_N)$, continuous in $(s_1, \ldots, s_N)$ and quasiconcave in s_i, has a symmetric pure-strategy equilibrium.[20]

Consider a formulation where the strategy space for each firm is the simplex $\Delta^{|S|}$. A firm's payoff, taking rival effort a_i as given, can be written as

$$u(x_i, x_{-i}) = \frac{\sum_{s' \in S} \alpha_{s'} x_{s'} + B(x_{-i,s'}, V_{is'})}{\sum_{s' \in S} \beta_{s'} x_{s'} + C(x_{-i,s'}, r)}.$$

We have continuous and quasiconcave payoffs in own strategy. Therefore there exists a symmetric pure equilibrium in the game. Uniqueness of the equilibrium is not guaranteed.

8.2 Mixing Equilibria

We say firms are *mixing* when they spread their scientists across multiple projects at a given time. By the usual mixed strategy condition, firms exert

[20] See, for example, Becker and Damianov (2006).

effort toward two different inventions only when these two inventions deliver the same payoff.

Let $f_s = w_s + V_{\mathcal{P}s}$. From the proof of Proposition 2, it is easy to see that a firm is indifferent between two states s' and ℓ iff

$$N(\lambda_{s'} f_{s'} - \lambda_\ell f_\ell) = \lambda s' f_{s'} \frac{(M\lambda_{s'} - M\lambda_\ell)}{r + M\lambda_{s'}} \quad \text{(Mix)}.$$

Obviously, when $\lambda_{s'} = \lambda_\ell$ and $f_{s'} = f_\ell$ condition (Mix) holds, because the inventions s' and ℓ are identical in terms of payoffs and simplicities.

Proposition 10. *Suppose inventions s' and ℓ are not identical. Condition (Mix)* **does not** *hold, i.e. there will be no mixing between s' and ℓ if*

1. $(f_{s'} - f_\ell)(\lambda_{s'} - \lambda_\ell) \geq 0$,
2. $(\lambda_{s'} - \lambda_\ell)(\lambda_{s'} f_{s'} - \lambda_\ell f_\ell) < 0$,

Proof. 1. Consider the first part of the proposition.

   (a) When $\lambda_{s'} = \lambda_\ell$ condition (Mix) reduces to $f_{s'} = f_\ell$. Therefore, if the inventions are not identical, there will be no mixing between s' and ℓ.

   (b) When $f_{s'} = f_\ell$ and $\lambda_{s'} \neq \lambda_\ell$ condition (Mix) reduces to $N = \frac{M\lambda_{s'}}{r+M\lambda_{s'}}$. Since $N > 1 > \frac{M\lambda_{s'}}{r+M\lambda_{s'}}$, this condition does not hold.

   (c) Condition (Mix) can be written as

$$N\left(1 - \frac{\lambda_\ell f_\ell}{\lambda_{s'} f_{s'}}\right) = \left(1 - \frac{\lambda_\ell}{\lambda_{s'}}\right)\frac{M\lambda s'}{r + \lambda s'}$$

   If $f_{s'} > f_\ell$ and $\lambda_{s'} > \lambda_\ell$, then since $N > 1$, $\frac{M\lambda s'}{r+\lambda s'} < 1$, $\lambda_\ell f_\ell < \lambda_{s'} f_{s'}$ and $\lambda_\ell < \lambda_{s'}$, then condition (Mix) cannot hold. Otherwise,

$$\left(1 - \frac{\lambda_\ell f_\ell}{\lambda_{s'} f_{s'}}\right) < N\left(1 - \frac{\lambda_\ell f_\ell}{\lambda_{s'} f_{s'}}\right) = \left(1 - \frac{\lambda_\ell f_\ell}{\lambda_{s'} f_{s'}}\right) < \left(1 - \frac{\lambda_\ell}{\lambda_{s'}}\right)$$

   implying $f_{s'} < f_\ell$, which is a contradiction. Similarly, if $f_{s'} < f_\ell$ and $\lambda_{s'} < \lambda_\ell$ we reach a contradiction.

2. In this case, the lhs of condition (Mix) is non positive and the rhs is strictly positive, and vice-versa.

$\square$

This proposition states that firms will never mix between states s' and ℓ if their simplicities are equal but one has higher payoff, or if their payoffs are the same but one is easier to discover than the other, or if one is easier and has higher payoff.

If one invention is easier and a second has a higher payoff inclusive of continuation value, then if firms best respond by mixing between the two, the flow payoff of the easier invention must be strictly higher than the flow payoff of the high payoff invention. In Figure 7, the gray area show inventions $(\lambda_{s'}, P_f(s'))$ that will never mix with the $(\lambda_{\bar{s}}, P_f(\bar{s}))$. This is all to say, large classes of invention graphs have no mixing equilibria.

Figure 7: Regions where simplicities and payoffs where firms will never mix with $(\lambda_{\bar{s}}, P_f(\bar{s}))$

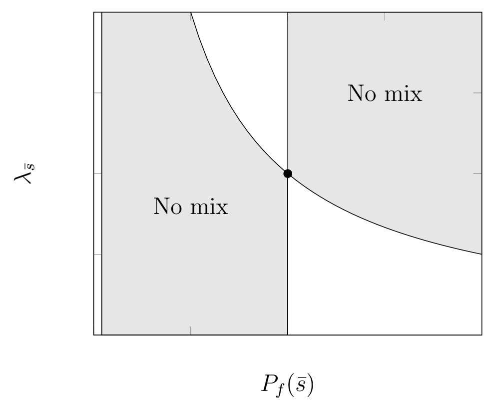

However, mixing equilibria can exist. It is easiest to see what causes them if we focus on states with no continuation value; in those cases, opponent actions only affect a firm through their cumulative discounted hazard rate, reflected in $\tilde{r} = Nr + N\sum_{z \in S(s)} a_{iz}$. Let $\tilde{r}_{min}$ correspond to all rivals exerting effort towards the hardest invention and $\tilde{r}_{max}$ the corresponding rate when all rivals work on the easiest invention. For any mixture we have $\tilde{r} \in [\tilde{r}_{min}, \tilde{r}_{max}]$.

A firm is indifferent between working on inventions k and ℓ iff

$$\frac{\lambda_k \pi_k}{\tilde{r} + \lambda_k} = \frac{\lambda_\ell \pi_\ell}{\tilde{r} + \lambda_\ell} \quad (MC)$$

Therefore if $\frac{\lambda_k \lambda_\ell (\pi_\ell - \pi_k)}{\lambda_k \pi_k - \lambda_\ell \pi_\ell} \in [\tilde{r}_{min}, \tilde{r}_{max}]$ there exists an (inefficient) symmetric mixing equilibrium. For example, if $\lambda_k = 4, \pi_k = 8, \lambda_\ell = 5, , \pi_\ell = 7, r =$

$1, N = 2$ and $M = 1$, then all firms exerting $1/3$ of the effort in k and $2/3$ in ℓ is a symmetric mixing equilibrium. By continuity, there is an open set of parameters values with these equilibria.

8.3 Asymmetric Equilibria

We can also construct an asymmetric equilibrium where firms are mixing. Let there be three inventions, and let $\tilde{r}_1$ and $\tilde{r}_2$ be the solutions to

$$\frac{\lambda_k \pi_k}{\tilde{r}_1 + \lambda_k} = \frac{\lambda_\ell \pi_\ell}{\tilde{r}_1 + \lambda_\ell} \quad \text{and} \quad \frac{\lambda_k \pi_k}{\tilde{r}_2 + \lambda_k} = \frac{\lambda_j \pi_j}{\tilde{r}_2 + \lambda_j}.$$

Let firm 1 mix between k and ℓ and firm 2 mix between k and j, accordingly. In this case, we also need to verify that firm 1 does not want to put effort towards j and firm 2 towards ℓ. For example, let $\lambda_k = 6, \pi_k = 3, \lambda_\ell = 12, \pi_\ell = 2, \lambda_j = 2, \pi_j = 6, r = 1, N = 2$ and $M = 1$. Here, firm 1 mixing between k and ℓ exerting $1/2$ of the effort in k, and firm 2 mixing between k and j exerting $1/3$ of the effort in j is an equilibrium.

8.4 Multiple Equilibria

There further exist small sets of parameters for which there exist multiple equilibria.

Proposition 11. *Consider only two inventions that are perfect substitutes. If $\lambda_k \neq \lambda_\ell$, then there is a region of parameters (π_k, π_ℓ) where there is multiplicity of equilibria with firms allocating effort only towards one invention.*

Proof. Let $M = 1$. All firms putting effort towards ℓ is a symmetric equilibrium if and only if

$$\frac{\lambda_\ell \pi_\ell}{\tilde{r}_\ell + \lambda_\ell} \geq \frac{\lambda_k \pi_k}{\tilde{r}_\ell + \lambda_k}$$

where $\tilde{r}_\ell = rN + (N-1)\lambda_\ell$.

Similarly, all firms putting effort towards k is a symmetric equilibrium iff

$$\frac{\lambda_k \pi_k}{\tilde{r}_k + \lambda_k} \geq \frac{\lambda_\ell \pi_\ell}{\tilde{r}_k + \lambda_\ell}$$

Combining the equations we obtain the inequalities, we obtain that both equilibria exist if and only if

$$\underbrace{\left(\frac{\lambda_\ell}{\lambda_k}\right)\frac{\tilde{r}_k+\lambda_k}{\tilde{r}_k+\lambda_\ell}}_{L_f} \leq \frac{\pi_k}{\pi_\ell} \leq \underbrace{\left(\frac{\lambda_\ell}{\lambda_k}\right)\frac{\tilde{r}_\ell+\lambda_k}{\tilde{r}_\ell+\lambda_\ell}}_{U_f}$$

Notice that $U_f - L_f = \frac{\lambda_\ell}{\lambda_k}(\lambda_k - \lambda_\ell)^2$. Also, we cannot have both $U_f > 1$ and $L_f < 1$.

The planner chooses invention k iff

$$\frac{\pi_k}{\pi_\ell} \geq \underbrace{\frac{\lambda_\ell}{\lambda_k}\left(\frac{r+\lambda_k}{r+\lambda_\ell}\right)}_{L_p}.$$

$\square$

Figure 8: Multiplicity of equilibria on perfect substitutes graph

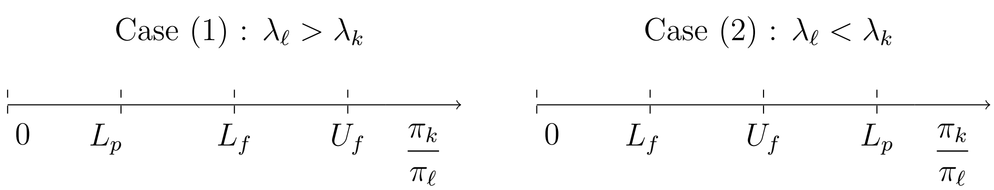

If the ratio $\frac{\pi_k}{\pi_\ell}$ is smaller than L_p (L_f), the planner (firm) works on invention ℓ. If the ratio $\frac{\pi_k}{\pi_\ell}$ is larger than L_p (U_f), the planner (firm) works on invention k. There are multiple firm equilibria if the ratio $\frac{\pi_k}{\pi_\ell}$ is in (L_f, U_f). The multiplicity is caused by the following tradeoff. If other firms are all working on the easy project, they are likely to make a discovery quicker than firm i deviating to the hard project. With perfect substitutes, if firm i does not discover first, it obtains a payoff of zero from the game. Although deviating can lead to a higher payoff conditional on succeeding first, the probability of being first is smaller. On the other hand, if all rivals are working on the hard project, the potential deviation is to work on an easy project with low payoff, foregoing the higher payoff of the harder project. When the ratios of payoffs and simplicities are structured such that $L_f \leq \frac{\pi_k}{\pi_\ell} \leq U_f$, it is both worth working on the hard project when everyone else does, and worth working on the easy project when everyone else does. As $N \to \infty$ we get $U_f \to L_f$ and the multiplicity disappears.