The Direction of Innovation

Kevin A. Bryan and Jorge Lemus*

April 4, 2017

Abstract

We construct a tractable and general model of the direction of innovation. Competition leads firms to pursue inefficient research lines both because they race toward easy yet less valuable projects, and because they work on difficult projects where they can appropriate a larger portion of a research line's value. Fixing this dual distortion requires policy to condition on properties of hypothetical inventions which are not discovered in equilibrium. Policies which do not do so, like patents and prizes, can therefore distort the direction of invention even if the aggregate quantity of R&D is optimal.

It has long been known that laissez faire markets may underproduce innovation due to indivisibilities, where the fixed cost of R&D is only fully paid by the initial inventor, and underappropriability, where research generates spillovers on subsequent inventions (e.g., Arrow (1962)). For this reason, existing theoretical work on innovation policy is largely

*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.

focused on evaluating how and when mechanisms like patents or research subsidies affect the quantity of R&D. However, firms do not simply choose how much research to perform, but also how to allocate their scientists across different research projects. 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, most critically, 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?

To understand when and what kind of directional inefficiencies arise, we construct an analytically tractable model of the direction of innovation. Firms are endowed with a set of researchers who can be allocated across a finite set of research projects. After an invention by any firm, a new set of potential research targets appears for all firms. We permit successful invention to affect the properties of future research targets, making future inventions easier (technological complementarities), harder (e.g., the invention reveals information about the difficulty of a research line), more valuable (market complements), or less valuable (market subsititutes). Researchers are fixed in number and can be costlessly deployed. By removing the intensive margin, we both isolate attention on the precise distortion caused by the existence of multiple research paths, and ensure that results hold *even if* the aggregate quantity of research is optimal.

Our model provides three main insights. First, two distinct classes of distortions simultaneously affect inventive direction in equilibrium: a "racing" and an "underappropriation" distortion. Second, policies designed to achieve the efficient rate of innovation, such as patents or subsidies, do not generally achieve the efficient direction of innovation, and can make directional distortions strictly worse. Third, directional inefficiency is generically a property of every innovation policy which both rewards inventors and does not condition on the characteristics 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 these results 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 undiscounted social value of their invention, and that once either invention is discovered, the marginal value of the other invention immediately falls to zero.[1] We make the latter assumption solely to provide stark intuition; as noted, the main model permits arbitrary links between inventions today and the value or difficulty of inventions available thereafter. 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.

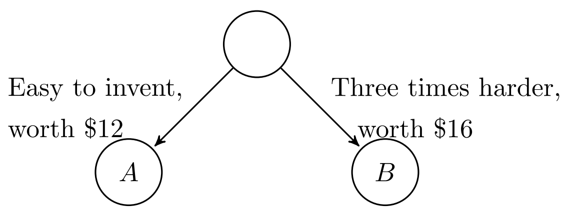

Figure 1: The racing distortion

If both firms work on A, $12 of value is created by its discovery, 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, this 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 account for the fact that its discoveries change rivals' continuation values. In order to profit from inventing B, a firm needs to invent it before other firms invent either A or B, an event likely to occur sooner when other firms work on the easier project A. In this example, the value of inventing B falls to zero after A is invented, but it should be clear (and will be formally shown) that the externality would be driven by less extreme changes in value, or in changes in the difficulty of future projects, or in changes in the existence or nonexistence of future research targets. Notice that this effect captures the intuition of the *racing distortion* of classic patent race models like Loury (1979) in a directional context, substituting the extensive margin of which

[1] For instance, let A and B be network goods.

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, again to keep intuition stark in these motivating examples, 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.

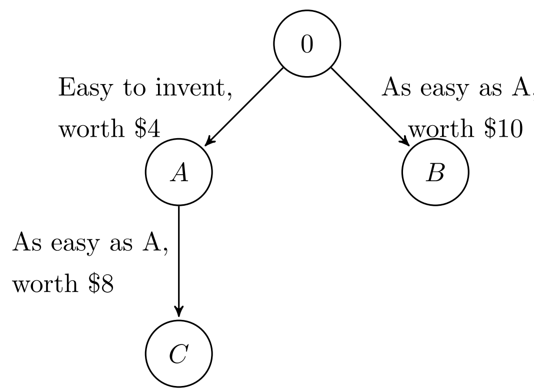

Figure 2: The underappropriation distortion

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. Inventors do not fully 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.

The fundamental problem of directionally efficiency is that both the underappropriation and the racing externalities may bind. Can policies like patents or prizes fix these distortions simultaneously? Patents fix underappropriation by causing firms to internalize the value of future inventions their work today makes possible, 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 the direction of invention 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. Correcting directional distortions using "automatic" policy will therefore be shown to be generically impossible.

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 equilibrium directional inefficiency is driven by a combination of racing and underappropriation distortions qualitatively similar to those in the examples above. Checking for an efficient firm equilibrium involves examining an inequality with a straightforward qualitative interpretation. In Section 3, we show that a baseline "laissez faire" policy, prizes, and patents of various strengths 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 also give two stylized applications of the theory to the questions of whether research is too incremental, and whether market size increases necessarily increase welfare via their effect on R&D. All proofs, and a number of generalizations, are left to the appendices.

Related Literature

Our results differ from the existing literature in focusing attention on cases when there are multiple projects available at any time, when simultaneous discovery does not occur, and when success on a project changes the nature of research targets available in the future.[2] That is, we study inefficiency 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)), 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)).[3] 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.

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

[3]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.

The most similar results to our model are found in three quite different papers. The research portfolio model in Dasgupta and Maskin (1987) investigates two firms choosing the correlation of their projects, showing in the main result that the market equilibrium involves projects with too much positive correlation in payoffs. In the baseline case, only the most valuable of the two inventions earns a reward, so when projects are highly correlated, the marginal value of the second invention is low. Firms do not properly account for the positive externality they provide to others by reducing that correlation in payoffs. Because firms only benefit from their own invention, they do not account for the societal "continuation value" generated by a later invention, hence they work in equilibrium on projects whose values are highly correlated.[4] This effect is the underappropriation distortion in our model, though we permit even the first inventor to redeploy their researchers to try for further inventions. We expand on Dasgupta and Maskin (1987) by allowing that redeployment, permitting arbitrary links between invention today and invention tomorrow, and allowing both the continuation value and the difficulty of inventions to vary. It is these features that permit a full investigation of our primary question: does there exist a simple policy under which decentralized research is efficient?

Second, Akcigit, Hanley and Serrano-Velarde (2013) present an endogenous growth model where firms can work on either basic or applied research. Their model is a steady state growth model which abstracts from our rich topology of research targets in order to study how innovation policy affects growth at large. Nonetheless, a main welfare result of theirs, that "neutral subsidies" like R&D tax credits operate by increasing the total amount of R&D without correcting the distortion toward applied research, holds in our model for a qualitatively identical reason.

Finally, Hopenhayn and Squintani (2016) construct an alternative model of pure innovation direction. In their basic model, there exist research lines with identical hazard rates of invention but varying payoffs. Researchers choose once how to allocate effort.

[4]In Dasgupta and Maskin (1987), following a tradition dating to R. Nelson in the 1950s, "parallelism need not imply waste" and hence diverse portfolios will be optimal unless "increasing returns to scale outweigh the benefits of diversification". However, we will show that in a properly dynamic model, the ability to work sequentially rather than simultaneously militates against diversity of research agendas. See the further discussion in Section 2.

Expected payoffs decline in the number of researchers working on a given project. A planner trades off getting high value payoffs quickly by using more researchers against causing those researchers to become idle once an invention is made on their research line. Competitive firms do not properly account for the fact that inventions they made will idle other researchers - or, in their more general model, cause other researchers to incur switching costs - and hence competitive firms are more likely to congest research areas with high expected payoffs. This result is similar in spirit to our racing externality: in competitive equilibrium, firms do account for the fact that their R&D affects the probability other firms succeed with a particular invention.[5] Our results permit a richer non-stationary distribution of inventions, leading to our main welfare result which depends critically on the interaction of the racing and underappropriation distortions.

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 first-order for government policy. Our results suggest fundamental limits on existing policy levers in generating directional efficiency.

1 A General Model of Direction of Innovation

Consider a 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: $\lambda : \Omega \times \Omega \to \mathbb{R}_+$, the *simplicity* to transition between any two states; and $\pi : \Omega \times \Omega \to$

[5]Though note that the relationship is only in spirit since the results in Hopenhayn and Squintani (2016) hold when all inventions are equally difficult. The reason is that they further assume either switching costs or diminishing returns to scale in the invention production function. These assumptions play the same role as "easy" inventions in our model in that they cause firms who pursue easy inventions (our model) or high value inventions (their model) to impose a negative externality.

$\mathbb{R}_+$, the *incremental immediate social payoff* from a transition.[6] We refer to the initial state of technology as s_0. To capture the idea of technology evolution—i.e., knowledge cannot be destroyed—we impose restrictions of λ such that (Ω, λ) defines an directed acyclic graph. For each state $s \in \Omega$, we define the set $S(s) \subseteq \Omega$ to be the set of all states than can be directly reached from s.

Definition 1. *An* invention graph *is a triplet* (Ω, λ, π) *where* (Ω, λ) *defines a directed acyclic graph with edges between states* s *and* s' *iff* $\lambda(s, s') > 0$.

The invention graph is common knowledge and all discoveries and inventive effort are publicly observed.[7]

As an example, consider the case of three inventions represented by the invention graph in Figure 3. The possible states of technology are given by $\Omega = \{s_0, \{1\}, \{2\}, \{1,2\}, \{1,3\}, \{1,2,3\}\}$. The transition simplicities 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. Our model allows for state contingent payoffs and simplicities. In Figure 3, the relationship between 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 market 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,

[6]The set of states may be countably infinite as long as π is bounded.

[7]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)).

research targets, or inventions.

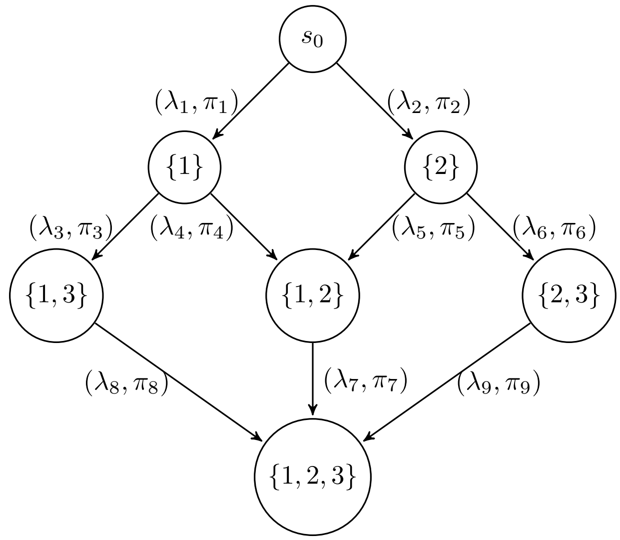

Figure 3: A simple invention graph.

It is important to note that we restrict to policies where payoffs π following inventions are one-time and fixed, rather than accrued as a flow. That is, there is no Arrow-style replacement effect where firms garnering flow rents from earlier inventions act asymmetrically from firms with no such flow rents. There is a methodological tradeoff if analytical tractability is desired: either the nature of links between past inventions and present opportunities must be highly restricted (for instance, to independent quality ladders) and steady state assumptions must be imposed (e.g., Chen, Pan and Zhang (2016)), or replacement effects must be jettisoned in exchange for richer possible topologies of present and future invention targets. We choose the latter as it more clearly isolates the externalities generated by direction choice, and because it permits easier comparison the microeconomic literature on patent races and sequential invention.

1.1 Discovery 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. At each state, a firm chooses how to allocate its research among the set of feasible research targets in $S(s)$. 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'.[8] Research is costless, so the problem 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.[9] 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: Efficient Allocation of Research

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 allocated researchers across states according to $x = (x_{s'})_{\{s' \in S(s)\}}$ 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

[8] 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.

[9] In Online Appendix B, 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.

along that line $(\pi_{s'}+V_{ps'})$. Simplifying the expression above, the social planner problem is to solve the recursive maximization problem:[10]

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

1.3 Firm Problem: Competitive Allocation of Research

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 $\mathcal{P}$ which specifies the payoffs each firm receives following any invention.

Definition 2. *A* transfer policy $\mathcal{P}$ *is a triplet* (Ω, w, z) *where* $w(s,s') : \Omega\times\Omega \to \mathbb{R}$ *is the transfer received by inventors and* $z(s,s') : \Omega\times\Omega\to\mathbb{R}$ *is the transfer received by noninventors following the invention which transitions the state from* s *to* s'*.*

Let $a_{-is'} = \sum_{j\neq i} x_{js'}$ be 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

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

where $V_{i\mathcal{P}s'}$ is the equilibrium continuation value for firm i following a transition to $s'\in S(s)$ given policy $\mathcal{P}$. Forcing the continuation value to be identical for inventing and noninventing firms is without loss of generality since w and z are unrestricted in sign. Simplifying this expression, firms solve:

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

Once any firm discovers some invention $s'\in S$, all firms reallocate effort across the new set of potential research targets $S(s')$. We restrict to Markov Perfect Equilibria.

[10]Blackwell's sufficiency condition guarantees the problem is well-defined as long as the payoff function $\pi_{s'}$ is bounded and the number of reachable inventions at a given state is finite (i.e., $|S(s)|<\infty$).

1.4 Common Transfer Policies

We permit general transfer policies $\mathcal{P}$, but four classes will be of special relevance: a baseline case, patents, neutral prizes, and information-constrained policies.

In our baseline case, inventors receive the full immediate social payoff π of their invention, but the continuation value accrues equally to all firms. This baseline can be interpreted either as a narrow patent which does not cover future inventions but permits the inventors to garner the full immediate incremental surplus, or as the "laissez faire" value of first mover advantage or other non-IP market power. Importantly, note that since π enters the firm value function linearly, a straightforward induction argument shows that if all immediate firm payoffs π are scaled by a common factor, the firm problem is unchanged, hence the baseline case can alternatively be interpreted as any innovation policy giving inventors a fixed percentage of the incremental social surplus.

Definition 3. *The transfer policy $\mathcal{P}_{BC}$ is* the baseline case *if inventing firms receive the full immediate social payoff of their invention, i.e. $w(s,s') = \pi(s,s')$, and if noninventing firms receive no immediate transfer following an invention but are equally able to build on today's invention ($z(s,s') = 0$).*

We model broader 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.[11] 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 the baseline case.

Definition 4. *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 $\gamma \in [0,1]$.*

We model a neutral prize q in state s as 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

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

$V_{i\mathcal{P}s'}$. 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." Restricting the definition of neutral prizes such that only inventions beyond a technological threshold garner prizes, where that unmodeled threshold does not completely depend on λ and π, will not change any of our welfare results.[12] 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 5. *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 the baseline case are all policies which do not condition transfers $w(s, s')$ and $z(s, s')$ on off-equilibrium-path parameters: the transfers following the invention s' do 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*. 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 6. *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 $\ell \neq s'$. In particular, for any two invention graphs, for all state transitions (s, s') where $\lambda(s, s'), \pi(s, s')$ and $V(s, s')$ are identical in both graphs, the policy $\mathcal{P}$ must assign the same transfers $w(s, s')$ and $z(s, s')$.*

[12] We return to this point in the discussion in Section 3.

2 Planner Optimum and Firm Equilibrium

In this section we explore when and why competing firms choose an inefficient research direction. To simplify the difference between the planner optimum and firm equilibrium, it will be useful to define $\tilde{r} = Nr + M(N-1)\lambda_{s'}$, a virtual discount rate for competing firms, and Θ, the expected discounted continuation value for firm i when it does not contribute to research until the next invention is completed by some other 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 optimum 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 ($r \neq \tilde{r}$).

Note that the planner and firm best response indices generically have unique maxima where full effort is exerted on a single project. It may seem surprising that there is no mixing across projects, particularly from the planner. However, we argue that there is a frequent misunderstanding in innovation research as to why society should pursue a diverse R&D portfolio. The intuition that a diverse research agenda provides "more lottery tickets" or "real options" (e.g., Nelson (1959)) is false in the absence of decreasing returns to scale in the research production function, since simultaneous diversified research with constant returns to scale can be replicated by sequentially exploration of those same projects. Sequential exploration allows a higher level of effort to be exerted first on projects which are believed to be more valuable, hence discounted payoffs increase.

For instance, in the canonical model of Dasgupta and Maskin (1987), the planner diversifies among projects whose values are not highly correlated since only the most valuable project that is eventually invented has any value. An alternative to working on two projects simultaneously, however, is to put full effort on whichever project has the highest flow payoffs inclusive of expected continuation value, earn that payoff more quickly, then switch to full effort on whichever remaining project has the highest expected flow payoffs. As Dasgupta and Maskin correctly show, firms insufficiently account for the continuation value their inventions create. But this has no direct implication for the diversity of the research portfolio at a single point on time, which requires a more direct rationale such as certain types of switching costs or decreasing returns to scale. This intuition will be familiar to some readers from the mathematics of the Gittins index in bandit problems—there can absolutely be differences between the social optimum and firm equilibrium *over time* in the frequency with which various arms are pulled, but at any single point in time there is an optimal arm to pull (Keller, Rady and Cripps (2005)). Of course, research portfolios are an empirical regularity, a fact we do not deny. We merely claim that competition per se is not the fundamental driver of this diversification, as opposed to the shape of R&D cost functions or projections functions.[13]

In addition to this economic interpretation, there is a geometric interpretation of Proposition 1 which may prove useful to innovation and growth theorists. Note that the plan-

[13] It is shown in Online Appendices B and C in variants 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.

ner and firm problems are both linear functionals. The Charnes-Cooper transformation converts linear functionals with linear inequality constraints into standard linear programs, whose maximands are generically the extreme points of a polyhedron (Charnes and Cooper (1962)). The proof of Proposition 1 notes that the extreme points of the transformed linear program are precisely the allocations where full effort is exerted on a single invention. Therefore, finding the planner optimum or firm best response involves comparing indices related to each of those allocations, and the qualitative difference between the planner and firms involves comparing the conditions under which the planner-optimal project maximizes the firm best response index when all other firms are exerting full effort on the optimal project. Such clean analytic tractability will not hold if the research production function is concave. In that case, we can show via a first order approach that the causes of firm inefficiency are qualitatively similar to the case where the research production function has constant returns to scale, but a completely analytical characterization of most of our results becomes impossible; see Online Appendix B for further details.

2.1 Inefficiency Direction of Innovation by the Firms

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 find the set of Markov Perfect Equilibria.[14]

Definition 7. *The direction of innovation of firms in state s is* efficient under policy $\mathcal{P}$ *if there exists a Markov perfect equilibrium involving full effort from all firms towards the planner's optimal direction $s' \in S(s)$.*

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

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

[14] We therefore rule out equilibria where firms collude and punish each other across states. Online Appendix C discusses the existence, multiplicity, and potential asymmetry of firm equilibria. We note in particular that equilibrium existence with a infinite invention graph is restricted to $\epsilon-$equilbria.

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.*

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 is a simple restatment of the planner optimum in Proposition 1 in terms of discounted flow payoffs. The second part 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 the socially optimal project s' to some other project ℓ receives a higher portion of the total social value of the invention. This is immediate since P_p is the social value (or planner payoff) and $w + V_{i\mathcal{P}}$ is the immediate payoff to firms plus their continuation value. $D_2(s', \ell)$, 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.[15] If s' is the planner optimum, any

[15] Recall that $\Delta(s', \ell) > 0$ if and only if invention ℓ is easier than s'.

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 in our setting can generate inefficiency in only three ways: either firms are overincentivized to race toward projects easier than the planner preferred ones; or 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 research 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.[16]

Consider, for example, the baseline 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 the baseline 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 baseline policy $\mathcal{P}_{BC}$:*

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

The firm equilibrium condition in Proposition 2 collapses to:

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

[16] 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.

Corollary 1 says that under the baseline, 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 $\lambda\pi$ than any potential deviation.

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

1. *If planner-optimal s' is not an equilibrium under the baseline policy 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 baseline policy firm equilibrium is inefficient.*

3 General Policy Solutions to Directional Inefficiency

In the previous section, we saw that the baseline case 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 the baseline case can each be more efficient than the other two, 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

Consider the neutral prize policy $\mathcal{P}_q$ which 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*

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

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

Neutral prizes of any size do not guarantee efficiency, and indeed can generate an inefficient outcome even when the baseline case is efficient. Note that neutral prizes still generate the underappropriation distortion of the baseline, 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 baseline 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.

A reasonable objection is that prizes ought be awarded only to inventions which meet some "standard", rather than to the first inventor of even low-value steps. Empirically, many prizes do not in fact possess such a restriction: for instance, scientific credit really does go to the first person to cross a technological threshold regardless of the market value or ease of building on that invention. However, consider prizes restricted either to inventions whose total surplus $\pi + V_p$ exceed a threshold, or to those that are sufficiently difficult for λ to be below some threshold. It is straightforward to show that, for a sufficiently large prize, there will be no effort on any project which does not clear one of those thresholds. However, the planner optimal project is that which maximizes the index $\frac{\lambda_s V_{ps}}{r+\lambda_s}$. Therefore, a large prize only given to very high value projects may distort effort away from a socially optimal project which is lower value yet easier, and a large prize given only to difficult inventions may distort effort away from socially optimal projects which are both higher value *and* easier. Incentivizing the "right" project involves giving prizes only to inventions which maximize the planner optimal index, but knowing which inventions those are requires conditioning the prize for an invention (λ_s, π_s) on the properties of other potential inventions, which may be particularly burdensome in terms of the information required by the prize-granting agency. We return to this point below when discussing general information-constrained policies.

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*

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

where

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

Suppose that invention in all future states is efficient ($V_{i\mathcal{P}_{BC}s} = V_{i\mathcal{P}_\gamma s} = \frac{V_{ps}}{N}, \forall s$), hence $\mathcal{V}(\gamma) = 0$. In this case, patents of maximal strength $\gamma = 1$ exactly cancel out the baseline underappropriation distortion $D_1(s', \ell)$, as might be expected. Patents, however, do not affect the baseline racing distortion $D_3(s', \bar{s})$. If the underappropriation distortion under the baseline policy is helping counteract the racing distortion—e.g., if firms are not deviating toward an inefficient easy project under the baseline policy 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 the baseline policy 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." Does there exist *any* policy which can generate directional inefficiency without conditioning on the parameters of off-path inventions?

Recall from Definition 6 that a policy is information-constrained if transfers do not condition on the parameters of inventions which are not ever invented in equilibrium. That is, if in state s there is an invention s' with difficulty $\lambda_{s'}$, immediate payoff $\pi_{s'}$, and social continuation value $V_{ps'}$, an information-constrained policy must give the same transfers $w(s, s')$ and $z(s, s')$ following the invention of s' regardless of the properties of potential inventions which are not invented in equilibrium.

Proposition 4 shows that any efficient information-constrained policy must reward inventing firms at least as much as non-inventing firms. 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 noninventors 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.[17] For many reasons aside from directional efficiency, however, we may wish to rule out policies that reward inventors and noninventors 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 intuition of the second part is more subtle, but is fundamentally related to the fact that the optimal policy response depends on whether the racing or underappropriation distortion dominates in the baseline case, and that the distinction between the two involves comparing the optimal invention to alternatives which are not invented in equilibrium. 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 yet more lucrative, or easier yet less lucrative, than an alternative ℓ. If s' is harder than ℓ, the racing distortion implies that firms will only work on s' in equilibrium if the relative transfer to inventors of s' compared to ℓ, inclusive of continuation value is strictly higher than the relative social payoff of s' compared to ℓ (e.g., $\frac{w_{s'}+V_{i\mathcal{P}s'}}{w_\ell+V_{i\mathcal{P}\ell}} > \frac{P_{ps'}}{P_{p\ell}}$). Likewise, if ℓ is harder than s', then the relative total payment to inventors of ℓ compared to s' must strictly exceed the relative social payoff of ℓ compared to s'. Note that only

[17]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.

one of the strict inequalities can hold for any given transfer policy $\mathcal{P}$. Note also that an arbitrarily small change in payoffs or simplicities will make the planner optimum unique without changing the nature of those strict inequalities.

The reader may wonder whether efficient policy can be restored by restricting the nature of the invention graph. It should be clear from the nature of the proof that the fundamental problem with information-constrained policy when it comes to directional distortion is that only by comparing inventions to their potential alternatives can a planner know whether firms race toward the easy-yet-unlucrative or whether they shifted toward difficult projects where most of the value is immediate rather than part of a continuation value potentially captured by other firms. That said, not every type of invention graph leads to inefficiency in the baseline case. In Sections 6.3 and 6.4 of Online Appendix B, we prove that directional inefficiency requires both the existence of multiple research targets in some states or 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, Section 6.3 shows that the baseline 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, as shown in Section 6.4 of the Online Appendix, 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 baseline firm equilibrium is no longer efficient in general.

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 policy on the full social value including continuation value of inventions when setting transfers, rather than on proxies such as the monopolist rents earned by an inventor.

3.4 Two Simple Applications

Though our primary results are theoretical, the model is general enough to apply to a broad set of policy problems, of which we consider two in a highly stylized form. First, we ask whether decentralized invention is excessively incremental. Second, we ask whether expansions in market size will increase the social welfare generated by the research sector. In both cases, we answer in the negative.

3.4.1 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. Let payoffs be baseline payoffs: inventors receive the full immediate social payoff of their invention, and noninventors nothing, but any firm can symmetrically build on any invention once it has been invented.*

1. *If the planner is indifferent between the incremental line and the radical line then the incremental line (radical line) is an 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 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.

Why do patents induce excessively incremental invention? 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.

3.5 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.[18] 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.

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

Assume that an expansion of trade increases the immediate social payoff π to all inventions by a constant factor ζ. Following the discussion at the start of Subsection 1.4, neither planner nor firm researcher allocation changes when all payoffs are equally scaled, as long as the number of firms N is held constant. 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 Conclusion

We provide three novel contributions to the study of innovation incentives.

First, we construct a dynamic model of the direction of innovation, with unrestricted links between past invention and the properties of future research targets. We show that it is possible to transform the planner and firm problems into linear programs, allowing us to characterize their maximands as a simple indices which can be tractably analyzed. Although this model rewards inventors only at the time of invention, and hence abstracts away from the replacement effect of quality ladder models, it does permit the investigation of invention topologies with much more complex links between past and future. We provide simple qualitative applications of our model, but note that the tractability of the model opens the door for fruitful further applied research on how policies like trade expansion affect the direction of innovation.

Second, we show that firms allocate their scientists inefficiently in equilibrium due 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. While in models of the rate of inventions, stronger patents and larger prizes unambiguously increase how much R&D is performed, in a model of direction strengthening patents or prizes can either push direction closer to or further from the optimum.

Third, there is a fundamental limit to decentralized "autopilot" innovation policy. Any policy which both rewards inventors and conditions these rewards 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, when some of the benefits of invention spill over to rival firms, and when some firms are "short-term" firms which live only one period. 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.

Chen, Yongmin, Shiyuan Pan, and Tianle Zhang. 2016. "Patentability, R&D direction, and cumulative innovation."

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.

Fudenberg, Drew, and David Levine. 1986. "Limit games and limit equilibria." *Journal of Economic Theory*, 38(2): 261–279.

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."

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.

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

Nelson, Richard R. 1959. "The Simple Economics of Basic Scientific Research." *The Journal of Political Economy*, 67: 297–306.

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.*

5 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.

5.1 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

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

5.2 Proof of Proposition 1

5.2.1 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 $h(x)$).[19]

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

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

[19] 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.

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 $v(s,s') = y(s,s')\left(1 + \frac{M}{r}\lambda(s,s')\right)$ 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:

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

5.2.2 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

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

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:

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

5.3 Proof of Proposition 2

5.3.1 Part 1: The Planner Optimum with Δ

The condition for the planner optimum is:

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

5.3.2 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:

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

5.4 Proof of Corollary 2 and 3

Let $D^*_{BC}(s',\ell) = D_1(s',\ell) + D_2(s',\ell) + D_3(s',\ell)$, the distortion under the baseline case. Straightforward algebra applied to Proposition 2 shows that a prize q changes $D^*_{BC}(s',\ell)$ 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.

5.5 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 $w_\ell = z_\ell = \pi_\ell$. If inventive effort in future states 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. Note that scaling all payoffs by a constant does not change the firm best response, hence any equal payments to inventors and noninventors which are a scale multiple of the immediate social payoff of the invention will induce the first best. That is, if we do not require inventors to be rewarded more than noninventors, there is a trivially, arbitrarily cheap method of inducing efficiency.

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.

5.6 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.

5.7 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.

6 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.

6.1 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 5.2.1 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.

6.2 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 6.1, 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.

6.3 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 in the baseline case does not result from the simple existence of multiple projects. Rather, in order to generate inefficiency in the baseline case, 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 8. *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, the baseline firm equilibrium 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 *firm equilibrium under the baseline policy.*

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. □

6.4 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 the baseline policy. 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 9. *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 $\pi = 0$.*[20]

[20]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.

With the social continuation value equal to zero, and inventing firms paid exactly the immediate social payoff of their invention, the baseline policy on the perfect substitutes invention graph generates distortions $D_1(s',\ell) = D_3(s',\ell) = 0$, 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 baseline policy $\mathcal{P}_{BC}$ 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 the baseline policy 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 directional inefficiencies under the baseline policy.

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.*

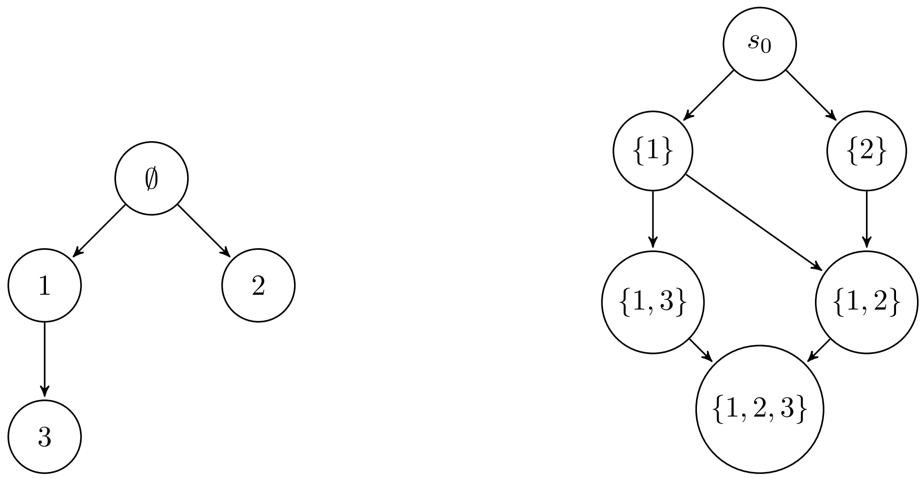

(a) Possible Inventions

(b) Graph of states.

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

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*

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

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 equilibrium under the baseline policy.

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. □

6.5 Spillovers

In the main results, under the baseline policy, 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 10. *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

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

where $\mathcal{V}(\alpha)$ is the distortion from the difference between the social continuation value under the baseline 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 baseline policy.

6.6 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.

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

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

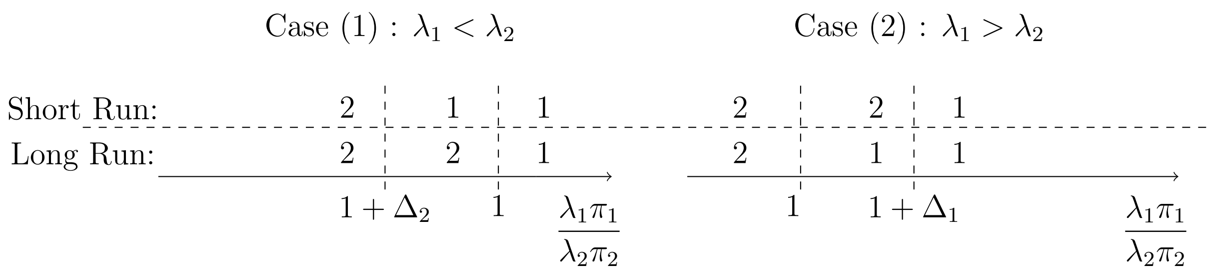

both projects are completed. Since the short run player receives no continuation value, he completely ignores the harm of delaying the completion of both projects. Note how extreme this effect is: short run firms can work on a project in equilibrium even when it has a strictly lower flow immediate payoff than the social optimum.

6.7 Laissez faire, patents and neutral prizes cannot be ranked

Laissez faire $\mathcal{P}_{BC}$, 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 $\mathcal{P}_{BC}, \mathcal{P}_\gamma$ 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 $A1$.

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 baseline 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 baseline faire equilibrium is inefficient, as is the equilibrium with any size

Figure 6: Ranking the baseline, patents and prizes

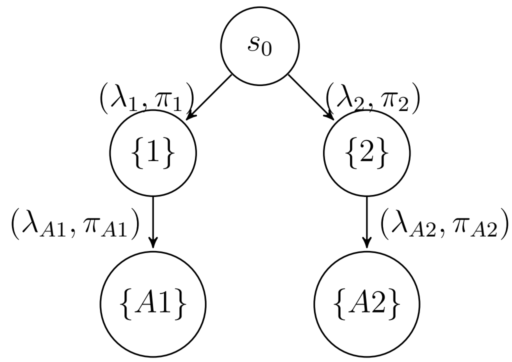

prize, but the equilibrium with maximal patents is efficient.

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 baseline 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 baseline 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 baseline 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 the baseline case, patents and neutral prizes in terms of their effectiveness at reducing directional inefficiency.

7 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 baseline firm transfers.

7.1 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.[21]

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.

Although much of our analysis is for a finite graph, we can extend the model to allow for infinite graphs as long as payoffs are bounded. Consider the finite truncation of an infinite invention graph with only T discoveries from the present state (note that the

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

invention graph is a directed acyclic graph by assumption, so this truncation is the tree beginning in the initial state with every branch having length T or less). In this finite truncation, we have already shown equilibria exist. If payoffs are uniformly bounded and there is discounting, then, for T large enough such that the maximal discounted continuation payoff after the T-th discovery is smaller than ε, any equilibria in the finite game with T inventions will also be part of an $\epsilon-$equilibria in the infinite game via the result in Fudenberg and Levine (1986).

7.2 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 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.

□

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.

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

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

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

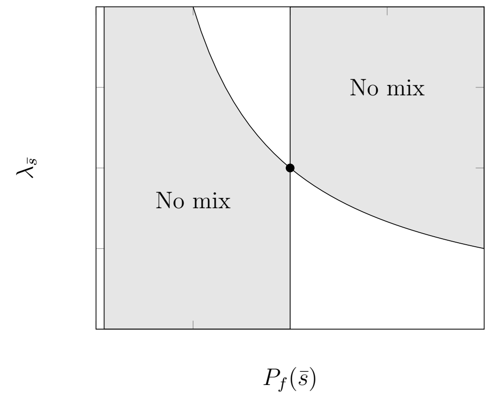

Therefore if $\frac{\lambda_k\lambda_\ell(\pi_\ell-\pi_k)}{\lambda_k\pi_k-\lambda_\ell\pi_\ell} \in [\tilde{r}_{min}, \tilde{r}_{max}]$ there exists an (inefficient) symmetric mixing equilibrium. For example, if $\lambda_k = 4, \pi_k = 8, \lambda_\ell = 5, , \pi_\ell = 7, r = 1, N = 2$ and $M = 1$, then 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.

7.3 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.

7.4 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

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

□

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

Figure 8: Multiplicity of equilibria on perfect substitutes graph

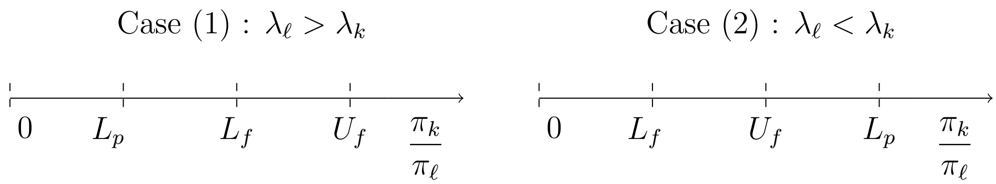

perfect substitutes, if firm i does not discover first, it obtains a payoff of zero from the game. Although deviating can lead to a higher payoff conditional on succeeding first, the probability of being first is smaller. On the other hand, if all rivals are working on the hard project, the potential 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.