

On the Direction of Innovation

Hugo Hopenhayn

University of California, Los Angeles

Francesco Squintani

University of Warwick

How are resources allocated across different R&D areas (i.e., problems to be solved)? As a result of dynamic congestion externalities, the competitive market allocates excessive resources into those of high return, being those with higher private (and social) payoffs. Good problems are tackled too soon, and as a result the distribution of open research problems in the socially optimal solution stochastically dominates that of the competitive equilibrium. A severe form of rent dissipation occurs in the latter, where the total value of R&D activity equals the value of allocating all resources to the least valuable problem solved. Resulting losses can be substantial.

I. Introduction

Innovation resources are quite unequally distributed across different research areas. This is true not only in the case of commercial innovations but also in our own fields of research. Some areas become more fashionable (“hot”) than others and attract more attention. A quick look at the distribution of patenting by different classes since the 1980s reveals significant changes in the distribution of patent applications: while early on the leading sector was the chemical industry followed closely by others, starting in 1995 the areas of computing and electronics surpassed all other areas in patent filings by an order of magnitude. The so-called dot-com

Hopenhayn acknowledges support from National Science Foundation grant 1757134.

Electronically published May 21, 2021

[*Journal of Political Economy*, 2021, vol. 129, no. 7]

© 2021 by The University of Chicago. All rights reserved. 0022-3808/2021/12907-0002\$10.00

bubble is an example of what many considered excessive concentration in the related field of internet startups. This example suggests that innovation resources might be misallocated across different areas and perhaps too concentrated on some, yet to date almost no economic theory has been devoted to this question.

This shift in innovative activity is likely the result of technological, demographic, and other changes that introduce new sets of opportunities to exploit and problems to solve. As new opportunities arise, firms compete by allocating innovation resources across these opportunities, solving new open problems and thus creating value. The process continues as new opportunities and problems arise over time and innovation resources get reallocated. We model this process and characterize the competitive equilibrium, as well as the socially optimal allocations. Our main finding is that the market disproportionately allocates researchers to hot R&D lines—characterized by higher expected rates of return per unit of research input—and leads to an excessive turnover of researchers.¹

We model this process as follows. At any point in time, there is a set of open problems (research opportunities) that, upon being solved, generate some social and private value v . This value is known at the time that research inputs are allocated and is the main source of heterogeneity in the model. The research side of the economy is as follows. There is a fixed endowment—inelastically supplied—of a research input to be allocated across problems, which for simplicity we call researchers. The innovation technology specifies probabilities of discovery (i.e., problem solution) as a function of the number of researchers involved. Ex ante, the expected value of solving a problem is split equally among the researchers engaged, consistently with a winner-takes-all rule as in patent races or with an equal sharing rule. Once a problem is solved, the researchers involved are reallocated to other problems at some cost. We consider both an environment where the set of problems is fixed and a steady state with exogenous arrival of new problems. Firms compete by allocating researchers to the alternative research opportunities to maximize value per unit input. As there is a large number of firms, we can equivalently assume that each researcher maximizes her value by choosing a research line. As a result, the value of joining any active research line is equalized.

The key source driving market inefficiency is differential rent dissipation resulting from competitive entry into research. This is due to the pecuniary externality imposed by a marginal entrant to all others involved in her research line. It is useful to contrast our results to models of patent races where there is a perfectly elastic supply of potential entrants in

¹ Hot R&D lines need not correspond to high-value innovations, because high value may often be associated with a low probability of success. Hot R&D may take the form of incremental innovations in a highly fertile R&D area.

the race and competitive forces drive average value down to the entry cost. With a concave discovery function, the average value exceeds the marginal value of an entrant, thus resulting in excessive entry. The gap between the average and the marginal value is a reflection of the fact that part of the return to an entrant comes from a decrease in the expected returns of the remaining participants, the pecuniary externality.

In contrast, in our model we assume that the total research endowment to be allocated is inelastically supplied and that entry costs are the same across all research lines, so there cannot be excessive entry overall. But as we find, there will be excessive entry in some areas and too little in others, as well as excessive turnover.

It is useful to divide the sources of this misallocation between static and dynamic ones. The static source of misallocation arises as the pecuniary externality changes with the number of researchers in a research line. To illustrate this, consider the case where the probability of innovation is linear up to a certain number of researchers \bar{m} and constant thereafter and there are two research lines: a hot one with high value and one with low value. Furthermore, suppose that given the total endowment of researchers, more than \bar{m} enter into the former while fewer than \bar{m} enter into the latter one, so the average values are equalized. It immediately follows that there is excessive entry into the hot area, where there are negative pecuniary externalities, and too little in the low-value one, where there are none.

This example suggests that the extent of pecuniary externalities can vary with scale and will do so in general. As total discovery probability is bounded, the results described in the example will occur in some parameter region, and as a result there will be excessive entry into the higher-value R&D areas. As we show in the paper, this distortion holds globally (there is excessive entry above a value threshold and too little below) for a canonical model of innovation considered in the literature.

We now turn to the dynamic sources of misallocation that can be orders of magnitude more important, as illustrated by our back-of-the-envelope calculations. The first dynamic source of misallocation arises from the cost of reallocation. When a researcher joins a research line and succeeds, this generates a capital loss to the remaining researchers who must incur a new entry cost to switch to a new, equally valuable research line. This externality grows with the number of researchers affected and thus with the value v of innovation, leading to excessive entry into hot areas. The second source is more subtle. As a consequence of rent dissipation, the value of entering any innovation line is equalized in the competitive equilibrium. In the eyes of competitors, there is no distinction between different open problems in the future, as they all give the same value. In contrast, a planner recognizes that better problems (i.e., those with higher v) have higher residual value and thus carry a higher future option value if they are not immediately solved; the planner is less rushed to solve them.

We analyze a steady-state allocation with an exogenous arrival of new problems and endogenous exit of existing ones resulting from the allocation of researchers. These two forces determine a stationary distribution for open problems. High-value problems are solved faster in the competitive equilibrium, owing to the biases indicated above, so the corresponding stationary distribution has a smaller fraction of good open problems. In addition, as the distribution of innovators is more skewed than in the optimal allocation, turnover is higher and so are reallocation costs. This leads to a severe form of rent dissipation, where in a competitive equilibrium the total value of R&D activity equals the value of allocating all resources to the least valuable problem solved. The magnitude of this distortion can be extremely severe, leading to very large welfare effects, as shown by our simple calculations.

Throughout our analysis, we assume that the private and social value of innovations is the same across research lines or, likewise, that the ratio of private to total value is identical. We do this to abstract from some other important but more obvious sources of misallocation. As patents attempt to align private incentives with social value, they are of no use in solving the distortions that we consider. The source of market failure in our model is the absence of property rights on problems to be solved, which are the source of R&D value. Patents and intellectual property are no direct solutions to this problem, as they entitle innovators to value once problems have been solved. Our research suggests that there might be an important role for the allocation of property rights at an earlier stage.²

The paper is organized as follows. The related literature is discussed in section II. Section III provides a simple example to illustrate the main ideas in the paper. Section IV describes the model and analyzes the static forces of misallocation. Section V considers the reallocation of researchers and the dynamic sources of misallocation for a fixed set of problems. Section VI considers the steady state with continuous arrival of new problems. Section VII concludes.

II. Literature Review

Early literature (e.g., Schumpeter [1911] 1934; Nelson 1959; Arrow 1962) pointed at limited appropriability of the innovations' social value

² As for policy considerations, our finding suggests the desirability of non-market-based incentives that rebalance remuneration across R&D lines, so as to subsidize R&D lines with less profitable or less feasible innovations. Existing R&D funding mechanisms include research grants, fiscal incentives on innovations or ongoing research, research prizes, and procurement. While often state-funded, R&D subsidization can also be funded by private consortia or donors (especially when taking the form of research grants and prizes), and the tenure system in academic institutions also entails R&D subsidization. Because subsidies can be at least partially funded with levies collected on patent monopoly profits, the kind of policy intervention suggested here contains elements of cross subsidization across R&D areas.

by innovators and at limited access to finances as the main distorting forces in R&D markets, both leading to the implication that market investment in R&D is insufficient relative to first best.³ A large academic literature has developed to provide policy remedies, often advocating strong innovation protection rights and the subsidization of R&D, trading off against the distortions resulting from market power.⁴

Another known source of market inefficiency is caused by the sequential, cumulative nature of innovations. This so-called sequential spillover problem arises when, without a “first” innovation, the idea for follow-on innovations cannot exist and the follow-on innovators are distinct from the first innovator (see Horstmann, MacDonald, and Slivinski 1985; Scotchmer 1991). An innovation in the sequence will typically reduce the rents of previous innovators (hence having a negative competitive effect) and contribute to the value of future ones, as well as that of consumers (hence having a positive spillover effect). In the absence of direct transfers, patent-like mechanisms can be used to trade off innovations at different points in the ladder.⁵

More generally, the misalignment between an innovation’s private and social returns is the result of negative competitive effects (“rent stealing”) and positive spillover effects to other innovators, firms, and consumers.⁶ Bloom, Schankerman, and Van Reenen (2013) show that technological positive spillovers tend to dominate, and as a result, social rates of return are twice as high compared with private ones. This implies that on average innovations are underprovided by the market, and this has been the driving concern of innovation policy discussions.

³ According to Bloom, Schankerman, and Van Reenen (2013), the social return to innovations is estimated to be twice as large as the private returns to innovators. Evidence of a funding gap for investment innovation has been documented by, e.g., Hall and Lerner (2010), especially in countries where public equity markets for venture capitalist exit are not highly developed.

⁴ Wright (1983) compares patents, prizes, and procurement as three alternative mechanisms to fund R&D. Patents provide incentives so that they exert R&D effort efficiently, as they delegate R&D investment decisions to innovators (i.e., to the “informed parties”), but they burden the market with the intellectual property monopoly welfare loss. Kremer (1998) suggests an ingenious mechanism, based on the idea of patent buyout, to design a prize system that provides efficient R&D investment incentives. Cornelli and Schankerman (1999) show that optimality can be achieved using either an up-front menu of patent lengths and fees or a renewal fee scheme. Boldrin and Levine (2008) provocatively challenge the views that patents are needed to remunerate R&D activity, when innovations are embodied in costly replicable capital or human capital.

⁵ For discussions about patent design in these settings, see Green and Scotchmer (1995), Scotchmer (1996), O’Donoghue (1998), O’Donoghue, Scotchmer, and Thisse (1998), and Denicoló (2000). For a mechanism design approach, see Hopenhayn, Llobet, and Mitchell (2006). Sequential innovations can also make the timing of innovation disclosure inefficient (see, e.g., Matutes, Regibeau, and Rockett 1996; Hopenhayn and Squintani 2016).

⁶ For example, these can also arise as a result of “horizontal” market value complementarities or substitutabilities among innovations (see, e.g., Cardon and Sasaki 1998; Lemley and Shapiro 2007).

These inefficiencies notwithstanding, these forces can also lead to biases in the direction of innovations because of differences in the degree of appropriability or financial needs, as argued by some recent papers in the literature. Most of these have centered on the scope of the innovations pursued, basic versus applied or extensive versus incremental. An early paper by Jovanovic and Rob (1990) considers the role of intensive and extensive search. Budish, Roin, and Williams (2015) investigate whether private research investments are distorted away from long-term projects. Akcigit, Hanley, and Serrano-Velarde (2021) consider the trade-off between basic and applied research in a general equilibrium model of technical change, while Akcigit and Kerr (2018) consider the trade-off between internal innovation by incumbents and external innovation of new entrants. Hopenhayn and Mitchell (2001) examine the case where innovations differ in terms of the prospects for follow-up innovations.

These considerations are obviously very important; however, our paper focuses on a different source of inefficiency that holds even when innovators receive the full social value of their innovations. This inefficiency arises from the fact that innovators pursue their research simultaneously, so the success of one crowds out the potential success of others. Our paper is thus closer to the literature on patent races (e.g., Loury 1979; Reinganum 1982).

A general conclusion from this literature is that there is excessive entry into innovation as a result of this negative spillover, driving to zero the rents of potential innovators (“rent dissipation”). Our research differs from this literature in two important ways. First, we focus on the allocation of a fixed set of innovators to alternative patent races, as opposed to a perfectly elastic supply of resources on a single race. Second, we examine sequences of patent races, as in the sequential innovation case. Our focus, of course, is on the allocation of these resources across different patent races. In line with the results of rent dissipation, we find that competition drives all rents to that of the marginal innovations, or what we call “differential rent dissipation,” as a result of overcrowding in certain areas of research and undercrowding in others.⁷

Our paper is also related to the literature on congestion. Our static misallocation force can be related to the study of the so-called price of anarchy in the congestion games developed by Rosenthal (1973). These

⁷ The idea of rent dissipation leading to excessive entry in models of product differentiation was considered in the seminal work by Spence (1976) and Dixit and Stiglitz (1977), in a paper by Mankiw and Whinston (1986), and more recently by Dhingra and Morrow (2019). The latter paper also examines the role of selection on productivity, which is somewhat related to the determination of the extensive margin of research areas that we also consider. Further distantly related to our work, there is also a literature studying the welfare effects of complementarities and substitutabilities among different research approaches to achieve the same innovation (e.g., Bhattacharya and Mookherjee 1986; Dasgupta and Maskin 1987; Letina 2016). Of course, this is different from the analysis of this paper, which considers several innovations without distinguishing different approaches to achieve any of them.

games model a traffic net, in which drivers can take different routes to reach a destination and routes get easily congested. In the optimal outcome, the drivers coordinate in taking different routes, whereas in equilibrium they excessively take routes that would be faster if they were not congested by suboptimal driving choices.⁸ Models of search with frictions have also focused on the role of congestion, mostly in connection to the single-market case. Directed search models (e.g., Shimer 1996; Moen 1997) consider the allocation of workers to heterogeneous firms with different productivities, which is closer to our setting. In contrast to our setting, the competitive equilibrium allocation is efficient. The key difference is that while firms have property rights for productive positions in models of directed search, there are no property rights for open problems in our setting.⁹

Bryan and Lemus (2017) provide a valuable general framework on the direction of innovation that encompasses the models cited here, as well as models of horizontal spillovers and sequential innovation. Building on the interaction across these different kinds of spillovers, they use their framework to assess when it is optimal to achieve incremental innovations versus large-step innovations and show that granting strong intellectual property rights to “pioneer patents” may lead to distortions in the direction of R&D. They also identify market distortions that are distinct from the market inefficiency identified here.¹⁰

Finally, there are some papers that build on our paper’s insights. Lee (2020) considers n innovators and two research lines to show that if the high-value innovation is more difficult, it may attract fewer researchers than in the first best. This is analogous to our results for the case of heterogeneous arrival rates in a model of Poisson arrival that we derive in the appendix (available online). Moraga-González, Motchenkova, and Nevrekar (2019) consider a market with a leader and n challengers. Each challenger attempts to become the market leader by achieving a quality innovation. Each allocates R&D effort between two projects that differ in terms of profits, difficulty, and social value. The winner of the challengers’ R&D race is determined according to a contest success function. They find that competitive equilibrium is inefficient for two reasons: (i) firms overinvest in the project with higher expected profitability, and (ii) firms underinvest

⁸ Similar results are obtained for the allocation of parking space in Anderson and De Palma (2004).

⁹ Mortensen (1982) had already proven the efficiency of allocations in the case of a single-patent race, when property rights over the innovation opportunity are assigned.

¹⁰ These distortions are demonstrated in a model with costless switching of researchers across R&D lines and without duplication of efforts in R&D races, so that it is optimal to concentrate all R&D resources on the most valuable R&D line, to then move on to the second-most valuable one after the first innovation is discovered, and so on and so forth. Under these assumptions, our paper’s market inefficiency that innovators overinvest in the most valuable R&D line may not arise.

in the more socially desirable project. A merger alleviates the former distortion but not the latter.

In addition, Chen, Pan, and Zhang (2018) consider an incumbent and n challengers. They study the effect of patentability standards on R&D efforts, entry decisions, and direction of innovation. Each challenger allocates effort between two possible R&D lines—one of known returns and one of uncertain returns. The winner of the challengers' race is determined by the first arrival of independent Poisson processes with arrival rates that are a function of each firm's effort. They find that R&D efforts and the number of entrants are too low in equilibrium relative to the first best. Interestingly, they find that firms are biased toward (against) innovation in the risky direction when the patentability standard is below (above) some threshold.

III. A Simple Example

There are two problems with private and social values $z_H > z_L$ and two researchers to be allocated to finding their solution. In any of the problems, the probability of success with one researcher is p and with two is $q > p$. We assume that $q - p < p$, capturing the idea that there is congestion or superfluous duplication of efforts. This assumption holds with slack in the case of independence, where $q = 2p - p^2 < 2p$.¹¹ We examine optimal and competitive allocations with one and two periods.

A. One-Period Case

Consider first the optimal allocation. Both researchers are allocated to H if and only if $qz_H \geq p(z_L + z_H)$ or, likewise,

$$(q - p)z_H \geq pz_L, \quad (1)$$

with a straightforward interpretation.

For the competitive case, we assume that if two researchers are allocated to H , then expected payoffs for each are $(1/2)qz_H$. This would happen, for instance, in a patent race where all value accrues to the first to solve the problem. The necessary and sufficient condition for both researchers to work on the H problem is that

$$\frac{1}{2}qz_H \geq pz_L. \quad (2)$$

¹¹ Our assumptions on congestion do not rule out correlation across innovation arrival rates or technological spillovers across R&D firms. We assume only that the probability of innovation with two competitors is less than twice the probability of success with one innovator. In other terms, a researcher is unhappy that a competitor starts researching on her same R&D project, turning her investigation into a race.

It is easy to verify that condition (43) implies (44), so the competitive allocation will always assign both researchers to H when it is optimal but might do so also when it is not.¹²

The difference between these two conditions can be related to the pecuniary externality (“market-stealing effect”) caused by entry into the H problem, which equals $(p - q/2)z_H$. Note that here the externality is not present when entering into the L problem, since there will be at most one researcher there. In the more general setting that we examine below with multiple research inputs, this externality will occur for more than one research line, and its relative strength is a key factor in determining the nature of the bias in the competitive allocation.

Another interpretation of this external effect is value burning. In the more general setup with many researchers that we examine below, the expected value of solving different research problems is equalized to the least attractive active one. All differential rents from solving more attractive problems are dissipated.

B. Two-Period Case

As above, in each period researchers can be allocated to the unsolved problems. In the case of both succeeding in the first period, there are then no more problems to solve. If one problem is solved in the first period, then in the optimal as well as in the competitive allocation both researchers are assigned to solve the remaining one.

To compare the equilibrium and optimal allocations, it is convenient to decompose total payoffs of the alternative strategies into first- and second-period payoffs. The second-period problem is a static one. If only one problem is left, then the two researchers will be assigned to it. If the two problems remain to be solved, we will assume for simplicity that condition (43) holds, so that both researchers are assigned to the H problem. Denote by w_{nt} the total expected payoffs in each period $t = 1, 2$ when assigning $n \in \{1, 2\}$ researchers to the H problem in period 1. We can write

$$\begin{aligned} w_{21} &= qz_H, w_{22} = q[(1 - q)z_H + qz_L], \\ w_{11} &= p(z_H + z_L), w_{12} = q[(1 - p)z_H + pz_L - p^2z_L]. \end{aligned}$$

The difference in first-period payoffs is identical to the calculation in the static case. Consider now second-period payoffs. The terms in brackets represent the expected value of the problems that remain to be solved. We call this the option value effect: as the planner has the option of solving

¹² We show in the appendix that this result generalizes if researchers’ ability is heterogeneous. This finding may be suggestive for future research on mergers and the direction of innovation, because mergers usually lead to improved economies of scale and efficiency.

problems in the second period, the planner recognizes that if they are not solved in the first period, then there is a residual value. This value is higher when the problem that remains to be solved is H . Ignoring the quadratic term (which becomes irrelevant in the continuous-time Poisson specification that follows), it is the case that $w_{12} > w_{22}$, so the incentives for initially allocating both researchers to H is weaker than in the static case.

Consider now the competitive allocation. Assuming that one player chooses H and letting v_{2t} represent expected payoffs in period t for the other player when also choosing H and v_{1t} when choosing L , it follows that

$$v_{21} = \frac{1}{2} qz_H, \quad v_{22} = \frac{1}{2} q[(1 - q)z_H + qz_L] = \frac{1}{2} w_{22},$$

$$v_{11} = pz_L, \quad w_{12} = q[(1 - p)z_H + pz_L - p^2 z_L] = \frac{1}{2} w_{12}.$$

Again, we are assuming here that in the second period, if both problems remain, then the two players will choose H . The difference $v_{21} - v_{11}$ is identical to the one for the one-period allocation. As shown above, and ignoring the quadratic term, the difference $w_{22} - w_{12}$ is negative, mitigating the gain from choosing H in the first period as in the optimal allocation. However, this difference here is divided by two. The reason is that the deviating agent does not internalize the value that leaving a better mix of problems to be solved for the second period has for the other researcher, while the planner does. In the more general setting that follows, as the number of players gets large, the dynamic effect vanishes from the competitive allocation condition, while it remains essentially unchanged in the planner's problem. The dynamic effect tilts the incentives in the competitive case toward the problem H , relative to the optimal allocation.

It is straightforward to find parameter values where (1) in the static allocation it is optimal to allocate both researchers to the H problem and (2) in the two-period case it is optimal to diversify, while specialization occurs in the competitive allocation. As an example, this will happen when $z_H = 3$, $z_L = 1$, $p = 3/8$, and $q = 1/2$.

The static allocation problem considered in section III.A can be reinterpreted as a multiperiod problem where researchers are fully specialized, so no reallocation takes place. Probabilities q and p should then be interpreted as those corresponding to final success.

IV. Assignment without Reallocation

In this section, we lay out the basic model used in the rest of the paper and consider a general form of the static allocation problem discussed in

section III.A. There is a continuum of problems or R&D lines, with one potential innovation each. Upon discovery, an innovation delivers value z that is distributed across research lines with cumulative distribution function F , which we assume to be twice differentiable. To isolate the findings of this paper from the well-known effects discussed earlier, we assume that the social value of an innovation coincides with the private value z .¹³ There is a mass $M > 0$ of researchers who are allocated to the different R&D lines according to a measurable function m . For each innovation of value z , we denote by $m(z)$ the mass of researchers competing for the discovery of that innovation. Hence, the following resource constraint needs to be satisfied:

$$\int_0^{\infty} m(z) dF(z) \leq M. \quad (3)$$

For each innovation z , the probability of discovery is $P(m(z))$. The function P is strictly increasing, concave, and such that $P(0) = 0$. The assumption that P is concave is the continuum analog of the congestion assumption $q < 2p$ of the example with two firms and two innovations of the previous section.¹⁴

The expected payoff of participating in an R&D line with value z and a total of $m(z)$ researchers is given by $U(z, m(z)) = P(m(z))z/m(z)$. As we show below, these payoffs can be interpreted as a winner-takes-all patent race where all participating researchers have equal probability of being first to innovate. Our model can thus be interpreted as an extension of the standard patent race to multiple lines. While that literature considers a single race with a perfectly elastic supply of researchers or firms with some entry or opportunity cost, here we consider the opposite extreme where a fixed supply of research inputs M must be allocated across multiple innovations.¹⁵

¹³ The expected present discounted value z_j of the patented innovation j does not necessarily equal the market profit for the patented product, net of development and marketing costs. It may also include the expected license fees paid by other firms that market improvements in the future or the profit for innovations covered by continuation patents. Thus, our model is compatible with standard sequential models of innovation, both those assuming that new innovations do not displace earlier ones from the market, as in the models following Green and Scotchmer (1995), and those assuming the opposite, as in the quality ladder models that follow Aghion and Howitt (1992).

¹⁴ We conjecture that our results generalize to an S-shaped function P with a convex region by taking the concave envelope. Our equilibria would select allocations in the concave part of the original function P , eliminating other sources of potential coordination failure that might lead to zero entry.

¹⁵ Obviously, with perfectly elastic supply the problem trivializes, as there would be no connection between entry decisions into different research areas. We consider the opposite extreme to emphasize the trade-off in allocating research inputs across different research lines, but our results should also hold for intermediate cases.

In a competitive equilibrium, expected payoffs $P(m(z))z/m(z)$ are equalized among all active lines, where $m(z) > 0$. We call this differential rent dissipation in analogy to absolute rent dissipation in the standard patent race literature. In contrast, in the optimal allocation \tilde{m} that maximizes $W(\tilde{m}) = \int_0^\infty zP(\tilde{m}(z)) dF(z)$, the marginal contributions $P'(\tilde{m}(z))z$ are equalized for all active lines.

In a competitive equilibrium, a marginal researcher contributes $P'(m)z$ to total value but gets a return $P(m)z/m$ that is greater as a result of the concavity of P . The difference $P(m)z/m - P'(m)z$ is the pecuniary externality inflicted on competing innovators. Relative to the value created, this externality is given by

$$\frac{P(m(z))}{P'(m(z))m(z)} - 1 = \frac{1}{\varepsilon_{Pm}} - 1, \quad (4)$$

where the first term corresponds to the inverse of the elasticity of discovery with respect to the number of researchers. It is immediate to see that the competitive allocation is optimal if and only if this external effect is the same across research lines. Given that differential rent dissipation implies an increasing function $m(z)$, this condition holds only when the elasticity of discovery is independent of m —that is, when the discovery function $P(m) = Am^\theta$ for some constant A .¹⁶

When this condition does not hold, the direction of the bias depends on how this external effect varies with m . Intuitively, when it increases (i.e., the elasticity of the discovery function is decreasing in m), there is excessive concentration in high- z areas, as we show below. We say that the competitive equilibrium is biased to higher- z (hot) research lines when the competitive and optimal allocations m and \tilde{m} satisfy the single crossing condition shown in figure 1. Formally, there exists a threshold \bar{z} such that $m(z) < \tilde{m}(z)$ for $z < \bar{z}$ and $m(z) > \tilde{m}(z)$ for $z > \bar{z}$. Further, when this condition holds, it is also the case that the smallest active R&D line innovation value is higher in equilibrium than in the first best—that is, that $\tilde{z}_0 = \inf_z \{\tilde{m}(z) > 0\} \leq z_0 = \inf_z \{m(z) > 0\}$. The following proposition gives conditions for this to hold.

PROPOSITION 1. In the absence of reallocation, the competitive equilibrium is biased to higher- z areas when the elasticity of discovery is decreasing in m .

Proof. See the appendix. QED

While the condition in this proposition might appear somewhat restrictive, it holds in the canonical model of innovation used in the patent race literature, as we show below. Moreover, as $P(m)/m$ is bounded by

¹⁶ Note, however, that since P is bounded by one, this function can hold only for a range where $m^\theta \leq 1/A$, and beyond this range the elasticity must be zero.




FIG. 1.—Bias to high- z areas. A color version of this figure is available online.

one, the elasticity must converge to zero as $m \rightarrow \infty$, so it must decrease in some region.

Stationary innovation process. The above setting can be embedded in a dynamic environment as follows. Let t denote the random time of discovery and $p(t; m)$ the corresponding density when m researchers are assigned from time zero to a research line of value z . The expected utility for each of them is given by

$$U(z; m) = \int_0^\infty \left(\frac{z}{m}\right) e^{-rt} p(t; m) dt.$$

Expected payoffs are divided by m since each innovator is equally likely to win the race and there are m researchers engaged in the race, and $p(t; m)$ denotes the density of discovery at time t . Letting $P(m) \equiv E[e^{-rt}; m] = \int_0^\infty e^{-rt} p(t; m) dt$, we can write $U(z; m) = zP(m)/m$, which is identical to the formulation given above. It is important to emphasize that while time is involved in the determination of payoffs, we are assuming here that once the discovery z is made, the m researchers involved become

idle and cannot be reallocated to other research lines. The following sections relax this assumption and explicitly consider the problem of reallocation.

We now specialize the setting to a stationary environment that is standard in the canonical models discussed in the introduction. Let $\lambda(m)m$ denote the hazard rate for discovery at any moment of time, so that $p(t; m) = \lambda(m)m e^{-\lambda(m)mt}$. Assume that $\lambda(m)m$ is increasing and concave and that $\lambda(0) = 0$. It easily follows that

$$P(m) = \frac{\lambda(m)m}{r + \lambda(m)m}$$

is concave and that the elasticity of P with respect to m is

$$\varepsilon_{Pm} = \left(\frac{r}{r + \lambda(m)m} \right) (1 - \varepsilon_{\lambda m}),$$

where $\varepsilon_{\lambda m}$ denotes the elasticity of λ with respect to m , $-\lambda'(m)m/\lambda(m)$. Proposition 2 then immediately follows:

PROPOSITION 2. If the elasticity $\varepsilon_{\lambda m}$ is weakly increasing in m , then the competitive allocation is biased to high- z lines.

We can interpret the elasticity $\varepsilon_{\lambda m}$ as the market-stealing externality per unit of value created: $\lambda(m)m/\lambda'(m)$. The condition given in proposition 2 then states that this externality increases with z . Note also that this is a sufficient but not necessary condition, as the first term is decreasing in m . Proposition 2 applies to the canonical R&D models, such as the ones discussed in the literature, where $\lambda(m) = \lambda m$ (i.e., discovery is independent across participants in a patent race) and the elasticity $\varepsilon_{\lambda m} = 1$. Each active research line can thus be interpreted as a patent race, where arrival rates are given by independent Poisson processes with rate λ and the first to innovate gets the rights to the full payoff z . More generally, the result applies for the constant elasticity case, where $\lambda(m) = \lambda m^{-\theta}$ for $0 \leq \theta < 1$. For the canonical model of patent races where $\theta = 0$, an explicit solution for the equilibrium and optimal allocations is given below.

PROPOSITION 3. Suppose that there is a continuum of R&D lines whose innovation discoveries are independent events, equally likely among each engaged researcher, with time-constant hazard rate λ . Then the equilibrium and optimal allocation functions are

$$m(z) = \frac{z - z_0}{\pi} \text{ for all } z \geq z_0 = r\pi/\lambda, \quad (5)$$

$$\tilde{m}(z) = \frac{r}{\lambda} \left(\sqrt{\frac{z}{z_0}} - 1 \right) \text{ for } z \geq \tilde{z}_0 = r\mu/\lambda, \quad (6)$$

where π represents the equilibrium profit of each R&D line and μ represents the Lagrange multiplier of the resource constraint. In equilibrium, innovators overinvest in the hot R&D lines relative to the optimal allocation of researchers: there exists a threshold \bar{z} such that $m(z) < \tilde{m}(z)$ for $z < \bar{z}$ and $m(z) > \tilde{m}(z)$ for $z > \bar{z}$.

Proof. See the appendix. QED

Importantly, this result demonstrates the market bias that is the theme of this paper (that competing firms overinvest in hot R&D lines) within a canonical dynamic model that may be related with the many R&D models since Loury (1979) and Reinganum (1981) that are built on the assumption of exponential arrival of innovation discoveries.¹⁷

To get a sense of the possible size of this distortion, we perform a simple back-of-the-envelope calculation. Suppose that innovation values are distributed according to a Pareto distribution of parameter $\eta > 1$, so that $F(z) = 1 - z^{-\eta}$ for $z \geq 1$. As proved in the appendix, when $\lambda(m) = \lambda$ and $\eta > 1$, in an interior allocation where $\tilde{z}_0 > 1$, the welfare gap is

$$\frac{W(\tilde{m})}{W(m)} = \frac{\eta}{\eta - 1} \left(\frac{\eta - 1}{2\eta - 1} \right)^{1/\eta}.$$

This ratio is plotted in figure 2. It is negligible for η close to one but quickly increases as η grows, so that $W(\tilde{m})/W(m) - 1$ reaches its maximum of about 20% for η close to 1.35 and then slowly decreases and disappears asymptotically as $\eta \rightarrow \infty$. Figure 3 gives the corresponding equilibrium and optimal allocations. The solid line corresponds to the optimal allocation and the dashed line to the equilibrium.

Heterogeneous arrival rates and flow costs. We have assumed here that arrival rates are the same for all research lines. Our results can be extended for heterogeneity where the attractiveness of R&D lines is determined not only by the innovations' expected market values but also by the ease of discovery. Letting λ_j be the discovery arrival rate of an innovation j with value z_j , we obtain $P_j(m_j) = m_j \lambda_j / [r + m_j \lambda_j]$. Ordering innovations j by the product $\lambda_j z_j$, it follows that the competitive equilibrium is biased to higher values, as shown in the appendix. In particular, this implies that if all R&D lines j have the same value z but differ in ease of innovation, there will be excessive entry into those with high λ_j . Further, the appendix shows how to generalize the analysis to allow for heterogeneous-flow research costs κ_j across innovations. This allows us to provide a general

¹⁷ Specifically, it is possible to formulate an "oligopolistic" version of our dynamic model with n R&D firms and each firm i hiring a mass $m_i(z)$ of researchers to allocate to an R&D line z . This results in an arrival rate of $\lambda m_i(z)$ for each innovation z to each firm i . The hiring choices of each firm i reproduce the same functional forms of the effort choices in Loury (1979) and subsequent papers built on the assumption of exponential arrival of innovations. Our model with a continuum of atomistic competitors can be understood as the limit case for $n \rightarrow \infty$ of this oligopolistic game.

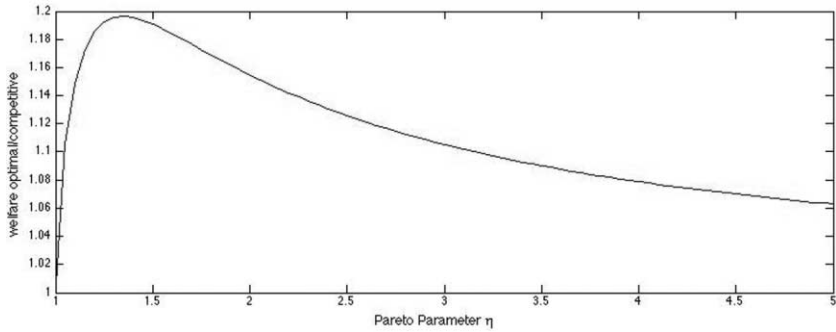


FIG. 2.—Plot of the welfare wedge $W(m)/W(\tilde{m})$.

definition of hot R&D areas, which includes easy problems or those that require a lower cost to solve. In that sense, our model is also consistent with a bias to smaller innovations or the low-hanging fruit. In sum, we say that the bias with which we are concerned is to high-return R&D lines—that is, lines j with high-flow expected return $\lambda_j z_j - \kappa_j$.

Quality ladders. Starting with the work of Grossman and Helpman (1991) and Aghion and Howitt (1992) and the subsequent model by Klette and Kortum (2004), quality ladders have been a workhorse model in the literature on sequential innovation. We show that our previous model can be easily inscribed in the context of a quality ladder. Suppose that goods or quality ladders are indexed by i in the unit interval and that for each good i there is an outstanding quality $q(i)$ interpreted as its social value flow. In addition, each quality ladder i has an opportunity of improvement $z(i)$ distributed according to $F(z)$ that, when attained, increases quality $q(i)$ by $rz(i)$, where r represents the common rate of discount. Given the cumulative nature of innovation, the present discounted social value of




FIG. 3.—Equilibrium and optimal allocation ($\eta = 1.35$). A color version of this figure is available online.

this improvement is $z(i)$. The technology for innovation is as described above, where each innovator assigned to a quality ladder makes the discovery with Poisson intensity λ and obtains a value $\alpha z(i)$, where $0 < \alpha \leq 1$. An allocation of innovators is an assignment $m(z)$ as a function of the quality of the innovation, where $\int m(z) dF(z) = M$, with an average arrival rate per quality ladder λM . If we further assume that step increases are i.i.d., setting $\alpha = r/(r + \lambda M)$ corresponds to the case where innovators appropriate the flow of value $rz(i)$, only until a new innovation occurs, as is commonly assumed in this literature.

V. Dynamic Allocation

We extend our previous analysis by allowing the mobility of researchers once a research line is completed. As before, we assume that there is a unit mass of research areas (the problems to be solved) with continuous distribution $F(z)$ and that there is an inelastic supply M of researchers. While here the set of problems is fixed, section VI considers a steady state with entry of new problems. Throughout this section, we assume that $\lambda(m)$ is twice continuously differentiable, $\lambda(m)m$ is strictly increasing and strictly concave, and $\lambda(0) = 0$. These assumptions imply that the arrival rate per researcher $\lambda(m)$ is decreasing—that is, that there is instantaneous congestion.¹⁸ Researchers are free to move across different problems, so the equilibrium and optimal allocations determine at any time t the number of researchers $m(t, z)$ assigned to each line of research z . This assignment, together with the results of discovery, implies an evolution for the distribution of open problems $G(t, z)$, where

$$\partial G(t, z)/\partial z = - \int^z \lambda(m(t, s))m(t, s)G(t, ds), \quad (7)$$

with $G(0, z) = F(z)$. An allocation is feasible if at all times the resource constraint

$$\int m(t, z)G(t, dz) \leq M \quad (8)$$

is satisfied.

¹⁸ Strict concavity implies that the arrival rate does not linearly scale with innovation, which can also capture duplication of innovation effort. In the process of achieving a patentable innovation, competing innovators often need to go through the same intermediate steps (see, e.g., the models of Fudenberg et al. 1983; Harris and Vickers 1985), and this occurs independently of every other innovators' intermediate results, which are jealously kept secret. Hence, the arrival rate of an innovation usually does not double if twice as many innovators compete in the same R&D race.

A. *Equilibrium*

Because the set of undiscovered innovations shrinks over time, it is never the case that innovators choose to move across R&D lines in equilibrium, nor that it is optimal to do so, unless the R&D line in which the researchers are engaged is exhausted as a consequence of innovation discovery. Indeed, the mass of researchers assigned to a particular line or research z will increase over time. Because mobility is free, the value of participating in any research line z at time t is equated to some value $w(t)$ whenever $m(t, z) > 0$. The value $v(t, z)$ of joining research line z at time t follows the Bellman equation:

$$rv(t, z) = \lambda(m)m\left(\frac{z}{m} + w(t) - v(t, z)\right) + v_t(t, z). \quad (9)$$

The first term represents the result of discovery that gives the researcher the value z with probability $1/m$ and the change in value $w(t) - v(t, z)$. The second term represents the change in value that occurs over time as the number of researchers allocated to every line increases. An equilibrium is given by an allocation $m(t, z)$ and distribution of open problems $G(t, z)$, together with values $v(t, z)$ and $w(t)$ such that:

1. The allocation m and distribution G satisfy equations (7) and (8);
2. The value function $v(t, z)$ satisfies the functional equation (9) and $v(t, z) \leq w(t)$ with strict equality when $m(t, z) > 0$.

Because the value of active research lines is equalized, $v(t, z) = w(t)$ and $v_t(t, z) = w'(t)$. As a result, equation (9) simplifies to

$$rv(t, z) = \lambda(m)m\left(\frac{z}{m}\right) + w'(t), \quad (10)$$

and since this value is equated across active research lines, it follows that $\lambda(m(t, z))z$ must be equal as well. This corresponds to the instantaneous value of participating in research line z , and because of free mobility, it must be the same across all active research lines. Differentiating this expression with respect to z , it follows that

$$m_z(t, z) = \frac{-\lambda(m(t, z))}{\lambda'(m(t, z))z}. \quad (11)$$

This equation can be integrated starting at a value $z_0(t)$, where $m(t, z_0(t)) = 0$ and $z_0(t)$ represents the unique threshold where the resource constraint

$$\int_{z_0(t)} m(t, z(t))G(t, dz) = M \quad (12)$$

is satisfied. As the mass of G decreases over time, it also follows that the threshold $z_0(t)$ decreases. We have proved the following result.

PROPOSITION 4. The equilibrium allocation is the unique solution $m(t, z)$ of equation (11) and is such that $m(t, z) > 0$ if and only if $z > z_0(t)$, where the threshold $z_0(t)$ is determined by equation (12).

B. Optimal Allocation

Consider an allocation $\tilde{m}(t, z)$. At time t , this gives a flow of value $\lambda(\tilde{m}(t, z))z$. Integrated over all active research lines and time periods, it gives the objective function

$$U = \max_m \int e^{-rt} \int \lambda(\tilde{m}(t, z)) \tilde{m}(t, z) z G(t, dz) dt. \tag{13}$$

The optimal allocation maximizes (13) subject to the resource constraint (8) and the law of motion (7). The latter is more conveniently expressed by the change in the density:

$$\partial g(t, z) / \partial t = -\lambda(m(t, z))m(t, z).$$

The formal expressions for the Hamiltonian are given in the appendix. Letting $u(t)$ denote the multiplier of the resource constraint and $v(t, z)$ the one corresponding to this law of motion, we can write the functional equation:

$$r\tilde{v}(t, z) = \max_m \lambda(\tilde{m}) \tilde{m} [z - \tilde{v}(t, z)] - u(t) \tilde{m} + \tilde{v}_t(t, z) \text{ for all } z \geq \tilde{z}_0(t). \tag{14}$$

Equation (14) represents the value of an unsolved problem of type z at time t . It emphasizes that problems are indeed an input to innovation, and as can be easily shown, the value of an open problem increases with z . Note the contrast to the private value $v(t, z)$, which is equal for all z as a result of the differential rent dissipation; in the eyes of competing innovators, all problems become equally attractive and valuable. The value function defined by (14) can also be interpreted as part of a decentralization scheme where property rights are assigned for each problem z and the owner of each open problem chooses the number of researchers to hire at a rental price $u(t)$. This interpretation highlights the source of market failure in our model precisely because of the lack of such property rights.¹⁹

¹⁹ This formulation can also be used to establish a connection with directed search. Consider the following market mechanism. A firm with value z offers a prize $p(z)$ to whomever first develops its research opportunity z . Joining this race gives a researcher the flow value $\lambda(m(z))p(z)$, which in equilibrium is equated at time t across all active areas to a value $u(t)$. This equivalence implicitly defines $m(p, t)$ by

$$\lambda(m(p, t))p = u(t).$$

The solution to the maximization problem in equation (14) gives

$$[\lambda(\tilde{m}(t, z))\tilde{m}(t, z) + \lambda(\tilde{m}(t, z))][z - \tilde{v}(t, z)] = u(t). \quad (15)$$

Comparing with the equilibrium condition where $\lambda(m(t, z))z$ is equalized reveals the two key sources of market failure that were illustrated in our simple example in section III. The first one is that the planner internalizes congestion (i.e., the market-stealing effect), which is why payoffs are multiplied not just by the arrival rate λ . It is useful to rewrite the term in brackets as

$$\lambda(\tilde{m}(t, z))[1 - \varepsilon_{\lambda m}]. \quad (16)$$

Note the parallel to the results in the case without reallocation considered in section IV, where the term in brackets corresponds to the wedge between the optimal and competitive allocation. The second term in brackets in equation (15) captures the fact that the payoff for discovery is smaller for the social planner, because it internalizes the fact that a valuable problem is lost as a consequence, as we found in our simple example. Taking the ratio of the two conditions gives

$$\frac{\lambda(m(t, z))}{\lambda(\tilde{m}(t, z))} = (1 - \varepsilon_{\lambda m}(\tilde{m}(t, z))) \left(\frac{z - v(t, z)}{z} \right) \frac{v(t)}{u(t)}.$$

For fixed t , the allocation functions cross when $\lambda(\tilde{m}(t, z)) = \lambda(m(t, z))$. Since the λ function is decreasing, a sufficient condition for $m(t, z)$ to remain higher after crossing is that this ratio decreases with z . This is the composition of two effects, represented by the two terms in brackets above. The first term is decreasing if the elasticity is increasing in z —that is, if the market-stealing effect increases. Because the value function is convex in z , the second term decreases in z . This corresponds to the option value effect that we found in our simple example. We have proved the following proposition.

PROPOSITION 5. Consider the model with free research mobility with individual arrival rate $\lambda(m)$. Suppose that the elasticity $\varepsilon_{\lambda m}$ is weakly increasing in m . Then, in equilibrium, innovators overinvest in the hot R&D lines: there exists a twice-differentiable threshold function \bar{z} such that $m(t, z) < \tilde{m}(t, z)$ for $z < \bar{z}(t)$ and $m(t, z) > \tilde{m}(t, z)$ for $z > \bar{z}(t)$.

The owner of research opportunity solves

$$rv(z, t) = \max_p \lambda(m(p, t))m(p, t)(z - p - v(z, t)) + v_i(z, t).$$

Substituting for $m(p, t)$ gives

$$rv(z, t) = \max_p \lambda(m(p, t))m(p, t)(z - v(z, t)) - m(p, t)u(t) + v_i(z, t),$$

which is equivalent to functional eq. (14). It immediately follows that the first-order conditions of this problem are identical to (15).

Note that even abstracting from the first effect (e.g., when the elasticity is constant) the competitive bias to hot areas still holds.

A borderline case occurs in the absence of instantaneous congestion—that is, when $\lambda(m) = \lambda$ and total arrival is linear in λ . Since there is no force to equalize rents in the competitive case, the solution is extreme and all researchers join the highest remaining payoff line at any point in time. The same turns out to be true in the optimal allocation, so in this knife-edge case the equilibrium is efficient. Reallocation costs provide an alternative rent-equalizing force that can lead to a nondegenerate equilibrium. These are examined in the following section.

C. Costly Reallocation

Assume that $\lambda(m) = \lambda$ in the stationary dynamic model presented above. Suppose that at any point in time each researcher can be moved across research lines by paying an entry cost $c > 0$. For every innovation of value z and time t , we denote the mass of engaged researchers as $m(t, z)$ and let $z_0(t)$ represent the smallest active R&D line innovation value at time t —that is, $z_0(t) = \inf_z \{m(t, z) > 0\}$. An equilibrium is defined in the same way as was done in the previous section.

Because the set of undiscovered innovations shrinks over time, there is positive entry into any active line of research, so it is never the case that innovators choose to move resources across R&D lines in equilibrium, nor that it is optimal to do so, unless the R&D line in which the researchers are engaged is exhausted as a consequence of innovation discovery.²⁰ Thus, we can approach the problem again using standard dynamic programming techniques. We express the equilibrium value $v(t, z)$ of a researcher engaged in an R&D line of innovation value z at t through the Bellman equation:

$$rv(t, z) = \lambda m(t, z) \left[\frac{z}{m(t, z)} + w(t) - v(t, z) \right] + \frac{d}{dt} v(t, z). \quad (17)$$

The flow equilibrium value $rv(t, z)$ includes two terms. The first one is the expected net benefit owing to the possibility of innovation discovery. The hazard rate of this event is $\lambda m(t, z)$; if it happens, each researcher gains z with probability $1/m(t, z)$ and experiences a change in value $w(t) - v(t, z)$, where $w(t)$ represents the value of being unmatched. The second term, $(d/dt)v(t, z)$, represents the time value change owing to the redeployment of researchers into the considered R&D line from exhausted research lines with discovered innovations.

²⁰ Further, as we show in proposition 6 below, there exists a time T after which researchers are not redeployed into other R&D lines, even when their research line is exhausted owing to innovation discovery.

For any time t , both the equilibrium value $v(t, z)$ and its derivative $(d/dt)v(t, z)$ are constant across all active R&D lines of innovation value $z \geq z_0(t)$.²¹ Let $v(t)$ and $v'(t)$ denote these values. In addition, note that since an unmatched researcher can join any research line at cost c , it follows that $v(t) = w(t) + c$. Substituting in (17), we obtain the no-arbitrage equilibrium condition:

$$\lambda[z - m(t, z)c] = rv(t) - v'(t) \text{ for all } z \geq z_0(t), \quad (18)$$

implying that $z - m(t, z)c$ is equated across all active research lines. Given that $m(t, z) \downarrow 0$ as $z \downarrow z_0$, it follows that payoffs of all active lines $z - m(t, z)c$ are equated to z_0 : differential rents are dissipated through higher-entry rates in higher-return areas, all being equated to the lowest active value line. Notice the parallel to the results in the patent race literature, where all rents are dissipated through entry. It follows that the flow value in the economy at time t is $\lambda Mz_0(t)$.

Solving for $m(t, z)$ using the above gives

$$m(t, z) = [z - z_0(t)]/c \text{ for all } z \geq z_0(t). \quad (19)$$

When the resource constraint $\int_{z_0(t)}^{\infty} m(t, z) dG(t, z) \leq M$ binds, the initial condition $z_0(t)$ is pinned down by the equation

$$cM = c \int_{z_0(t)}^{\infty} m(t, z) dG(t, z) = \int_{z_0(t)}^{\infty} (z - z_0(t)) dG(t, z), \quad (20)$$

where $G(t, z)$ is again the cumulative distribution function of innovations not yet discovered at time t .

We also note that because active R&D lines with innovation value $z \geq z_0(t)$ get exhausted over time, more researchers engage in the remaining lines (i.e., $m_i(t, z) > 0$ for all $z \geq z_0(t)$), less valuable lines become active (i.e., $z'_0(t) < 0$), and each active research line becomes less valuable (i.e., $v'(t) < 0$). Indeed, the value $v(t)$ decreases over time until the time T such that $v(T) = c$. At that time, redeployment of researchers stops at the end of the R&D race in which they are engaged. By then, active research lines have become so crowded that their value is not sufficient to recover the entry cost c any longer.²²

²¹ These conditions are akin to value matching and smooth pasting conditions in stopping problems (e.g., see Dixit and Pindyck 1994). Because R&D firms are competitive and labor is a continuous factor, the equilibrium dissipates all value differences from discovery of different innovations, through congestion and costly redeployment of researchers. This is similar to the phenomenon of rent dissipation in models of patent races with costly entry.

²² The characterization of the allocation $m(t, z)$ of researchers on undiscovered R&D lines at any time $t \geq T$ is covered by the earlier analysis of the canonical dynamic model without redeployment of researchers (cf. proposition 3). In our setup with a continuum of R&D lines distributed according to the twice-differentiable function G , arguments invoking laws of large number suggest that the allocation $m(t, z)$ would smoothly converge to the allocation $m(t)$ described in proposition 3.

The next proposition summarizes the above equilibrium analysis.

PROPOSITION 6. Assume that $\lambda(m) = \lambda$ and that researchers can be moved across R&D lines at cost $c > 0$. The equilibrium allocation is

$$m(t, z) = \frac{z - z_0(t)}{c} \text{ for all } z \geq z_0(t), \tag{21}$$

where the boundary $z_0(t)$ solves equation (20). Researchers are redeployed into different active R&D lines until the time T such that $z_0(T) = rc/\lambda$ and only if their research line is exhausted because of innovation discovery. The flow value in the economy at time t is $\lambda Mz_0(t)$.

We now consider the optimal allocation that is defined as in the previous section after subtracting total entry costs. Following the same approach, we solve for the optimal allocation \tilde{m} using the Bellman equation defined by the costate dynamic condition in the Hamiltonian. The details are provided in the appendix. The value of a research line z at time t satisfies

$$\begin{aligned} r\tilde{v}(t, z, 0) &= \max_{\tilde{m} \in \mathbb{R}} \lambda \tilde{m} [z - \tilde{v}(t, z, \tilde{m})] - r\tilde{m}c - u(t)\tilde{m} \\ &+ \tilde{v}_t(t, z, \tilde{m}) \text{ for all } z \geq \tilde{z}_0(t). \end{aligned} \tag{22}$$

There are several comments to make about this equation. First, because of the irreversible entry cost c , the value function includes as an additional argument the state variable \tilde{m} , representing the number of researchers allocated to this research line. On the left-hand side, we consider the flow value of an empty research line with $\tilde{m} = 0$. On the right-hand side, we consider the optimal choice of \tilde{m} . The entry cost $\tilde{m}c$ is expressed in flow terms consistently with the formulation of the value function. As before, $u(t)$ is the multiplier for the resource constraint. Finally, note that $\tilde{v}(t, z, \tilde{m}(t, z)) = \tilde{v}(t, z, 0) + \tilde{m}(t, z)c$ at the optimal choice. This also implies that $v_t(t, z, \tilde{m})$ is independent of \tilde{m} , which is used below. Substituting in (22), we obtain

$$\begin{aligned} r\tilde{v}(t, z, 0) &= \max_{\tilde{m} \in \mathbb{R}} \lambda \tilde{m} [z - \tilde{v}(t, z, 0) - \tilde{m}c] - r\tilde{m}c - u(t)\tilde{m} \\ &+ \tilde{v}_t(t, z, \tilde{m}) \text{ for all } z \geq \tilde{z}_0(t). \end{aligned} \tag{23}$$

This can also be interpreted as the Bellman equation for a firm that is assigned the property rights to a problem z at time t and needs to choose the initial amount of researchers to hire, paying the initial entry cost $\tilde{m}c$ and rental price $u(t)$. The solution of program (22) leads to the first-order conditions

$$\lambda [z - \tilde{v}(t, z, 0) - 2\tilde{m}(t, z)c] = rc + u(t) \text{ for every } z \geq \tilde{z}_0(t). \tag{24}$$

Equating these first-order conditions leads to the differential equation

$$\tilde{m}_z(t, z) = \frac{1 - \tilde{v}_z(t, z, 0)}{2c}. \quad (25)$$

By comparison, the differential equation for the equilibrium allocation obtained by differentiating (19) gives

$$m_z(t, z) = \frac{1}{c}. \quad (26)$$

It immediately follows that the derivative \tilde{m}_z of the optimal allocation function \tilde{m} is smaller than m_z , the derivative of the equilibrium allocation function m .²³ Because both functions m and \tilde{m} need to satisfy the same resource allocation constraint, this implies that the competitive equilibrium is biased toward high-return areas.

Further, the comparison of equations (25) and (26) allows us to single out two separate effects that lead to this result. The first one is the option value effect described earlier, where the marginal value of a better research line is $1 - v_z$, which is less than one, the marginal value in the competitive equilibrium. The derivative v_z captures the fact that a better problem also has more value in the future, in contrast to the equalization owing to rent dissipation that occurs in the competitive case. As a result, when engaging in an R&D line of value z , competing firms do not internalize the negative externality $\tilde{v}_z(t, z)$, the change in the continuation value owing to the reduced likelihood of discovering the innovation later. This leads the competing firms to suboptimally anticipate investment in the hot R&D lines, leading to overinvestment at every time t .

To see the second effect, note that the additional social marginal cost for engaging an additional researcher in a marginally more profitable line, $2\tilde{m}_z(t, z)c$, is twice the private additional expected cost $m_z(t, z)c$ incurred by the individual researcher. On top of this private cost, the society also suffers an additional redeployment cost. This cost is incurred in expectation by all researchers already engaged in the more profitable R&D line, in case the additional researcher wins the R&D race. This additional redeployment cost is not internalized by the competing firms, and it also pushes toward equilibrium overinvestment in the hot R&D lines.²⁴

PROPOSITION 7. Assume that $\lambda(m) = \lambda$ and that there is a cost of entry $c > 0$ to engage in any new problem. Then, in the competitive equilibrium, innovators overinvest in high-return R&D lines at every time t

²³ We prove in the appendix that $0 < \tilde{v}_z(t, z) < 1$.

²⁴ The result that R&D firms overinvest in hot R&D lines fails to hold only when $c = 0$ (the case of perfectly costless redeployment of researchers). In this case, assuming that the innovation value distribution has bounded support, all researchers will be first engaged in the most valuable R&D lines. When these innovations are discovered, the researchers will

there exists a threshold function $\bar{z}(t)$ such that $m(t, z) < \tilde{m}(t, z)$ for $z < \bar{z}(t)$ and $m(t, z) > \tilde{m}(t, z)$ for $z > \bar{z}(t)$.

This result completes our analysis of the stationary dynamic model in which the set of available problems to solve is fixed over time. The following section extends this model by considering the arrival of new R&D lines.

VI. Steady-State Economy

Consider the model analyzed in the last section, and suppose in addition that new problems arrive with Poisson intensity α and returns z distributed according to an exogenous distribution $F(z)$. We focus on the R&D line replacement that keeps the economy in steady state.²⁵

In steady state, the equilibrium allocation m is independent of time t and thus calculated with obvious modifications of the analysis presented earlier in this section. The expression $\lambda[z - m(z)c]$ is constant for all $z \geq z_0$, so that the equilibrium solves the differential equation $m'(z) = 1/c$, which gives the solution

$$m(z) = (z - z_0)/c \text{ for every } z \geq z_0. \quad (27)$$

Likewise, obvious modifications of the Bellman equation (22) show that the social planner problem takes the following form, in steady state:

$$r\tilde{v}(z) = \max_{\tilde{m} \in \mathbb{R}} \lambda \tilde{m} [z - \tilde{v}(z) - \tilde{m}c] - r\tilde{m}c - u\tilde{m} \text{ for all } z \geq z_0, \quad (28)$$

under the constraint that u satisfies the resource constraint. The associated first-order conditions are

$$\lambda [z - \tilde{v}(z) - 2\tilde{m}(z)c] = rc + u \text{ for every } z \geq z_0. \quad (29)$$

Inspection of the equilibrium condition and equation (28) reveals the same two forces identified earlier, leading to excessive equilibrium investment in hot areas, so there exists a threshold \bar{z} such that $m(z) < \tilde{m}(z)$ for $z < \bar{z}$ and $m(z) > \tilde{m}(z)$ for $z > \bar{z}$.

With simple manipulations presented in the appendix, we obtain

$$\lambda \left[z - c \frac{\lambda \tilde{m}(z)^2}{r} - 2\tilde{m}(z)c \right] = rc + u \text{ for every } z \geq z_0. \quad (30)$$

These equations are analogous to the ones obtained in the first-order conditions (24) for the model redeployment that we solved earlier. The

be redeployed to marginally less valuable research lines, until also these innovations are discovered, etc. This unique equilibrium outcome is socially optimal.

²⁵ For simplicity, we assume that when an innovation is discovered, the cost c for redeploying researchers is the same for all R&D lines, including the follow-up lines of the innovation discovery. Our results would extend to a more complicated model in which the redeployment cost is smaller for these lines as long as they are not exactly equal to zero.

only difference is that the term $\tilde{v}(z)$ takes the constant form $\lambda\tilde{m}(z)^2c/r$ here, which is the discounted cost of all future redeployment of the mass $\tilde{m}(z)$ of researchers engaged in the considered R&D line—the term $\lambda\tilde{m}(z)c/r$ is the individual discounted cost. Thus, we can identify as $\lambda\tilde{m}(z)^2c/r + \tilde{m}(z)c$ the “redeployment cost externality” that an additional researcher imposes on the $\tilde{m}(z)$ researchers engaged in the R&D line. As shown in the appendix, equation (30) can be used to obtain an explicit solution for the optimal allocation as a function of z_0 :²⁶

$$\tilde{m}(z) = \frac{r}{\lambda} \left(\sqrt{\lambda \frac{z - \tilde{z}_0}{rc} + 1} - 1 \right) \text{ for all } z \geq z_0. \quad (31)$$

To further the comparison between the equilibrium and first-best allocation functions m and \tilde{m} , we continue the analysis under the assumption that, in the steady-state economy, the value distribution of the new R&D lines is independent of the values of the R&D lines that they replace. Under this assumption, both allocation functions m and \tilde{m} satisfy the simple steady-state conditions:

$$\lambda m(z)g(z) = \alpha f(z) \text{ for all } z \geq z_0, \quad (32)$$

$$\lambda \tilde{m}(z)\tilde{g}(z) = \alpha f(z) \text{ for all } z \geq \tilde{z}_0, \quad (33)$$

where g and \tilde{g} denote the stationary equilibrium densities of undiscovered innovation values associated with m and \tilde{m} , respectively, f denotes the density of the innovation values of the new R&D lines, and $\alpha \leq \lambda M$ denotes the flow arrival rate of R&D lines. The equilibrium densities are defined for values of z above the respective thresholds. The R&D lines below the threshold become untouched and thus grow unboundedly.²⁷

Conditions (32) and (33) imply that, for any innovation value z with active R&D lines, the total mass of researchers allocated in the steady-state equilibrium and optimal allocations—respectively, $m(z)g(z)$ and $\tilde{m}(z)\tilde{g}(z)$ —are both equal to $(\alpha/\lambda)f(z)$, the net inflow of R&D lines of

²⁶ The net benefit of an additional researcher in the R&D line equals this researcher’s discovery hazard rate λ , multiplied by the innovation value z , minus the current and future discounted switching costs $\tilde{m}(z)c + \lambda\tilde{m}(z)^2c/r$ borne by the other $\tilde{m}(z)$ researchers, minus the cost $\tilde{m}(z)c$ of redeploying this marginal researcher. The latter, grouped with z , gives the expression $\lambda[z - \tilde{m}(z)]c$, which is the private marginal net benefit of researchers in the R&D line, as reported earlier.

²⁷ In the appendix, we consider the general case in which the values distribution of new R&D lines is not independent of the values of the discovered innovations. Unless the R&D line value distribution support also changes with the values of the discovered innovations, there exists an equilibrium that also satisfies eqq. (32) and (33). In the extreme case in which each discovery leads to an R&D line with the same innovation value, the option value effect identified comparing program (22) with eq. (17) disappears, but our main result that competing researchers overinvest in the hot R&D lines persists. In every other case, both the option value effect and our main result persist.

innovation value z . Of course, this does not mean that the mass of researchers engaged in each R&D line is also the same: it need not be that $m(z) = \tilde{m}(z)$ for any active R&D line of innovation value z because the stationary distributions g and \tilde{g} will differ.

We now turn to the determination of the thresholds. Assuming that the resource feasibility constraint $\int_{z_0}^{+\infty} m(z)g(z) dz \leq M$ binds in both allocations, the threshold z_0 is pinned down by plugging the stationarity condition (32) into the binding resource constraint, so as to obtain the equation

$$\alpha(1 - F(z_0)) = \lambda M. \tag{34}$$

Again, λM is the outflow of solved problems, which must equal the inflow of new relevant problems in steady state. Remarkably, this implies that the threshold z_0 is determined independently of the allocation function m , so in particular $\tilde{z}_0 = z_0$. Equations (27), (31), and (34) can be used to solve for the equilibrium and optimal allocations.²⁸

Because the thresholds z_0 and \tilde{z}_0 coincide, $m_z > \tilde{m}_z$, and $\lim_{z \downarrow z_0} m(z) = \lim_{z \downarrow z_0} \tilde{m}(z) = 0$, it follows that $m(z) > \tilde{m}(z)$ for all $z > z_0$. This is consistent with both allocations integrating to total resources M precisely because the stationary distribution of open problems \tilde{G} in the optimal allocation stochastically dominates G , the one in the stationary competitive equilibrium.

In words, the density of the R&D lines with undiscovered innovations is very large for small innovation values and very few researchers are engaged on these R&D lines; hence, innovation discoveries arrive with a very low rate. As the innovation value grows larger, the density of R&D lines with undiscovered innovations decreases. The rate of decrease is larger for the competitive equilibrium than for the optimal allocation function. Thus, the market suboptimally exhausts too many high-value R&D lines too early and leaves too few for future discovery. As a consequence of this, the number of researchers per project is always higher in equilibrium than in the social optimum, because there are more high-value R&D lines in the social optimum and these high-value R&D lines take up more researchers than low-value R&D lines.

Welfare.—In any allocation $m(z)$, the flow of value is

$$\begin{aligned} rV &= \lambda \int_{z_0} m(z)(z - m(z)c)g(z) dz \\ &= \alpha \int_{z_0} zf(z) dz - \alpha c \int_{z_0} m(z)f(z) dz, \end{aligned} \tag{35}$$

because of the steady condition $\lambda m(z)g(z) = \alpha f(z)$.

²⁸ If the resource constraint is satisfied with a strict inequality, $\int_{z_0}^{\infty} m(z) dG(z) < M$, then the economy cannot support entry by all firms; the participation constraint $\bar{v} \geq c$ binds and pins down z_0 through the equality $c = \bar{v} = (\lambda/r)z_0$.

The first term in equation (35) is the same in any allocation, and it is precisely the value of the outflow of problems solved that in a stationary equilibrium equals the corresponding inflow. Since the latter is independent of the allocation, so is the former. The second term corresponds to the total flow costs of redeployment, which differs across the two allocations. In the competitive allocation, $cm(z) = (z - z_0)$. Substituting in equation (35) gives

$$\begin{aligned} rV &= \alpha \int_{z_0}^{\infty} zf(z) dz - \alpha \int_{z_0}^{\infty} (z - z_0)f(z) dz \\ &= \alpha z_0(1 - F(z_0)) = \lambda M z_0. \end{aligned}$$

This represents a value equivalent to the flow of all innovations equalized to the lowest-value one, again reflecting differential rent dissipation. Note that as z_0 is independent of c , this value is the same for all c , within a range where all researchers are employed in the steady state. In particular, it holds surprisingly even as $c \downarrow 0$ because of an unboundedly increasing concentration in high-return areas.

Consider now the flow value of the optimal allocation. Using (31), it follows that

$$\alpha \tilde{m}(z)f(z)c = \frac{r\alpha}{\lambda} \left(\sqrt{c\lambda \frac{(z - \tilde{z}_0)}{r} + c^2} - c \right) f(z).$$

Substituting in (35) and using our previous result proves the following proposition.

PROPOSITION 8. When $\lambda(m) = \lambda$, the cost of entry is $c > 0$, and new problems arrive exogenously at rate α with quality distribution given by density f , aggregate equilibrium and optimal welfare are given, respectively, by

$$W(m) = (\lambda/r)z_0M \tag{36}$$

and

$$W(\tilde{m}) = \frac{\alpha}{r} \int_{z_0}^{\infty} zf(z) dz + cM - \frac{\alpha}{\lambda} \int_{z_0}^{\infty} \sqrt{\frac{c\lambda}{r}(z - z_0) + c^2} \cdot f(z) dz. \tag{37}$$

These closed-form expressions make welfare assessments simple and precise. Welfare is dissipated in the equilibrium allocation with excess researcher turnover to equal the flow of the lowest active research area. The welfare is not dissipated in the optimal solution, because the social planner spreads out researchers more evenly and leaves a larger number of hot R&D lines for later so that the society does not pay as much in terms of relocation costs.

When the switching costs are small, the optimal welfare expression simplifies further:

$$\begin{aligned} \lim_{c \rightarrow 0^+} W(\tilde{m}) &= (\alpha/r) \int_{z_0}^{\infty} zf(z)dz = (\alpha/r)E(z|z \geq z_0)[1 - F(z_0)] \\ &= (\lambda/r)E(z|z \geq z_0)M. \end{aligned}$$

Comparing this expression with the one derived above makes transparent the extent of rent dissipation in the competitive equilibrium allocation. For small switching cost c , the welfare ratio $W(m)/W(\tilde{m})$ takes the form

$$\lim_{c \rightarrow 0^+} \frac{W(m)}{W(\tilde{m})} = \frac{z_0}{E(z|z \geq z_0)}.$$

In words, the welfare ratio converges to the innovation value of the smallest active R&D line z_0 , divided by the average innovation value. This ratio can be very small for empirically plausible cumulative distributions F of innovation values.²⁹

VII. Final Remarks

Research on the efficiency of innovation markets is usually concerned with whether the level of innovator investment is socially optimal. This paper has asked a distinct, important question: Does R&D go in the right direction? In a simple dynamic model, we have demonstrated that R&D competition pushes firms to disproportionately engage in areas with higher expected rates of return. As far as we can tell, the identification of this form of market failure is a novel result. The competitive bias toward high-return areas comes from three distortions: (1) the cannibalization of returns of competing innovators, (2) excessive turnover and duplication costs, and (3) excessive entry into high-return areas because the market does not take into account the future value of an unsolved problem while a social planner does. In our steady-state analysis, the allocation of resources to problem-solving leads to a stationary distribution over open problems. The distribution of the socially optimal solution stochastically dominates that of the competitive equilibrium. A severe form of rent dissipation occurs in the latter, where the total value of R&D activity equals the value of allocating all resources to the least valuable problem solved.

²⁹ We performed a back-of-the-envelope calculation of the welfare ratio $W(m)/W(\tilde{m})$ under the assumption that the distribution F is lognormal with mean equal to seven and standard deviation equal to 1.5, consistently with the estimates provided by Schankerman (1998). With cost $c = 1$ million, the welfare ratio $W(m)/W(\tilde{m})$ is approximately 0.28. As the cost c vanishes, the ratio $W(m)/W(\tilde{m})$ converges to approximately 0.17.

The source of market failure in our model is the lack of property rights on problems. Standard forms of intellectual protection are not the immediate solution, as they grant property rights over the solutions and not the original problems. However, patent policy and other ways of rewarding innovation might still serve indirectly to offset the distortion we identified, by reducing private appropriation in the high-return areas.

The main sources of research funding are grants and fiscal incentives in the form of subsidies or tax breaks. Prizes, procurements, and the funding of academia also serve to subsidize R&D. These funding sources could serve to mitigate the bias to high-return areas when considerations other than return are taken into account. Indeed, prizes and direct subsidies have been used in the past to stimulate research into areas with lower returns, such as the development of orphan drugs to treat rare diseases. However, it is still possible that some prize, procurement, and career concerns in academia exacerbate the market inefficiency singled out in this paper. Plausibly, they may bias the incentives of individual researchers so that they disproportionately compete on a small set of high-profile breakthroughs, instead of spreading their efforts more evenly across valuable innovations.³⁰

The modeling framework that we present in this paper can be elaborated in several directions. In this paper, we have taken the arrival of innovation opportunities (problems) as an exogenous process, as we focused our attention on the market allocation of resources to solve the problems. But it is quite natural that new questions can arise in the process of solving older ones, so that the two processes are interrelated. One of the inefficiencies we find in our steady-state analysis is that good ideas are exhausted too fast in the market allocation, leading to a poor distribution of outstanding problems to solve in the steady state. This could be mitigated in part if, in the process of solving problems, new ones arise that are positively correlated with the quality of those being solved. Moreover, while we have assumed that the set of open problems is a public good, the discovery of some of these opportunities might be private and remain protected through secrecy by the firms or agents involved in this R&D process. We leave the investigation of these elaborations of our model to future research.

References

- Aghion, Philippe, and Peter Howitt. 1992. "A Model of Growth through Creative Destruction." *Econometrica* 60 (2): 323–51.
- Akcigit, Ufuk, Douglas Hanley, and Nicolas Serrano-Velarde. 2021. "Back to Basics: Basic Research Spillovers, Innovation Policy and Growth." *Rev. Econ. Studies* 88 (1): 1–43.

³⁰ Another policy of frequent application is fiscal incentives to R&D, often advocated on the basis that they leave the choice of the direction of R&D to the informed parties: the competing innovators. However, this is exactly the source of the market inefficiency that we have identified in this paper.

- Akcigit, Ufuk, and William R. Kerr. 2018. "Growth through Heterogeneous Innovations." *J.P.E.* 126 (4): 1374–443.
- Anderson, Simon P., and André De Palma. 2004. "The Economics of Pricing Parking." *J. Urban Econ.* 55 (1): 1–20.
- Arrow, Kenneth K. 1962. "Economic Welfare and the Allocation of Resources to Invention." In *Rate and Direction of Inventive Activity: Economic and Social Factors*, edited by Richard R. Nelson, 609–26. Cambridge, MA: NBER.
- Bhattacharya, Sudipto, and Dilip Mookherjee. 1986. "Portfolio Choice in Research and Development." *RAND J. Econ.* 17 (4): 594–605.
- Bloom, Nicholas, Mark Schankerman, and John Van Reenen. 2013. "Identifying Technology Spillovers and Product Market Rivalry." *Econometrica* 81:1347–93.
- Boldrin, Michele, and David Levine. 2008. *Against Intellectual Monopoly*. Cambridge: Cambridge Univ. Press.
- Bryan, Kevin A., and Jorge Lemus. 2017. "The Direction of Innovation." *J. Econ. Theory* 172:247–72.
- Budish, Eric, Benjamin N. Roin, and Heidi Williams. 2015. "Do Firms Underinvest in Long-Term Research? Evidence from Cancer Clinical Trials." *A.E.R.* 105 (7): 2044–85.
- Cardon, James H., and Dan Sasaki. 1998. "Preemptive Search and R&D Clustering." *RAND J. Econ.* 29 (2): 324–38.
- Chen, Yongmin, Shiyuan Pan, and Tianle Zhang. 2018. "Patentability, R&D Direction, and Cumulative Innovation." *Internat. Econ. Rev.* 59 (4): 1969–93.
- Cornelli, Francesca, and Mark Schankerman. 1999. "Patent Renewals and R&D Incentives." *RAND J. Econ.* 30 (2): 197–213.
- Dasgupta, Partha, and Eric Maskin. 1987. "The Simple Economics of Research Portfolios." *Econ. J.* 97:581–95.
- Denicoló, Vincenzo. 2000. "Two-Stage Patent Races and Patent Policy." *RAND J. Econ.* 31 (3): 488–501.
- Dhingra, Swati, and John Morrow. 2019. "Monopolistic Competition and Optimum Product Diversity under Firm Heterogeneity." *J.P.E.* 127 (1): 196–232.
- Dixit, Avinash K., and Robert S. Pindyck. 1994. *Investment under Uncertainty*. Princeton, NJ: Princeton Univ. Press.
- Dixit, Avinash K., and Joseph E. Stiglitz. 1977. "Monopolistic Competition and Optimum Product Diversity." *A.E.R.* 67 (3): 297–308.
- Fudenberg, Drew, Richard Gilbert, Joseph E. Stiglitz, and Jean Tirole. 1983. "Preemption, Leapfrogging and Competition in Patent Races." *European Econ. Rev.* 22:3–31.
- Green, Jeremy R., and Suzanne Scotchmer. 1995. "On the Division of Profit in Sequential Innovation." *RAND J. Econ.* 26:20–33.
- Grossman, Gene M., and Elhanan Helpman. 1991. "Quality Ladders in Theory of Growth." *Rev. Econ. Studies* 58 (1): 43–61.
- Hall, Bronwyn H., and Josh Lerner. 2010. "The Financing of R&D and Innovation." In *Handbook of the Economics of Innovation*, vol. 1, edited by Bronwin H. Hall and Nathan Rosenberg, 609–39. Amsterdam: North-Holland.
- Harris, Christopher, and John Vickers. 1985. "Perfect Equilibrium in a Model of a Race." *Rev. Econ. Studies* 52:193–209.
- Hopenhayn, Hugo, Gerard Llobet, and Matthew Mitchell. 2006. "Rewarding Sequential Innovators: Prizes, Patents, and Buyouts." *J.P.E.* 114:1041–68.
- Hopenhayn, Hugo, and Matthew Mitchell. 2001. "Innovation Variety and Patent Breadth." *RAND J. Econ.* 32 (1): 152–66.
- Hopenhayn, Hugo, and Francesco Squintani. 2016. "Patent Rights and Innovation Disclosure." *Rev. Econ. Studies* 83 (1): 199–230.

- Horstmann, Ignatius, Glenn MacDonald, and Alan Slivinski. 1985. "Patents as Information Transfer Mechanisms: To Patent or (Maybe) Not to Patent." *J.P.E.* 93 (5): 837–58.
- Jovanovic, Bojan, and Rafi Rob. 1990. "Long Waves and Short Waves: Growth through Intensive and Extensive Search." *Econometrica* 58 (6): 1391–409.
- Klette, Tor Jakob, and Samuel Kortum. 2004. "Innovating Firms and Aggregate Innovation." *J.P.E.* 112 (5): 986–1018.
- Kremer, Michael. 1998. "Patent Buyouts: A Mechanism for Encouraging Innovation." *Q.J.E.* 113:1137–67.
- Lee, Kangoo. 2020. "The Value and Direction of Innovation." *J. Econ.* 130:133–56.
- Lemley, Mark A., and Carl Shapiro. 2007. "Patent Holdup and Royalty Stacking." *Texas Law Rev.* 85:1990–2049.
- Letina, Igor. 2016. "The Road Not Taken: Competition and the R&D Portfolio." *RAND J. Econ.* 47 (2): 433–60.
- Loury, Glenn C. 1979. "Market Structure and Innovation." *Q.J.E.* 93 (3): 395–410.
- Mankiw, Gregory N., and Michael D. Whinston. 1986. "Free Entry and Social Inefficiency." *RAND J. Econ.* 17 (1): 48–58.
- Matutes, Carmen, Pierre Regibeau, and Katharine Rockett. 1996. "Optimal Patent Design and the Diffusion of Innovations." *RAND J. Econ.* 27:60–83.
- Moen, Espen R. 1997. "Competitive Search Equilibrium." *J.P.E.* 105:385–411.
- Moraga-González, José Luis, Evgenia Motchenkova, and Saish Nevrekar. 2019. "Mergers and Innovation Portfolios." Working Paper no. TI 2019-085/VII, Tinbergen Inst., Amsterdam.
- Mortensen, Dale. 1982. "Property Rights and Efficiency in Mating, Racing and Related Games." *A.E.R.* 72 (5): 968–79.
- Nelson, Richard R. 1959. "The Simple Economics of Basic Scientific Research." *J.P.E.* 49:297–306.
- O'Donoghue, Ted. 1998. "A Patentability Requirement for Sequential Innovation." *RAND J. Econ.* 29:654–79.
- O'Donoghue, Ted, Suzanne Scotchmer, and Jacques-Francois Thisse. 1998. "Patent Breadth, Patent Life, and the Pace of Technological Progress." *J. Econ. and Management Strategy* 7:1–32.
- Reinganum, Jennifer. 1982. "A Dynamic Game of R&D: Patent Protection and Competitive Behavior." *Econometrica* 50:671–88.
- Rosenthal, Robert W. 1973. "A Class of Games Possessing Pure-Strategy Nash Equilibria." *Internat. J. Game Theory* 2:65–67.
- Schankerman, Mark. 1998. "How Valuable Is Patent Protection? Estimates by Technology Field." *RAND J. Econ.* 29 (1): 77–107.
- Schumpeter, Joseph. (1911) 1934. *Theory of Economic Development*. Translated by Redvers Opie. New York: McGraw Hill.
- Scotchmer, Suzanne. 1991. "Standing on the Shoulder of Giants: Cumulative Research and the Patent Law." *J. Econ. Perspectives* 5:29–41.
- . 1996. "Protecting Early Innovators: Should Second-Generation Products Be Patentable?" *RAND J. Econ.* 27:322–31.
- Shimer, Robert. 1996. "Contracts in Frictional Labor Markets." Manuscript, Massachusetts Inst. Tech.
- Spence, Michael. 1976. "Product Differentiation and Welfare." *A.E.R.* 66 (2): 407–14.
- Wright, Brian D. 1983. "The Economics of Invention Incentives: Patents, Prizes, and Research Contracts." *A.E.R.* 73 (4): 691–707.