

Academic Freedom, Private-Sector Focus, and the Process of Innovation

Philippe Aghion
Harvard University and NBER

Mathias Dewatripont
Universite Libre de Bruxelles and CEPR

Jeremy C. Stein
Harvard University and NBER

July 13, 2005

Abstract

We develop a model that clarifies the respective advantages and disadvantages of academic and private-sector research. Our model assumes full protection of intellectual property rights at all stages of the development process, and hence does not rely on lack of appropriability or spillovers to generate a rationale for academic research. Instead, we focus on control-rights considerations, and argue that the fundamental tradeoff between academia and the private sector is one of creative control versus focus. By serving as a precommitment mechanism that allows scientists to freely pursue their own interests, academia can be indispensable for early-stage research. At the same time, the private sector's ability to direct scientists towards higher-payoff activities makes it more attractive for later-stage research.

This work is supported by the National Science Foundation. We are grateful to seminar participants at [FILL IN], and to Nicola Lacetera, Josh Lerner and Fiona Murray [OTHERS?] for their comments. Thanks especially to Bob Gibbons and Rebecca Henderson for extremely helpful conversations during the early stages of this research.

1 Introduction

Many important innovations, in industries ranging from pharmaceuticals to computer technology, have their origins in publicly funded research conducted at universities, foundations, and other non-profit institutions. The traditional case for government funding of such academic research is a familiar one: Because of knowledge spillovers and imperfect intellectual-property-rights (IPR) protection, the economic value associated with certain kinds of ideas cannot be fully appropriated by the developers of these ideas, leading to private-sector underinvestment in “basic” research (Nelson (1959), Arrow (1962)).¹

In recent years, there has been a substantial expansion of formal IPR protection for early-stage research. This trend is in part a result of the Bayh-Dole Act of 1980, which gives academic institutions the right to patent and commercialize discoveries made with government-sponsored research support, and which has helped to launch a boom in the creation of university technology-transfer offices. Lach and Schankerman (2004) report that the number of patents granted to university scientists increased from 500 in 1982 to more than 3,100 in 1998.²

According to the traditional view, one might expect this trend toward increased IPR protection to be a largely beneficial one. After all, if academia is thought of as a second-best solution to the underinvestment problem caused by insufficient appropriability, then increased appropriability and an accompanying shift of research activity to the private sector should be efficiency-enhancing. However, the trend has been highly controversial, particularly in the fields of biotechnology and pharmaceuticals. Many authors have expressed the concern that innovation in these fields is ultimately held back, rather than encouraged, when IPR protection is granted to the sorts of ideas that have traditionally been left in the public domain. Heller and Eisenberg (1998, p. 698) talk of an “anti-commons” effect associated with early IPR protection, arguing that: “A proliferation of intellectual property rights upstream may be stifling life-saving innovations further downstream in the course of research and product development.”

In this paper, we develop a model that: i) clarifies the respective advantages and disadvantages of academic and private-sector research; and ii) allows one to say when—in the process of developing an idea from its very earliest stages to a finished commercial product—it is normatively optimal to make the transition from academia to the private sector. Importantly, we assume *full IPR protection at all stages of the development process*. We do so not because we think that this is the most realistic scenario in all cases. Rather, it is the scenario where we think the tension between academic and private-sector funding of innovation is the most interesting, and least well-understood. Certainly, there are many sorts

¹It should be noted that while arguments based on spillovers make clear the need for basic research to be subsidized by the government, they are less clear on why this subsidy needs to happen in a different organizational form—i.e., in a university, as opposed to in a private corporation.

²See also Henderson, Jaffee, and Trajtenberg (1998).

of early-stage ideas where IPR protection is simply not feasible, and hence where, for the traditional reasons, academic research is the only option. Nevertheless, in fields like biotechnology, it is becoming increasingly evident that a lack of available IPR protection is not the only relevant factor in thinking about the merits of academia vs. the private sector.

Our model is based on authority and control-rights considerations, and emphasizes what we believe to be a fundamental tradeoff between academia and the private sector, namely the tradeoff of creative control versus focus. We take the defining characteristic of academic research to be that individual scientists retain the decision rights over what specific projects to take on, and what methods to use in tackling these projects (within certain bounds, of course). Indeed, in our model, academia effectively boils down to nothing more than a commitment mechanism that ensures scientists that these decision rights will not be abrogated.³ In contrast, the defining characteristic of private-sector research is that decision rights inevitably reside with the owner/manager of the firm, who can (and will) largely dictate project choice and methods to the individual scientists who work for the firm.⁴

By modeling academia as a commitment device, we deliberately abstract from a variety of other issues. These include the difficult question of what the objective function of academia is or should be, as well as many other job-design aspects that are relevant to the life and performance of universities (e.g., explicit and implicit incentive schemes, promotion systems, multitasking between research and teaching, etc.). Instead, we concentrate on analyzing the costs and benefits of delegating authority to scientists, and on showing how these costs and benefits vary over the life cycle of a research program.

More specifically, following Aghion and Tirole (1997), and in the general spirit of the property-rights literature (Grossman and Hart (1986), Hart and Moore (1990), Hart (1995)), we argue that scientists value creative control, and will have to be paid a wage premium in order to give it up. This assumption receives striking support in recent empirical work by Stern (2004), who studies the job market for recent PhDs in biology. By using multiple job offers, Stern is able to control for differences in ability across job candidates. After doing so, he finds that wages are substantially lower in jobs that promise scientists either some freedom to pursue their own individual research agendas, or that encourage the publication of this work.⁵

Thus one advantage of academia is that scientists can be hired more cheaply than in the private sector. The disadvantage of academia, however, is that they

³Lacetera (2005) also emphasizes the precommitment function of academia. As we argue below, the ability of academia to precommit not to interfere with scientists is intimately connected to its non-profit nature. In practice, this function is further reinforced by the tenure system; see, e.g., McPherson and Schapiro (1999).

⁴Howitt (2003) stresses that "...the kind of eccentricity that leads one down the path (of individual creativity) is rarely tolerated in private corporations, whose culture in this respect is a collective one that rewards employees for subordinating their private goals to that of the organization".

⁵Moreover, this result continues to hold when the sample is restricted to job offers from private-sector firms.

may end up working on projects that they find interesting, but that have little immediate economic value. In contrast, firms can, by virtue of their control rights, direct scientists to work on those projects that have the highest economic payoffs.

It turns out that the resolution of this tradeoff depends crucially on how far from commercialization a particular line of research is. To be concrete, imagine a line of biotech research which consists of ten distinct stages, and which will yield a drug worth \$10 billion if and only if all ten stages are successfully completed. At the final stage, so close to an enormous potential payoff, the wages of individual scientists are relatively insignificant, and the most important consideration is simply ensuring that every available scientist is working on the task at hand, as opposed to on some other pet project. Thus the directedness advantage of the private sector looms large, and it is optimal to have the project be privately owned at the last stage.

Now consider things from the perspective of the very first stage of the research line. It may be that even if this first stage is successful, there is only a one-in-a-thousand chance that all nine of the subsequent stages will be also. So loosely speaking, the value of succeeding in the first stage is only on the order of \$10 million. In this case, it becomes much more important to cede creative control, so as to economize on scientists' wages: if private-sector scientists cost \$200,000 each per year, and academic scientists cost only \$100,000 each per year, it may well be better to locate the project in academia, even if this entails some probability of the scientists wandering off in other directions.

Thus our first contribution is to provide a simple account of why it can be socially optimal to have earlier-stage, more "basic" research take place in academia, without relying on spillovers, imperfect IPR protection, or any of the other frictions that are usually invoked to rationalize a role for public funding of research. Somewhat more strongly, we also show that some lines of research can only get started in academia, and are simply not initially viable in the private sector.

Our second contribution is to ask whether, with full IPR protection, it is possible for ideas to get privatized too early. The answer is yes. Ideas will in general be viable in the private sector—i.e., will be able to attract non-zero bids from private firms—before it is socially optimal for them to migrate away from academia. In other words, if university technology transfer offices follow a policy aimed at maximizing their individual revenues from selling patents to private firms, ideas will tend to get into the private sector too soon.⁶ In this sense, the model validates the "anti-commons" concerns about there being such a thing as too much IPR protection for early-stage ideas: one can come up with examples where, in the context of the model, social welfare would be increased by eliminating early-stage IPR protection, and thereby forcibly delaying the

⁶ Anecdotal evidence suggests that some universities have been very aggressive in pursuing revenue maximization based on royalties from their patents. See, e.g., "Columbia's Pursuit of Patent Riches Angers Companies," by Bernard Wysocki Jr., (*Wall Street Journal*, December 11, 2004) for a discussion of Columbia University's management of its lucrative portfolio of biotechnology patents.

transition from academia to the private sector.

As noted earlier, much of the attention in the policy debate has been devoted to the biotechnology industry, where the ability of firms to patent individual genes has been a particular source of worry to some. Murray and Stern (2004) discuss an interesting case:

“Anecdotal evidence also suggests that whole gene patents (such as those for the breast cancer genes BRCA-1 and BRCA-2) have also stifled innovation and rendered more complex knowledge accumulation around breast cancer diagnostics and therapeutics...In the breast cancer case, a portfolio of over a dozen US and international patents gave biotechnology firm Myriad exclusive rights to commercialize laboratory testing services, diagnostic test kits and therapeutic products that use the BRCA1/2 DNA sequences. However, because the discovery of the genes was based not only on Myriad’s commercially funded research efforts but also built upon internationally generated public knowledge, private ownership of the breast cancer gene diagnostic kit was highly controversial.”

However, in spite of the visceral pull of examples like these, as a matter of basic theory it is not immediately clear why a firm like Myriad, once granted exclusive access to the breast cancer genes, would not have very strong incentives to pursue all economically viable applications of these genes. Moreover, even if Myriad does not have the in-house scientific expertise to go down all the different positive-NPV research paths, why can’t it either hire the right scientists, or license its patents to other firms who are better positioned to do so?⁷

Building on our control-rights foundation, we go on to develop two complementary explanations for why a firm like Myriad might not pursue all the economically viable applications of the breast cancer genes that it owns. The first story flows directly from our insight that private-sector firms’ cost structures make them excessively averse to relatively early-stage research. Suppose that there are two potentially legitimate research projects that make use of Myriad’s gene patents. One is an “applied” project that is only two stages away from a commercial payoff, while the other is a more “basic” project that is five stages away from any payoff. It is possible that the ultimate payoff on the latter, more basic project is sufficiently high that, *evaluated at academic-sector wages*, it is not only positive-NPV, but of greater NPV than the applied project. At the same time, it is also possible that, *evaluated at private-sector wages*, the basic project is negative-NPV, for the reasons described above. If this is the case, then when a private-sector firm like Myriad has the decision rights, it will allocate all of its scientists to the applied project, and completely ignore the basic project. Note that this conclusion does not rest on any exogenously-imposed

⁷Heller and Eisenberg (1998) argue that the central problem is the fragmentation of property rights across different private firms. If, for example, multiple patented inputs need to be combined to generate a new product, and each patent is owned by a different firm, various inter-firm transactions costs can be a barrier to efficiency.

constraints regarding the scope of private firms: it is not that Myriad cannot manage multiple projects at once, it simply finds it uneconomic to do so in this particular example, given its wage structure.

By contrast, if Myriad had never acquired the patents in the first place, and the ideas were left freely available to academic scientists, government funding of breast cancer research would naturally tend to lead to some progress on both projects, as individual scientists followed their own interests.⁸ Moreover, the government's investment in each of the two projects would be positive-NPV, because they would now each be evaluated at academic-sector wages. There would still be some loss of direction inherent in this outcome—without the ability to direct scientists, academia can never ensure the *optimal* allocation of scientists across the two projects—but this might be better than the private solution, which simply shuts down the more basic project.

It is easy to see where full IPR protection can go wrong in this example. Ex ante, even knowing that it will pursue only the applied project, the patents for the breast cancer genes may be quite valuable to Myriad, and it might be willing to bid aggressively to acquire them from a university technology-transfer office. The problem is that in setting its reservation price, the tech transfer office may not internalize the fact that privatization denies society the opportunity to finance more early-stage research in breast cancer at academic-sector wages.

Another reason why private-sector firms may pass up certain positive-NPV research projects that would be undertaken in academia has to do with the informational limitations that are inherent in directed research. Recall that the advantage of locating research in the private sector is that a principal (i.e., an owner, or a designated supervisor) can effectively force scientists to work on those projects with the highest commercial payoffs, rather than on worthless puzzle-solving. However, in reality it may be hard for any one principal to always discriminate between good and bad projects. For example, suppose that, starting with a given patent, there are two distinct research strategies that scientists can pursue, and that the relative payoffs to these strategies depend on the realization of a random variable that only scientists (and not the principal) can observe, with neither strategy dominating the other in all states of the world.

In this setting, there is value to letting scientists freely experiment in the early stages of a research program. Experimentation may reveal the underlying state of the world—i.e., which general approach to the problem at hand is the better one—and this is information which will be useful in future stages. Moreover, we show that under certain circumstances, it is impossible to sustain such experimentation in the private sector. Instead, private-sector scientists will simply be forced to follow whichever strategy looks more promising according

⁸See, e.g., “Cancer Research Has Many Healthy Side Effects,” by Raja Mishra (*Boston Globe*, February 8, 2005). This article discusses the many contributions to other medical fields—including AIDS, cardiovascular disease, kidney disease, multiple sclerosis, and psoriasis—that have come from research on cancer at the Dana-Farber Cancer Institute. The president of Dana-Farber is quoted in the article as saying that the non-profit institute encourages such spin-off research, which is largely funded by federal research grants.

to the principal's priors, in which case no information is produced that can be used in later stages. The advantage of academia in this version of the model is that it encourages experimentation in the early stages, which allows private-sector principals to bring better information to bear when they direct scientists in the later stages.

Although much of our analysis treats academia and the private sector as extreme organizational forms—the former one in which scientists face no incentives whatsoever, the latter one in which research strategies are always dictated from above—the model can also rationalize the existence of hybrid governance structures that lie between these two extremes. In particular, we examine private firms that behave in a less than fully authoritarian manner, leaving some control to their scientists.⁹ Such firms may be the optimal organizational form at an intermediate stage in a research program, at which point they effectively strike a compromise between the extreme creative freedom of academia and the absolute directedness of more fully authoritarian firms. We also consider the effect of introducing small monetary incentives (e.g., prizes) into an academic setting. While these variations enrich our basic model, and provide a more nuanced view of the roles of both academia and the private sector, they do not alter the main messages of our theory.

The remainder of the paper is organized as follows. In Sections 2 and 3, we introduce and then solve the most basic version of the model, in which at any given stage there is only one economically legitimate way to work with an existing idea carried forward from the previous stage. In this version, the only downside to the private sector is its more expensive wage structure. Thus, when we say that an idea gets taken private “too early”, the inefficiency is manifested in the fact that—because of the high wages—a private firm may employ too few scientists to work with the idea (relative to what would happen in academia). In Section 4, we add the possibility that, in addition to pushing forward an idea along an existing chain, a scientist may instead prefer to branch off and work on something that is more basic in nature, but still economically legitimate. Here, there is an additional inefficiency associated with going private too early, namely an excessive aversion on the part of the private sector to “stepping back” in this fashion. In Section 5, we examine hybrid governance structures. We consider in turn: i) the case where some real authority over the choice of research strategy may be left to scientists in a private firm; and ii) the possibility that incentive schemes are used to focus the interests of academic scientists. In Section 6, we assume that scientists may be better informed than private-sector principals. This yields one more cost of early privatization: it shuts down experimentation, and thus prevents the information of early-stage scientists from being made available to later-stage principals. Section 7 discusses some of the model's empirical content, as well as the connection to related literature, and Section 8 concludes.

⁹The role of Celera Genomics in the human genome project is a good example of such a hybrid structure (see Howitt (2003)).

2 Basic Framework

2.1 Technology

The development of an economically valuable product (e.g., a new drug) starts with an initial idea I_0 . This idea can be built on by subsequent scientists, in stages. If stage 1 is successful, there is a refined idea I_1 ; this refined idea can be further worked on to potentially generate an even-more-refined idea I_2 , etc. There are a total of k stages after the initial idea. If and only if all k stages are successful, there is a final idea I_k which generates a marketable product with value V .

The probability of success at any given stage depends on: i) the number of scientists who are active at that stage; and ii) the research strategies that they pursue. When a scientist is first exposed to an idea that has been brought forward from the previous stage, he must decide what strategy he wants to adopt in working with it. In the simplest version of the model, there are two options. First, the scientist can follow a “practical” strategy, which maximizes the probability that the current idea will be refined, and hence move on to the next stage. In particular, if there are n scientists at stage j who begin with the idea I_{j-1} and who all follow a practical strategy, there is a probability $\phi(n)$ that the idea will be refined and yield the new idea I_j .

We consider two different specifications of the function $\phi(n)$: i) $\phi(n) = p$ for all $n \geq 1$, and $\phi(0) = 0$; and ii) $\phi(n) = (1 - (1 - p)^n)$. The first specification corresponds to the assumption that all scientists working on the practical strategy have a perfectly correlated draw from the same success/failure distribution. This makes things especially simple—since it implies that in equilibrium there will always be exactly one scientist active at each research stage—and hence provides a useful way to illustrate the intuition for some of our results. At the same time, it can be *too simple* for some purposes, not allowing us to see the effects that arise when n is meaningfully endogenous. Hence the second specification, which corresponds to the assumption that scientists working on the practical strategy have independent draws from the same success/failure distribution, with each individual having a success probability of p , so that the probability of at least one success among a group of n is given by $(1 - (1 - p)^n)$.

Instead of the practical strategy, any given scientist may choose to follow the “alternative” strategy in working with an existing idea. In this case, the scientist has a zero individual probability of success, and hence contributes nothing to a group’s chances of a breakthrough. The simplest interpretation is that the alternative strategy is fundamentally worthless: i.e., it amounts to the scientist spending his time on puzzle-solving activities with no hope of an economic payoff. However, another possibility—which we explore below—is that even if the alternative strategy does not advance the *current* line of research, it may spawn an *entirely new* line of inquiry instead.

2.2 Scientists' preferences

There is an infinite pool of potential scientists, each of whom has an outside option R that he can obtain by working in another profession, e.g., as a taxi driver. This outside option sets a floor on the wages that scientists must earn. Beyond this, our key assumption is that scientists value creative independence—i.e., they value the right to choose how to pursue a particular idea. Specifically, after being exposed to idea I_{j-1} , each scientist at stage j decides whether he would better enjoy following the practical strategy or the alternative strategy. If he is able to undertake his favored strategy, he suffers no disutility from working. In other words, if a scientist could be promised *ex ante* that he would always be able to follow his favored strategy, he would be willing to work for a wage of exactly R . However, if the scientist has to undertake the strategy that he likes less, he suffers disutility of z . So if the scientist is certain that he is going to be forced to follow the less attractive strategy, he will set a reservation wage of $R + z$. In between these two extremes, scientists behave in a risk-neutral fashion, and require a wage premium that is proportional to the probability that they will have to undertake the less desirable strategy.

Importantly, neither the scientists themselves, nor their potential employers, know the scientists' preferences over the two strategies *ex ante*—i.e., before the scientists have had a chance to look at the previous-stage idea and think about it. That is, scientists' preferences for the practical vs. alternative strategies depend on the specifics of what kind of work these strategies will entail, and these specifics in turn depend on the nature of the previous-stage idea. To take a concrete example: a particular scientist may like to do a certain very specific kind of experimental work. *Ex ante*, it is unclear to him how this kind of experimental work will tie in with the project at hand. But after he has digested the previous-stage idea, it will become apparent whether his preferred experimental techniques are actually useful for pushing the idea to the next stage (in which case it will turn out that he is a practical type) or not (in which case it will turn out that he is an alternative type).

We assume that the *ex ante* probability that a scientist prefers to follow the practical strategy is given by α . Finally, we assume perfect correlation across all scientists at a given stage in terms of their preferences over the two strategies. In other words, either all scientists at a given stage prefer the practical strategy, or all prefer the alternative strategy. This strong perfect-correlation assumption is not necessary for our results—any non-zero positive correlation will do—but it greatly simplifies the exposition. Moreover, positive correlation in preferences can be thought of as reflecting the natural idea that some types of research are simply more fun for most scientists than others.

2.3 Academia

As noted in the introduction, we present an extremely rudimentary and stripped-down rendition of academia. We take the defining characteristic of this organizational form to be that it represents a precommitment to leave control over the

choice of research strategy in the hands of individual scientists. Although this assumption would appear to be empirically well-motivated, a natural question is *why* academia is uniquely able to make this commitment. Although we do not model it formally, we suspect that the non-profit nature of academia plays a central role. In particular, if one thinks of supervisory effort (the resources devoted to monitoring and directing scientists) as endogenous, it is plausible that academic administrators have much lower incentives to exert such effort than, e.g., a corporate CEO, whose compensation can be linked to the share price.¹⁰

Except in Section 5 below, we set aside the possibility of incentives schemes (either implicit or explicit) in academia. In Section 5, we effectively endogenize this assumption. We show that if z is non-stochastic, it is generally optimal not to use incentives in academia. Intuitively, if it is desirable to induce all scientists—irrespective of their preferences—to follow the practical strategy, this is more cheaply done in the private sector, where they can simply be directed to do so, as opposed to in academia, where they have to be promised unconditional bonuses for voluntarily choosing this option.

Finally, we assume that if the results from academic research are not sold to the private sector, these results are published and freely disseminated to other academic scientists. These assumptions are the only features that distinguish academia from the private sector in our model. In particular, we abstract from the question of what the exact mission of universities is or should be, or whether universities should be private or public.

The outcome of any stage j that takes place in academia is easy to describe. Suppose there are n scientists active at this stage. Each scientist is paid a wage $w_a = R$, and always works on his preferred strategy. This implies that with probability α , all n scientists work on the practical strategy, and with probability $(1 - \alpha)$, all n work on the alternative strategy. Thus the ex ante probability of advancing to the next stage is given by $\alpha\phi(n)$.

2.4 The private sector: firms and property rights

At some stage, a firm may acquire exclusive rights to an idea. Thus the following transaction is contractually feasible: an academic scientist (or the institution for which he works) may sell his idea to an entrepreneur, and promise not to publish the idea or in any other way share the idea with anybody else. The model therefore incorporates the potential for a strong form of IPR protection at all stages of the innovation process.

An important assumption is that only an entrepreneur has the funds to pay for an idea, as no scientists have funds of their own. This implies that if an entrepreneur is to take the idea forward, she will have to hire scientists as her agents at each subsequent stage of the development process. The entrepreneur can only derive utility from monetary sources, so she will only pay for an idea if

¹⁰See Hart, Shleifer and Vishny (1997) and Acemoglu, Kremer and Mian (2003) for related discussions about the role of non-profits.

she can earn a monetary return from it. Unlike the scientists, the entrepreneur has no innate preferences over research strategies. Nor does she get any utility from an idea being widely disseminated.

At the time the entrepreneur hires a team of scientists to work on a given stage, she cannot know the scientists' preferences over the practical vs. alternative strategies—these preferences only become evident once the scientists are inside the firm and have been given access to the idea by the entrepreneur. Yet *ex post*, the entrepreneur has the authority to force the scientists to work on whichever strategy she (the entrepreneur) deems to be most profitable. Indeed, it is impossible for the entrepreneur to precommit to doing otherwise—this is the defining characteristic of private-sector research. For example, once it becomes clear that the practical strategy requires a specific type of experimental work (call it type A), while the alternative strategy involves a different kind of experimental work (call it type B), the entrepreneur will force the scientists' hands by only buying the type of laboratory equipment that is compatible with type A work.

It follows that scientists will demand a wage of $w_p = R + (1 - \alpha)z$ in order to work in the private sector. The $(1 - \alpha)z$ markup over the academic wage represents compensation for loss of creative freedom—the fact that scientists now always have to adopt the practical strategy, whether this turns out to coincide with their preferences or not.¹¹

3 The Case of a Single Research Line

The initial point of departure for our analysis is the case of a single research line, as described above. First, we start with the perfectly-correlated draws assumption, which ensures that $n = 1$ at all stages. Next, we consider the independent-draws alternative, in which n is endogenous, and in which $\phi(n) = (1 - (1 - p)^n)$.

3.1 Perfectly correlated draws: $n = 1$

3.1.1 The basic trade-off between academia and private research

Consider a project which involves k stages, and imagine that the first $k - 1$ stages have been successful, so that we are now at stage k , with only one more success required to generate a payoff of V . If the last round of research is conducted in the private sector, and one scientist is hired, the expected payoff is:

$$E(\pi_k^p) = pV - w_p. \tag{1}$$

If instead the last round of research is conducted in academia, and one

¹¹If scientists were risk averse, the markup over the academic wage would be increased, because of the uncertainty that a private-sector scientist faces as to whether or not he will be able to pursue his favored strategy.

scientist is hired, the expected payoff is:

$$E(\pi_k^a) = \alpha pV - w_a. \quad (2)$$

Thus there is a simple tradeoff: on the one hand, wages are lower in academia. On the other hand, the inability to direct scientists in academia means that the probability of success is lower than in the private sector. Comparing the two equations, it is easy to see that the private sector will yield a higher payoff than academia if and only if:

$$(1 - \alpha)pV > (w_p - w_a), \quad (3)$$

or

$$pV > z. \quad (4)$$

An important first piece of intuition is that the private sector looks relatively more attractive when p and V are high, i.e., when the expected payoff to research is greater.

3.1.2 Why do we need academia?

Denote the maximum of $E(\pi_k^p)$ and $E(\pi_k^a)$ by Π_k . Folding back to stage $k - 1$, we can now compare:

$$E(\pi_{k-1}^p) = p\Pi_k - w_p. \quad (5)$$

and:

$$E(\pi_{k-1}^a) = \alpha p\Pi_k - w_a. \quad (6)$$

This implies that the private sector will yield a higher payoff than academia at stage $k - 1$ if and only if:

$$p\Pi_k > z. \quad (7)$$

Since $\Pi_k < V$, it follows that if the private sector is value-maximizing at stage $k - 1$, it is also value-maximizing at stage k . This recursive logic can be extended backwards, so that at any earlier stage i , we have:

$$E(\pi_i^p) = p\Pi_{i+1} - w_p. \quad (8)$$

and:

$$E(\pi_i^a) = \alpha p\Pi_{i+1} - w_a. \quad (9)$$

Moreover, the private sector will generate a higher payoff than academia at stage i and all future stages if and only if:

$$p\Pi_{i+1} > z. \quad (10)$$

Observe that as i declines—i.e., as we move backward to earlier and earlier stages—it becomes progressively harder for the private sector to outperform academia, since Π_{i+1} falls. This immediately implies:

Lemma 1: *It cannot be value maximizing to have academia operate at later stages than the private sector.*

Next, we can show that academia may become indispensable at the earlier stages of a line if the total length k of the line is sufficiently large. To see this, note that if the entire line is located in the private sector, its ex ante value, $\Omega(\text{allprivate})$, is given by:

$$\Omega(\text{allprivate}) = p^k V - (1 + \dots p^{k+1})w_p. \quad (11)$$

But this expression clearly becomes negative for k sufficiently large since $p^k V$ converges to zero whereas the expected wage bill $(1 + \dots p^{k+1})w_p$ remains bounded away from zero and increasing in k . We thus have:

Proposition 1: *A research program with a sufficiently large number of stages k will not be viable if located exclusively in the private sector.*

The proposition by itself does not fully establish the necessity of academia—we still need to show that for a non-empty set of parameter values, a research line that is not viable if located exclusively in the private sector can be viable if started in academia. But this latter point is easy to demonstrate. For example, suppose that $w_a = R = 0$, and consider the ex ante value $\Omega(\text{allacademic})$ of a line that is located exclusively in academia:

$$\Omega(\text{allacademic}) = (\alpha p)^k V, \quad (12)$$

which is obviously positive for all k .

We should emphasize at this point that the notions of “basic” and “applied” research implicit in our analysis are very different than what is usually seen in the literature on innovation. The tradition in this work is to think of the “basic-ness” of a particular line of research as being a function of how inappropriable are the ideas that it generates. In contrast, in our setting there are never any appropriability problems, and basic-ness—i.e., the necessity of locating research in an academic setting—is instead a function of the number of remaining steps until a commercial payoff can be realized.

3.1.3 The socially optimal transition point

Using our recursive approach, it is straightforward to calculate the point at which it is socially optimal for a research line to make the transition from academia to the private sector. The following lemma is an immediate consequence of (10):

Lemma 2: *From the perspective of a social planner, there is a unique transition point i^* , such that it is optimal for stage i^* to be the first stage conducted in the private sector. This transition point i^* is the smallest value of i that satisfies:*

$$p\Pi_{i+1} = p^{k-i+1}V - (p + \dots p^{k-i})w_p > z. \quad (13)$$

The comparative statics properties of the optimal transition point follow from this lemma. They are intuitive, and can be summarized as follows:

Proposition 2: *Holding fixed the number of stages k in a research line, it is optimal to have the transition to the private sector occur earlier if: i) V is greater; ii) α is smaller; or iii) z is smaller.*

Given an optimal transition policy, we still need to check that the research project is ex ante positive NPV—i.e., that it is socially worthwhile to fund the stages prior to i^* in academia. If the line is managed optimally throughout, its ex ante value, which we denote by $\Omega(i^*)$, is given by:

$$\Omega(i^*) = \alpha^{i^*-1} p^k V - (1 + \alpha p + \dots (\alpha p)^{i^*-2}) w_a - \alpha^{i^*-2} (p^{i^*-1} + \dots p^{k-1}) w_p, \quad (14)$$

where i^* is the optimal transition point determined in the previous lemma.

The ex ante feasibility (henceforth, EAF) constraint for the research line is then simply the condition that $\Omega(i^*) > 0$. This condition is always satisfied if $w_a = 0$, so that academic research is costless, and in much of what follows we use this assumption to keep things simple.¹² However, we will also briefly consider what happens when $w_a > 0$.

3.1.4 Comparison with early privatization

As noted in the Introduction, a number of authors have expressed the concern that, in a world with full IPR protection, privatization of a research line may occur sooner than is socially optimal. To provide a concrete way to think about this issue, imagine that the decision of whether to sell an academic idea to a private-sector firm rests in the hands of the university technology transfer office, so that the transition to the private sector occurs as soon as the value of the line under private management exceeds the reservation value T of the tech transfer office. If we denote by $i(T)$ the first stage conducted in the private sector under this scenario, we have that $i(T)$ is the smallest value of i that satisfies:

$$p^{k-i+1} V - (1 + p + \dots p^{k-i}) w_p > T. \quad (15)$$

Under this transition policy, the EAF constraint is modified, with i^* being replaced everywhere by $i(T)$:

$$\alpha^{i(T)-1} p^k V - (1 + \alpha p + \dots (\alpha p)^{i(T)-2}) w_a - \alpha^{i(T)-2} (p^{i(T)-1} + \dots p^{k-1}) w_p > 0. \quad (16)$$

This condition is obviously more restrictive than the EAF condition corresponding to socially optimal transition. That is, it is harder for the research program

¹² A less literal way of thinking about the assumption that $w_a = 0$ is that, for whatever reason, certain research lines are always able to get funding in academia, irrespective of NPV considerations. This could be because the government agency responsible for funding these lines has other objectives besides value maximization.

to be ex ante positive-NPV if it is managed suboptimally than if it is managed optimally.

In what follows, we consider the limiting case where the reservation price T is set at an arbitrarily low positive value. This implies that, absent any countervailing government policy, an idea transitions from academia to the private sector as soon as any private sector firm finds it economic to make a non-zero bid for it. We refer to this outcome as “early privatization”, and denote the associated transition point $i(0)$ by i^e .¹³ Clearly this is an extreme case, and not necessarily the most realistic one. However, it provides the starkest possible illustration of the concerns implicit in the “anti-commons” view: if there are ever going to be deviations from social optimality—in terms of ideas winding up in the private sector too soon—we will see these deviations most starkly by focusing on the early-privatization case.

As a logical matter, the early-privatization outcome can be rationalized by thinking of a revenue-maximizing tech transfer office that recognizes that if an idea that it controls in the current stage is instead allowed to stay in academia for one more stage, there is essentially a zero probability that the next successful breakthrough will take place in the same institution (because, say, there are a large number of competing universities working on pushing forward the same line of research).¹⁴

What are the welfare costs associated with early privatization? In the current version of the model, there are two possible effects. First, supposing that the EAF constraint is always satisfied, regardless of the timing of the transition to the private sector (this will be the case if $w_a = 0$), then early privatization does not prevent a research program from getting started in the first place. Consequently, its only downside is that it leads to inefficiently high labor costs. Indeed, in this case, early privatization necessarily *raises* the ex ante odds that the research program will ultimately bear fruit, but it does so at a labor cost that is too high relative to the benefit.

However, if we do not take for granted that the EAF constraint is always satisfied (because $w_a > 0$), then there can be a second cost of early privatization. In particular, a project that would initially get funded in academia under the socially optimal transition policy may no longer be worth funding if it anticipated that the transition will happen too soon. In this case, the expectation of early privatization has a more drastic effect, since it completely kills off an otherwise positive-NPV research line.¹⁵

¹³It is easy to show that a necessary condition for academia to viable one stage earlier than the private sector—i.e., for academic research to be positive-NPV at stage $i^e - 1$, conditional on a transition to the private sector at stage i^e —is that $\alpha z > R$. Intuitively, when αz is large relative to R , this tends to make academia viable at an earlier stage than the private sector because: i) academic scientists choose the practical strategy fairly often even without being directed to do so; and ii) there is a proportionally large wage premium in the private sector.

¹⁴Again, anecdotal evidence—such as the Columbia University story mentioned above—suggests that revenue maximization has become a very important consideration for some university technology transfer offices.

¹⁵We should be clear about the nature of the thought experiment we have in mind when we say that early privatization may lead to a violation of the EAF constraint and hence

3.1.5 Two examples

To conclude this subsection, we present two numerical examples to illustrate our above discussion:

Example 1: Set $R = w_a = 0$, $z = 2$, $\alpha = 0.5$ (implying that $w_p = 1$), $p = 0.45$, $V = 100$, and $k = 6$. Under the optimal transition policy, the first two stages are in academia, and the last four are in the private sector. This optimal policy yields an ex ante expected payoff of 0.12, and an ex ante probability of success of 0.0021. Under early privatization, the first stage is in academia, and the last five are in the private sector. Early privatization yields an ex ante expected payoff of 0.01, and an ex ante probability of success of 0.0042.

Example 2: Keep all the other parameters the same as in the previous example, but set $w_a = 0.05$, while keeping $w_p = 1$. (This corresponds to $R = 0.05$ and $z = 1.90$.) Under the optimal transition policy, the first two stages are again in academia, and the last four are in the private sector. This optimal policy yields an ex ante expected payoff of 0.06, and an ex ante probability of success of 0.0021. Under early privatization, the expectation that the project will move to the private sector after the first stage leads to a violation of the EAF constraint, so the research program never gets initiated.

3.1.6 Summary

Let us summarize the main messages from the model thus far. First of all, academia has an important role to play in the early stages of a research line: it is in general neither socially optimal, nor even feasible, to have early-stage research conducted in the private sector. Second, relative to a socially optimal transition policy, there are two types of costs associated with early privatization: i) inefficiently high labor costs; and ii) an increased likelihood of violating the EAF constraint, thereby preventing the research line from ever getting off the ground.

We should note one effect which is conspicuously absent from this version of the model. Conditional on the EAF constraint being satisfied, early privatization can never reduce the ex ante probability of success. That is, conditional on the project getting started in the first place, early privatization is necessarily a force in favor of innovation, with the only downside being that this comes at an inefficiently high labor cost. It turns out that this particular feature is an artifact of our simplifying assumption that the number of researchers at each stage is always equal to one. As we now show, when n is made endogenous in a more reasonable way, early privatization can stymie innovation even conditional

deter the initiation of a given research line. In this case, we are implicitly assuming that a social planner makes value-maximizing funding decisions in academia, but *takes as given* the inefficiency associated with early privatization. That is, the social planner can be thought of as a government agency that funds academic research, but that has nothing to say about either IPR policy, or the behavior of university tech transfer offices, and hence is unable to do anything about the timing of transition to the private sector.

on the project getting off the ground. This is because the higher labor costs associated with early privatization can now lead to a reduction in the number of scientists employed at a given stage in the private sector.

3.2 Independent draws: $\phi(n) = (1 - (1 - p)^n)$

3.2.1 Analysis

As before, to solve the social planner's problem, we work backward from stage k . If the last round of research is conducted in the private sector, and n scientists are hired, leading to a success probability of $\phi(n) = (1 - (1 - p)^n)$, the expected payoff is:

$$E(\pi_k^p) = (1 - (1 - p)^n)V - nw_p. \quad (17)$$

Ignoring integer problems, the firm's first order condition implies that the optimal number of scientists, n_k^p , is given by:

$$n_k^p = (\log(\beta V/w_p))/\beta, \quad (18)$$

where we have defined $\beta = -\log(1 - p) > 0$. It follows that at the optimum, expected stage- k private sector profit is:

$$E(\pi_k^{p*}) = V - (w_p/\beta)(1 + \log(\beta V/w_p)). \quad (19)$$

If instead the last round of research is conducted in academia, and n scientists are hired, the expected payoff is:

$$E(\pi_k^a) = (1 - (1 - p)^n)\alpha V - nw_a. \quad (20)$$

If we imagine that the number of academic scientists is also set at an optimal level—i.e., a well-intentioned government agency chooses the aggregate level of research funding across all universities—then we have:

$$n_k^a = (\log(\alpha\beta V/w_a))/\beta. \quad (21)$$

Note that the optimal number of academic scientists n_k^a can be either greater than or less than the optimal number of private sector scientists, n_k^p . This is because academic scientists are simultaneously cheaper, but less productive. At the optimum, expected stage- k profit in academia is:

$$E(\pi_k^{a*}) = \alpha V - (w_a/\beta)(1 + \log(\alpha\beta V/w_a)). \quad (22)$$

Denote the maximum of $E(\pi_k^{p*})$ and $E(\pi_k^{a*})$ by Π_k . Proceeding recursively, it follows that at any earlier stage i , we have:

$$E(\pi_i^{p*}) = \Pi_{i+1} - (w_p/\beta)(1 + \log(\beta\Pi_{i+1}/w_p)), \quad (23)$$

and

$$E(\pi_i^{a*}) = \alpha\Pi_{i+1} - (w_a/\beta)(1 + \log(\alpha\beta\Pi_{i+1}/w_a)). \quad (24)$$

In the appendix, we prove the following analog to Lemma 2:

Lemma 3: *Suppose that $E(\pi_k^{p*}) > E(\pi_k^{a*})$, so that it is optimal to locate the last stage (i.e., stage k) in the private sector. From the perspective of a social planner, there is a unique transition point i^* , such that it is optimal for stage i^* to be the first stage conducted in the private sector. This transition point i^* is the smallest value of i that satisfies $E(\pi_i^{p*}) > E(\pi_i^{a*})$, where these two quantities are defined by the recursive equations (23) and (24).*

By contrast to the social optimum, to solve for the transition point under early privatization, we simply keep folding backwards to earlier stages, always staying in the private sector. If k is large enough, we will eventually hit a stage i^e such that $E(\pi_{i^e}^{p*}) > 0$, but $E(\pi_{i^e-1}^{p*}) = 0$. The latter condition obtains when $(\beta\Pi_{i^e}/w_p) < 1$, so that at stage $i^e - 1$, a private sector firm is at a corner solution, with $n_k^p = 0$. It then follows that i^e is the earliest stage at which an idea is viable in the private sector—if an idea were to somehow accidentally wind up in the private sector earlier, no firm would ever invest positive resources in it.¹⁶

3.2.2 An additional benefit of academia

The following example—which is spelled out in more detail in Table 1—illustrates the additional positive scale effect associated with academic research that arises when n is endogenous. The example also shows how this scale effect alters the welfare comparison between the socially optimal transition policy and early privatization.

Example 3: *Set $R = w_a = 0$, $z = 1$, $\alpha = 0.5$ (implying that $w_p = 0.5$), $p = 0.10$, $V = 100$, and $k = 8$. Under the optimal transition policy, the first three stages are in academia, and the last five are in the private sector. This optimal policy yields an ex ante expected payoff of 2.08, and an ex ante probability of success of 0.078. Under early privatization, the first stage is in academia, and the last seven are in the private sector. Early privatization yields an ex ante expected payoff of 0.06, and an ex ante probability of success of 0.044.*

[INSERT TABLE 1 HERE]

The key feature of the example is that early privatization now not only lowers the ex ante expected payoff (by definition) it also lowers the ex ante probability that, conditional on the research line getting started, it will ultimately bear fruit. This is because now, with variable n , when the idea is privatized early,

¹⁶Even when an idea is not viable in the private sector at stage $i^e - 1$, it can nevertheless be viable in academia at this stage, provided that $(\alpha\beta\Pi_{i^e}/w_a) > 1$. Thus a necessary condition for academia to be viable at an earlier stage than the private sector is that $(\alpha/w_a > 1/w_p)$, which can be re-written as $\alpha z > R$. This is the same necessary condition that we stated above for the version of the model with $n = 1$ at all stages.

relatively few scientists are hired to work on it in the initial private-sector stages, as compared to the number that would be hired in academia. Consequently, the success probabilities for the initial private-sector stages are relatively low. In the context of the example, if the idea moves to the private sector early, at $k = 2$, only $n = 2.1$ scientists are hired, yielding a probability of success at this stage of 0.20. By contrast, if the idea stays in academia for the second stage, an infinite number of researchers are hired (since $w_a = 0$), yielding a probability of success at this stage of 0.50.

From a policy perspective, the example suggests that increased early-stage IPR protection need not always be beneficial to the innovation process. In particular, if it was legally impossible to protect the IPR associated with stage-2 and stage-3 ideas, then privatization could never happen before the socially-optimal transition point, and the ex ante probability of obtaining a successful product would necessarily be increased. This general feature—a non-monotonicity of innovation probabilities with respect to the degree of IPR protection—also recurs in several variations of the model that we present below.¹⁷

4 Alternative Strategies Create New Lines

Thus far, we have assumed that at each stage there is only one economically legitimate research strategy—namely the practical strategy, which has the potential to advance the project to the next stage along the chain. In contrast, the alternative strategy has been taken to be nothing more than worthless puzzle-solving. Now we modify this assumption. While we keep the restriction that only the practical strategy helps to advance the current line of research, we now allow the alternative strategy to yield new insights which may spawn wholly different lines of research. The interpretation is that when scientists turn away from the applied task of pushing the current line forward, they may not be shirking per se, but rather taking a useful step back that may ultimately generate fundamental breakthroughs.

To embed this notion into our model, we proceed as follows. We keep all the same assumptions as before, with one modification. Now, if at any stage of the original research line, a scientist works on the alternative strategy, there is a probability p_r of a revolutionary new idea which will form the basis for γ entirely new “offspring” research lines, with $\gamma \geq 1$. Each of these offspring lines has the same properties as the single lines analyzed above, although we allow for the possibility that the offspring have a greater number of stages than the original line, i.e., that $k_o \geq k$. Moreover, for computational simplicity but without any major loss of insight, we assume that the offspring lines are themselves sterile, and cannot give rise to further generations of revolutionary

¹⁷If there was no IPR protection at any stage, the entire line would have to be located in academia, leading to an ex ante probability of success of only 0.004 for the parameters in Example 3, much lower than in even the early-privatization case. So clearly, while some relaxation of early-stage IPR protection may be socially valuable, IPR protection in the later stages of a research line is absolutely essential.

ideas. That is, revolutionary ideas that yield offspring can only come from the alternative strategy applied at some stage of the original parent line.

In order to make things interesting, we assume that an offspring line has sufficiently many stages k_o that it is not viable if it is born into the private sector. This just means that the early-privatization point for an offspring line, i_o^e , exceeds one, or that:

$$p^{k_o}V - (1 + p + \dots p^{k_o-1})w_p < 0 \quad (25)$$

This assumption ensures that private-sector entrepreneurs will continue to direct scientists to stay away from the alternative strategy, and to focus all their efforts on the practical strategy, no matter how large γ is. Simply put, the assumption implies that the private sector never has any use for the offspring lines generated by the alternative strategy, because these lines are so early-stage that they are negative-NPV *when evaluated at private-sector wages*.¹⁸

In contrast, we assume that an offspring line is viable if it is born in academia, and managed optimally from that point on. That is, denoting the ex ante value of an offspring line under optimal management by Ω_o^* , we assume that $\Omega_o^* > 0$.

Using a logic similar to that above, we can derive the socially-optimal transition point for the parent line, i_p^* . We do this first for the perfectly correlated draws case where $n = 1$, and then for the independent draws case where $\phi(n) = (1 - (1 - p)^n)$.

4.1 Perfectly correlated draws: $n = 1$

Lemma 4: *From the perspective of a social planner, there is a unique transition point i_p^* , such that it is optimal for stage i_p^* to be the first stage of the parent line conducted in the private sector. This transition point i_p^* is the smallest value of i that satisfies:*

$$p\Pi_{i+1} > z + p_r\gamma\Omega_o^*. \quad (26)$$

The logic is identical to that of Lemma 1 in the basic model, and the expression for the optimal transition point is the same, except that a $p_r\gamma\Omega_o^*$ term has been added to the right-hand side of the inequality. The intuition is straightforward. In the basic model, it is optimal to make the transition to the private sector once the increase in value that comes from a higher probability of moving to the next stage is sufficient to outweigh the private-sector wage premium. Now, in addition to this wage premium, there is a second cost of moving to the private-sector—the fact that offspring lines are never developed. Or said differently, academia now has the added benefit of letting many more flowers bloom,

¹⁸For the purposes of this section, we are implicitly assuming that the private sector is unable to sell the rights to offspring lines back to academia. This could be, e.g., because no single university, acting as a revenue maximizer, would be willing to pay much for an idea: each university knows that it is going to take multiple stages for the idea to be sufficiently advanced as to be ready for resale to the private sector, at which point the final academic breakthrough is likely to occur in another university.

which makes it optimal to wait longer before moving to the private sector, all else equal.

The early-privatization transition point, i_p^e , is identical to that in the basic model. This is because once privatized, the remaining payoffs from the parent line are unchanged from before, as scientists are still always assigned to the practical strategy. Comparing the socially-optimal transition point and the early-privatization transition point, we have:

Proposition 3: *For the case of $n = 1$, the gap between the socially-optimal transition point for the parent line and the early-privatization transition point, given by $(i_p^* - i_p^e)$, is greater than in the basic model, and is increasing in the productivity $p_r\gamma$ of the of the alternative strategy.*

Example 4: *Set $R = w_a = 0$, $z = 2$, $\alpha = 0.5$ (implying that $w_p = 1$), $p = 0.45$, $V = 100$ for both the parent and offspring lines. Further assume that the parent line has $k = 5$ stages, that all offspring have $k_o = 6$ stages, and that $p_r = 1$ and $\gamma = 20$ for the parent line. Under the optimal transition policy, the first two stages of the parent line are in academia, and the last three are in the private sector. (By contrast, with $\gamma = 0$, only the first stage of the line is in academia under the optimal transition policy.) The optimal policy yields an ex ante expected payoff of 1.85, and an expected number of successfully completed products equal to 0.1189. Under early privatization, all five stages of the parent line are in the private sector, which implies that there are never any offspring lines. Early privatization yields an ex ante expected payoff of 0.06, and an expected number of successfully completed products equal to 0.0185.*

The example demonstrates two points. First, having $\gamma > 0$ increases the divergence between the optimal transition point and the early-privatization transition point; this is just a concrete illustration of Proposition 3. Second it is now possible to have a substantially lower rate of innovation—defined in terms of the ex ante expected number of research lines that reach successful completion—under early privatization than under the optimal policy, even when we restrict ourselves to the case where $n = 1$. Once again, this latter result suggests that it may be possible to increase the overall level of innovation in the economy by relaxing early-stage IPR protection. In the context of the example, innovation would be boosted if it were impossible to protect stage-1 ideas, so that stage-1 research could only be done in academia. This version of the model thus offers a particular rationalization of the “anti-commons” effect pointed out by Heller and Eisenberg (1998).

From an empirical perspective, what is perhaps most interesting about the current version of the model is that it implies that once an idea becomes the property of a private firm, it will be developed along relatively narrow lines. That is, the private sector’s ownership of a given idea *will not yield as diverse an array of useful next-generation ideas as would be generated in academia*. We discuss this prediction in more detail in Section 7.

4.2 Independent draws: $\phi(n) = (1 - (1 - p)^n)$

The effect of adding offspring lines into the independent-draws version of the model is qualitatively similar to that in the correlated-draws case. In the appendix, we prove a direct analog to Proposition 3:

Proposition 4: *For the case of $\phi(n) = (1 - (1 - p)^n)$, the gap between the socially-optimal transition point for the parent line and the early-privatization transition point, given by $(i_p^* - i_p^e)$, is greater than in the basic model, and is increasing in the productivity $p_r\gamma$ of the of the alternative strategy.*

Example 5: *Set $R = w_a = 0$, $z = 1$, $\alpha = 0.5$ (implying that $w_p = 0.5$), $p = 0.10$, $V = 100$, and $k = 8$. (Note that these are all the same parameter values as in Example 3.) Further assume that all offspring also have $k_o = 8$ stages, and that $p_r = 0.10$ and $\gamma = 10$ for the parent line. Under the optimal transition policy, the first five stages of the parent line are in academia, and the last three are in the private sector. This optimal policy yields an ex ante expected payoff of 21.58, and an expected number of successfully completed products equal to 2.02. Under early privatization, the first stage of the parent line is in academia, and the last seven are in the private sector. Early privatization yields an ex ante expected payoff of 0.36, and an expected number of successfully completed products equal to 0.26.*

[INSERT TABLE 2 HERE]

The contrast between Example 3 and Example 5—the details of which appear in Table 2—provides a clean illustration of the impact that offspring lines have on our normative comparisons. In Example 3, early privatization was seen to reduce the rate of innovation by 44 percent, as compared to the socially optimal transition policy (0.044 vs. 0.078). In Example 5, with everything else the same but for the addition of the offspring lines, early privatization reduces the rate of innovation by 87 percent (0.26 vs. 2.02).

5 Hybrid Organizational Forms

Thus far, our renditions of academia and the private sector have been extreme caricatures. At one end of the spectrum, we have cast the private sector as an organizational form in which scientists have no freedom of choice whatsoever, and are always forced to follow a research strategy dictated by an entrepreneur/supervisor, whether they like this strategy or not. At the other end of the spectrum, we have cast academia as a setting in which researchers not only have absolute creative control, but also face no incentives—either explicit or implicit—that might encourage them to pursue a relatively more practical strategy.

We now discuss how each of these extreme assumptions might be relaxed. Doing so leads to a more nuanced and realistic view of both organizational forms, though it does not alter the main message of our model. In what follows, we

focus on the basic version of the model from Section 3, and on the expositionally simpler case where $n = 1$ at all stages.

5.1 Less authoritarian private-sector firms

Many private-sector firms are known for taking explicit measures to give some of their research-oriented employees a degree of creative independence. For example, both 3M and Google apparently allow certain employees one day a week to pursue their own research interests.¹⁹ In an effort to model this behavior, we follow Aghion and Tirole (1997) and draw a distinction between formal and real authority in private-sector firms. The idea is that while the entrepreneur in a private firm always retains the formal right to direct her employees—by, e.g., choosing the type of lab equipment they work with—she may in fact choose not to exercise this right if she is not sufficiently informed to know which is the better strategy.

Suppose that a project is located in the private sector at stage i . The timing of events is now as follows. First, the entrepreneur hires a scientist, and agrees to pay him a wage of w_{pi} . Next, the entrepreneur invests effort in trying to become informed about the project. For an effort cost of $\theta\lambda^2/2$, the entrepreneur has a probability λ of becoming informed. If she is informed, she is then able to force the scientist to follow the practical strategy, as we have been assuming above. However, if she is uninformed, the entrepreneur is unable to direct the scientist, and the scientist is thus free to do what he wants, just as in academia.

Observe that the basic version of the model in Section 3 is just a special case of this one, in which $\theta = 0$, so that the entrepreneur always chooses to become informed with probability $\lambda = 1$. It is also worth noting that this modeling framework could be used to explicitly link the precommitment function of academia to its non-profit nature: if academic administrators (e.g., deans, or a university president) do not get profit-linked compensation, they will be unwilling to expend any effort on becoming informed, leading to a situation in which $\lambda = 0$ —i.e., in which real authority always rests with individual scientists.

Given our assumptions, the payoff to the entrepreneur if she is informed at stage i is:

$$E(\pi_i^p \mid \text{informed}) = p\Pi_{i+1} - w_{pi}. \quad (27)$$

The payoff to the entrepreneur if she is uninformed at stage i is:

$$E(\pi_i^p \mid \text{uninformed}) = \alpha p\Pi_{i+1} - w_{pi}. \quad (28)$$

Therefore, the marginal value of being informed at stage i is $(1 - \alpha)p\Pi_{i+1}$, and the entrepreneur's equilibrium probability of becoming informed at this stage is:

$$\lambda_i = (1 - \alpha)p\Pi_{i+1}/\theta. \quad (29)$$

¹⁹We thank Rebecca Henderson for suggesting these examples.

It follows that the unconditional expected payoff at stage i in the private sector is given by:

$$E(\pi_i^p) = (\lambda_i + \alpha(1 - \lambda_i))p\Pi_{i+1} - w_{pi} - \theta\lambda_i^2/2, \quad (30)$$

where the wage w_{pi} is itself determined as:

$$w_{pi} = R + \lambda_i(1 - \alpha)z. \quad (31)$$

With equation (30) taking the place of equation (8), the rest of the analysis from Section 3.1 continues to apply as stated. In particular, both the socially optimal transition point and the early privatization point are determined using the same approach as before. A couple of new empirical implications are worth noting, however. First, from (29), as the project moves closer to completion, the likelihood that the entrepreneur becomes informed and imposes her will on the scientist increases, since Π_{i+1} goes up. It then follows from (31) that the scientist's wage also increases, to compensate for the fact that he has less de facto creative control. In other words, private-sector firms endogenously become more authoritarian—and less like academia—as research projects move into their later stages.

Moreover, the model makes it clear that, even though we may observe some private-sector firms behaving in a less-than-fully authoritarian manner (e.g., the 3Ms and Googles of the world), it does not follow that there is no role for academia. Since the private-sector wage is set before the entrepreneur exerts effort to become informed, there is still a commitment problem in the private sector: while the probability of authority being exercised may be less than one, it can still be inefficiently high in the early stages of a research program. Thus the precommitment associated with academia remains valuable.

Example 6: Set $R = w_a = 0$, $z = 2$, $\alpha = 0.5$, $p = 0.45$, $V = 100$, and $k = 6$. In addition, assume that for a cost of $c = 4$, the entrepreneur has a probability $\lambda = 1$ of becoming informed, while for zero cost the entrepreneur has a probability $\lambda = 0.5$ of becoming informed. Under the optimal transition policy, the first three stages are in academia, the third stage is in the private sector with “moderate” authoritarianism ($\lambda = 0.5$) and a wage of $w_p = 0.5$, and the final two stages are in the private sector with full authoritarianism ($\lambda = 1$) and a wage of $w_p = 1$. This optimal policy yields an ex ante expected payoff of 0.04. Under early privatization, the first two stages are in academia, the next two stages are in the private sector with moderate authoritarianism, and the final two stages are in the private sector with full authoritarianism. Early privatization yields an ex ante expected payoff of 0.03.

5.2 Low-powered incentives in academia

In modeling academia, we have assumed that scientists are completely free to follow their preferences, and face no incentives—either explicit or implicit—that might push them in the direction of the practical strategy. Although the tenure

system can certainly be thought of as blunting the implicit incentives associated with career concerns (Holmstrom (1999)), it is nevertheless hard to argue that tenured academics face no incentives whatsoever. For example, scientists can earn both professional prestige and monetary prizes if their work is highly cited, and citations in turn are likely to have some relationship (albeit a noisy one) to the underlying usefulness of the research.

To introduce a meaningful role for incentives in academia, we generalize the model slightly, so that the disutility that a scientist experiences from following his less-favored strategy is now a random variable that can take on one of two values: z_L with probability ω ; and $z_H > z_L$ with probability $(1 - \omega)$. The outcome of this random variable is independent of everything else in the model, and it has a mean of $\omega z_L + (1 - \omega)z_H = z$.

Now suppose we want to design an incentive scheme that induces a scientist to follow the practical strategy when his disutility from doing so is z_L , but not when it is z_H . This scheme will have to have two properties. First, incentive compatibility requires that the scientist receive an expected bonus equal to z_L (say in the form of expected prize winnings) whenever he follows the practical strategy, which happens with probability $(\alpha + (1 - \alpha)\omega)$. Second, the scientist's participation constraint requires that the ex ante expected wage be equal to at least $R + (1 - \alpha)\omega z_L$, to compensate him for the $(1 - \alpha)\omega$ probability that he winds up following the practical strategy in a state of the world when it is not his favorite.

Putting it all together, the ex ante expected wage bill for an academic scientist is now given by:

$$w_a = \max\{R + (1 - \alpha)\omega z_L, (\alpha + (1 - \alpha)\omega)z_L\}. \quad (32)$$

And as noted, with this incentive scheme the probability that an academic scientist pursues the practical strategy is now increased to $(\alpha + (1 - \alpha)\omega)$. Everything in the private sector remains exactly as before: the wage $w_p = R + (1 - \alpha)z$, and the scientist always follows the practical strategy.

It is clear that if z_L is close to zero, academia with this particular incentive scheme can be preferred to academia without incentives: the added wage cost is minimal, but there can be a meaningful increase (by an amount $(1 - \alpha)\omega$) in the probability that the practical strategy is undertaken. At the same time, it can never make sense to try to use more powerful incentives to induce academic scientists to *always* follow the practical strategy, even when their aversion to it is strong (i.e., given by z_H instead of z_L). This is because the ex ante cost of such higher-powered incentives would be:

$$w_a = \max\{R + (1 - \alpha)z, z_H\} \geq w_p. \quad (33)$$

with the inequality being strict if $z_H > R + (1 - \alpha)z$, which is the case so long as R is not too large.

In other words, if we are in the late stages of a research program, and it is important to always have scientists working on the practical strategy, this is more efficiently accomplished in the private sector, where they can simply

be compelled to do so, rather than in academia, where this behavior has to be elicited by a system of high-powered incentives.²⁰ At the same time, the combination of an academic environment with some relatively low-powered incentives can be more efficient than the private sector in the early stages of a research program. The appeal of such a combination is that it gently nudges those with only a mild aversion to the practical strategy in the right direction, while leaving creative independence to those who value it the most.

The bottom line from this exercise is that our basic conclusions are robust to the introduction of some form of incentives in academia. At the same time, the model also suggests that to the extent that such incentives exist, it probably makes sense for them to be relatively low-powered.

6 How Academia Fosters Experimentation

In this section, we explore the hypothesis that academia encourages economically valuable experimentation by scientists. The basic story is as follows. Individual scientists know more than entrepreneurs about the right way to attack a given problem. Moreover, their incentives are inherently well aligned, in the sense that, all else equal, they get more utility from adopting an approach that is likely to be successful. Hence if left alone in the early stages of a research line, scientists will naturally tend to adopt research strategies that are responsive to their private information. This in turn has benefits for the later stages of the line, particularly if these later stages are located in the private sector, since now private-sector entrepreneurs can piggyback on the previously-generated information and thereby do a better job of directing the scientists that they manage.

However, this line of argument raises an obvious question: if later-stage private-sector research benefits from the information generated by earlier-stage experimentation, why wouldn't a private-sector firm that owns an idea at an early stage simply adopt a hands-off policy, and let scientists do what they want? In other words, if experimentation is value-enhancing, why doesn't the private sector have the right incentives to pursue experimentation, thereby rendering academia unnecessary?

The answer again has to do with the inability of the private sector to pre-commit to not interfering with scientists. In particular, suppose that some research strategies are simply easier for an entrepreneur to meddle with once they are underway, perhaps because the entrepreneur is more familiar with the underlying technology. In this case, a scientist operating in the private sector faces a tension in deciding which strategy to pursue: he can either follow his private information and choose the strategy that he thinks is most likely to succeed, or he can choose the strategy that is least likely to lead to meddling on the part of the entrepreneur. As a result, even if private-sector firms allow

²⁰Note that this same logic also implies that if z is non-stochastic, there is no role for incentives in academia. Thus we have effectively endogenized the no-incentives assumption that we made about academia in the course of developing the basic model.

scientists to choose their research strategies, these choices may no longer reflect the scientists' private information in the desired fashion.²¹

6.1 Basic setup

Consider again the variant of our basic model with $n = 1$, and to keep things especially simple, set the number of stages $k = 2$. As before, scientists can pursue two alternative strategies at each stage, but now either strategy may be practical with some probability. Let S_1 and S_2 denote the two strategies. There are two possible states of nature, state 1 and state 2, and, importantly, the same state persists across both stages of the research line. The state is observed by individual scientists, but not by private-sector entrepreneurs.

Both the probability of success of a given strategy, as well as the utility it yields to the scientist, depend on two factors: i) the state of nature; and ii) the extent to which an entrepreneur intervenes in the strategy once it is underway. Moreover, there is an asymmetry across the two strategies, in that the entrepreneur has the ability to intervene in S_1 , but not in S_2 . This could be, e.g., because strategy S_1 is more conventional, and uses a technology that is more widely familiar to non-scientists.

In state 1, which occurs with probability $s > \frac{1}{2}$, strategy S_1 is the generally more promising of the two. However, once a scientist gets started on strategy S_1 in state 1, the entrepreneur can further increase the likelihood of success by meddling—i.e., by forcing the scientist to go about strategy S_1 in a very particular way. While such intervention raises the odds of success, it also reduces the utility that the scientist derives from his work, to the point that he will prefer to adopt strategy S_2 instead, since the entrepreneur never meddles in S_2 .

Denote the probabilities of success in state 1 under strategy S_1 with and without meddling by $prob(S_1, m \mid state1)$ and $prob(S_1, nm \mid state1)$ respectively, and denote the probability of success in state 1 under strategy S_2 by $prob(S_2 \mid state1)$. Similarly, denote the associated utilities derived by the scientist by $u(S_1, m \mid state1)$, $u(S_1, nm \mid state1)$ and $u(S_2 \mid state1)$. We assume that:

$$prob(S_1, m \mid state1) = h > prob(S_1, nm \mid state1) = p > prob(S_2 \mid state1) = q \quad (34)$$

and

$$u(S_1, nm \mid state1) = 0 > u(S_2 \mid state1) = -z/2 > u(S_1, m \mid state1) = -z. \quad (35)$$

In state 2, which occurs with probability $(1 - s)$, strategy S_2 is the one that maximizes both the probability of success and the scientist's utility, whereas S_1 is both less promising and less desirable for the scientist, all the more so when

²¹This logic is similar to that in Stein (2002), who argues that individual agents will have less incentive to produce valuable soft information when they work in a hierarchical setting and do not have control rights.

accompanied by meddling. Thus we have:

$$\text{prob}(S_2 \mid \text{state2}) = p > \text{prob}(S_1, m \mid \text{state2}) = \text{prob}(S_1, nm \mid \text{state2}) = q \quad (36)$$

and

$$u(S_2 \mid \text{state2}) = 0 > u(S_1, nm \mid \text{state2}) = -z/2 > u(S_1, m \mid \text{state2}) = -z. \quad (37)$$

Again, absent meddling, the scientist prefers to work on strategies that are likely to be successful (S_1 in state 1 and S_2 in state 2) but he also values independence enough that he would rather pursue the less promising strategy S_2 in state 1 in order to avoid meddling by the entrepreneur.

6.2 Strategy choice in academia and the private sector

If the project is located in the private sector, the timing of the stage game between the entrepreneur and the scientist can be described as follows. First, the entrepreneur decides whether to allocate authority over the choice of strategy to herself or to the scientist. Next, if the choice of strategy is delegated to the scientist, he observes the state of nature and then makes his decision. Finally, if the scientist opts for strategy S_1 , the entrepreneur chooses whether or not to meddle. The key assumption here is that the entrepreneur cannot commit not to meddle once the scientist has chosen strategy S_1 .

With these assumptions in place, it is easy to describe the solution to the stage game.

Lemma 5: *Suppose that stage 1 is located in the private sector. Then even when granted the authority to pick his research strategy, the scientist will not reveal the true state of nature through his choice.*

Proof: Suppose that the true state is state 1. If the scientist reveals the state by selecting strategy S_1 , the entrepreneur will meddle, since this is the ex post optimal thing for her to do—it maximizes the probability of success—and she cannot commit to do otherwise. But then the scientist’s utility is $-z$. If the scientist instead selects strategy S_2 , her utility is higher, at $-z/2$. So independent of the state of nature, the scientist always chooses strategy S_2 when he works in the private sector.

Corollary: *Suppose that stage 1 is located in the private sector. Then it is optimal for the entrepreneur to retain authority over strategy choice, to always impose strategy S_1 , and then to meddle.*

Proof: The corollary follows immediately from the preceding lemma and from the fact that by unconditionally imposing S_1 followed by meddling, the entrepreneur achieves a probability of success of $(sh + (1 - s)q)$, which is greater than the probability $(sq + (1 - s)p)$ that comes from unconditionally imposing S_2 .

Thus it is impossible to sustain experimentation in the private sector. Knowing that the scientist will never condition his strategy on his private information when faced with the threat of meddling, the best the entrepreneur can do is to unconditionally impose the strategy with the higher prior probability of success, i.e., S_1 . In contrast, academia is more friendly to experimentation:

Proposition 5: *The scientist will reveal the state of nature through his choice of research strategy at stage 1 if and only if stage 1 is located in academia.*

Proof: If stage 1 is located in academia, then the scientist never has to worry about meddling. Therefore, he maximizes his utility by picking strategy S_1 in state 1 and strategy S_2 in state 2.

Again, the key advantage of academia is precommitment—in this case, the precommitment not to meddle once the scientist has selected strategy S_1 . The private sector’s inability to precommit in this way is its one weakness as an organizational form.

6.3 The optimal organization of research

With only two stages in the research line, there are four possible organizational structures to be considered: (i) private sector at both stages 1 and 2 (which we refer to as pp); (ii) academia at both stages 1 and 2 (which we refer to as aa); (iii) academia at stage 1 followed by the private sector at stage 2 (which we refer to as ap); and (iv) private sector at stage 1 followed by academia at stage 2 (which we refer to as pa). To further simplify the analysis, we consider the case where the scientist’s disutility parameter z is negligibly small compared to the value V of a successful line. This allows us to think of wages in both academia and the private sector as being effectively equal to zero, and therefore to ignore wages in what follows—unlike in prior sections of the paper, wages no longer play a critical role in this version of the model.

Under structure pp , we know from the previous subsection that the entrepreneur always imposes strategy S_1 followed by meddling. If it is state 1, this leads to a probability of success of h in each of the two stages, so that the overall probability of success is h^2 . If it is state 2, the probability of success in each of the two stages is q , so that the overall probability of success is q^2 . Since state 1 occurs with probability s , we have that the ex ante value of the line under structure pp is equal to:

$$\Omega^{pp} = (sh^2 + (1 - s)q^2)V. \quad (38)$$

Under structure aa , the scientist chooses strategy S_j in state j at both stages. According to our above assumptions, this means that the probability of success at any stage, and in either state of nature, is given by p . So the ex ante value of the line is equal to:

$$\Omega^{aa} = (sp^2 + (1 - s)p^2)V = p^2V. \quad (39)$$

Note that under the first two structures, there is no informational spillover across the two stages. This is no longer the case under structure ap , since the revelation of the true state of nature at stage 1 under academia now increases the expected value of the line when it moves from academia to the private sector. If state 1 prevails at stage 1, the scientist chooses S_1 —since academia protects him against meddling—and without meddling this choice of strategy leads to success with probability p . Then, having learned the true state from observing the academic strategy in stage 1, the entrepreneur at stage 2 will choose to impose strategy S_1 followed by meddling, which in turn leads to success at stage 2 with probability h . The overall probability of success corresponding to this sequence of events is thus ph . If instead state 2 prevails at stage 1, the scientist chooses S_2 , which leads to success with probability p . Having learned the true state, the entrepreneur will then impose S_2 at stage 2, which again leads to success with probability p . The overall probability of success corresponding to this sequence of events is p^2 . So the ex ante value of the line under structure ap is equal to:

$$\Omega^{ap} = (sph + (1 - s)p^2)V. \quad (40)$$

Finally, under structure pa , logic similar to that above implies that the ex ante value of the line is given by:

$$\Omega^{pa} = (sph + (1 - s)qp)V. \quad (41)$$

Comparing the expected payoffs across the four structures, we obtain the following results. First, the ap structure strictly dominates both the pa structure (since $p > q$), as well as the aa structure (since $h > p$). The reason is that ap makes it possible to first learn the state of nature under academia at stage 1, and then to use this information to direct researchers at stage 2. In contrast, it can never make sense to locate the project in academia at the second stage, because this precludes ex post efficient intervention by the entrepreneur, and because there is no further value to information beyond stage 2.

Second, the ap structure is preferred to the pp structure whenever:

$$sh(h - p) < (1 - s)(p^2 - q^2). \quad (42)$$

The left hand side of this inequality reflects the value of having the entrepreneur meddle in strategy S_1 at stage 1. The right hand side reflects the value of the information produced in academia at stage 1. When the latter exceeds the former, it is better to locate the first stage of the research line in academia. This is more likely to be the case as s falls towards $\frac{1}{2}$, since a smaller s means that there is more ex ante uncertainty about the state of nature, and hence more to be learned in academia.

It should be noted that the inequality is independent of V , and does not have anything to do with wage differentials, since wages in both academia and the private sector have been set to zero. In other words, the rationale here for locating the first stage of the research in academia is distinct from that emphasized in previous sections. Nevertheless, the broad theme of the paper

continues to apply: the fundamental tradeoff between academia and the private sector is the tradeoff between creative control and directedness. And creative control tends to be more valuable in the early stages of a research line, while directedness tends to be more valuable later on.

With wages set to zero, the pp structure will always be positive-NPV, even if it is dominated by ap . Or said differently, a private-sector firm can always afford to make a non-zero bid for the research line prior to stage 1. Thus we have:

Proposition 6: *Suppose that $sh(h - p) < (1 - s)(p^2 - q^2)$. Then the socially optimal policy is to have the first stage of the research line located in academia, and the second stage located in the private sector. However, under early privatization, both stages are located in the private sector.*

7 Discussion

7.1 Are wage differentials quantitatively important?

With the exception of the model in Section 6 just above, all of our results hinge on there being a wage differential between academia and the private sector. In particular, academic scientists must be willing to work for lower wages than their private-sector counterparts, because they value creative freedom. While this assumption fits qualitatively with both casual observation, as well as with the evidence in Stern (2004) mentioned in the Introduction, it can reasonably be asked whether real-world wage differentials are *quantitatively* large enough to justify making them the centerpiece of our theory.

Stern's (2004) estimates—which, again, are based on multiple job offers to entry-level PhD scientists—suggest differentials on the order of roughly 20 to 30 percent of salary. These are certainly economically significant differences, though perhaps not enormous ones. However, for several reasons, we believe that a superficial glance at these sorts of numbers may lead one to underestimate the actual importance of academic vs. private-sector wage differentials.

First, consider Stern's multiple-job-offer methodology. This approach is attractive, in that it allows one to control for differences in aptitude across job candidates. But by its nature, it only reveals the wage differential for those types who are “on the cusp”, in the sense of being willing to entertain both academic and private-sector jobs. It seems likely that there are more extreme types for whom the required wage premium to go the private sector would be much higher, but who are never observed receiving private-sector offers.

In the terminology of our model, this amounts to saying that there is heterogeneity across scientists in the disutility parameter z . In the presence of such heterogeneity, the benefits associated with academia are not adequately summarized by the observed wage differential. To take an extreme example, suppose that there are two types of scientists: some who are moderately talented and who are willing to go to the private sector at a 25 percent wage

premium, and others who are extremely talented but who would not be willing to go to the private sector at any wage—i.e., who have an infinite value of z . The observed wage differential will be 25 percent, but this does not fully capture the benefits of academia, since in addition to allowing for the hiring of moderately-talented scientists at a discount, it also represents the only way to ever hire extremely-talented scientists.

A second observation is that, as we have cast it, the model understates the total wage bill associated with the private-sector form, because we have assumed away any wages paid to private-sector research supervisors. Since it is impossible to have the benefits of focus in the private sector without such supervision, this is a significant omission if one wants to begin taking the magnitudes in the model seriously.

Finally, it should be noted that the “offsprings” version of the model in Section 4 has the potential to greatly lever up the effects associated with even relatively modest wage differentials. In this setting, the costs of higher wages in the private sector show up not only directly, but also indirectly, in the form of all the foregone basic research opportunities that might have been undertaken in an academic environment.

7.2 Empirical implications

Both the offsprings version of the model in Section 4, as well as the experimentation version in Section 6, suggest the following empirical implication of our theory: once an idea moves into the private sector, it will be developed along relatively narrow lines. That is, the private sector’s ownership of a given idea *will not yield as diverse an array of useful next-generation ideas as would be generated in academia.*

This implication seems to fit with the broad spirit of recent empirical work by Kaplan, Sensoy and Stromberg (2005). They study the life-cycle evolution of 49 venture-capital-backed firms, beginning with their first business plans, and continuing until three years after these firms have gone public. The firms in their sample are largely in high-technology industries, with the vast majority coming from either the biotechnology or software/information-technology sectors. Perhaps the most striking of Kaplan et al’s findings is the extent to which these innovative firms stick to their original business plans, and do not branch off into other lines of business. As they put it:

“While the companies grow dramatically, their business models or core businesses are remarkably stable. Only one firm changes its core line of business over the sample period.”

Although the results of Kaplan, Sensoy and Stromberg (2005) are suggestive, they hardly represent a sharp test of the theory. In order to create such a test, one would ideally like to find a source of exogenous variation in the extent to which ideas get taken private. One could then use patent-citation data to see if, e.g., ideas which get taken private earlier are cited by a narrower set of next-generation patents. We hope to explore this hypothesis in future work.

7.3 Connection to the literature

What is the role of academia in the innovation process? One common answer is that because of knowledge spillovers and imperfect IPR protection, the economic value associated with certain kinds of ideas cannot be fully appropriated by the developers of these ideas, leading to private-sector underinvestment—hence the need for public funding of such “basic” research (Nelson (1959), Arrow (1962)).²² While this story certainly has a good deal of merit, it is becoming increasingly difficult to draw an unambiguous connection between the “basicness” of a line of research and the degree of appropriability of the resulting output. For example, Howitt (2003) mentions an NSF survey which finds that more than 22 percent of all basic research in the US during the period 1993-1997 was—according to the NSF’s definition—performed by private enterprises. In our model, by contrast, the relevant notion of basicness has nothing to do with appropriability, but rather corresponds to the number of stages remaining until a commercial payoff can be realized.

Going beyond the traditional appropriability arguments, Dasgupta and David (1994) take a broader and more institutional view of the role of academia. They emphasize that, as compared to the private sector, academia has a variety of distinctive rules, norms and incentives that reward the production and rapid diffusion of knowledge. These include peer review, priority rules, and rewards based on the impact of publications (as measured, e.g., by citations).

Our work can be thought of as fitting in to the sort of institutional framework advocated by Dasgupta and David (1994). However, we focus almost exclusively on a single institutional attribute of academia, namely the commitment that it embodies to allowing individual scientists to pursue their own preferred research strategies. In so doing, we largely set aside many of the other features highlighted by Dasgupta and David, including incentive schemes.²³

By emphasizing the commitment role of academia, our model implicitly offers a particular rationale for the tenure system, which has been prevalent for more than half a century, especially in more research-oriented universities (McPherson and Schapiro (1999)). Our interpretation of tenure differs from that in Carmichael’s (1988) well-known contribution: in Carmichael’s model, only incumbent scientists are well-enough informed to evaluate potential new hires, and tenure serves to reassure the incumbents that the new hires will not ultimately displace them.²⁴

The focus on academia as a commitment device can also be found in the recent work of Lacetera (2005), developed contemporaneously with this paper.

²²Nelson (1959) also argues that the public sector has a longer time horizon, which makes it more willing to undertake longer-term and more uncertain projects.

²³In the spirit of Hart and Holmstrom (1987), one might ask the following: if all that distinguishes academia from private firms are their respective incentive systems, why do we need two separate institutional entities to solve the underlying contracting problem, when a more sophisticated incentive scheme would presumably also do the job?

²⁴Of course, the goals of the two papers are very different. Unlike us, Carmichael does not seek to understand the role of academia, but rather takes its existence as given and tries to rationalize one of its distinctive features.

Like we do, Lacetera adopts a control-rights perspective. However, unlike us, Lacetera does not model research as a multi-stage process. For our purposes, this multi-stage feature of the model is crucial, because it allows us to show that academia is most useful in the early stages of a research program, while the private sector tends to do better in the later stages. It also allows us to compare the socially optimal transition policy to early privatization, and to draw out the associated welfare implications.

Another recent paper which does model research as a multi-stage process is Hellmann and Perotti (2004). They contrast the free flow of ideas in academia with the more controlled informational exchange that occurs in private firms. In particular, they model a commercially attractive new research program as consisting of two stages. The open exchange of ideas in academia maximizes the probability of completing the second stage, and therefore of innovating once the first stage has been successful. But this open structure also raises the risk of the first-stage idea being stolen. Hellmann and Perotti view the private firm as an institution that guards against such stealing, by carefully recording the property rights attached to first-stage ideas. Their paper shares with ours the goal of endogenizing the choice of academic versus private-sector research. However, instead of focusing as we do on control allocation, it emphasizes incentives to share information. And, in contrast to our analysis, it stresses the commitment powers of the private sector (in terms of its ability to restrict information flows) rather than those of academia (in terms of academic freedom).

Finally, seeing organizations as differing in terms of the allocation of authority is of course not new when talking about private firms. A key element here is that individuals value creative control and are therefore ready to work at lower wages in return for more authority. This latter element is not entirely new either: Hart and Holmstrom (2002) stress it when comparing focused firms with conglomerates, arguing that the former are able to pay lower wages because of their greater commitment to pursuing the goals of their employees.

8 Concluding remarks

This paper has provided a framework for evaluating the pros and cons of academic as opposed to private-sector research. We have argued that even in a world with full IPR protection at all stages of the research process, academia—by virtue of its commitment to leaving creative control in the hands of scientists—can play a valuable role in fostering multi-stage research lines that would not be viable entirely in the private sector. Moreover, we have shown that it is possible for ideas to be privatized sooner than is socially optimal, with negative consequences for the overall rate of innovation. This latter point echoes some of the concerns raised in the policy literature about the potential for an “anti-commons” effect due to excessive early-stage IPR protection.

In terms of directions for further research, it might be interesting to study incentives more carefully in our framework. To take just one example, what are the pros and cons of citation-based rewards and promotions in a world where

citations are only a noisy indicator of the value of a research contribution? Do such citation-based incentives help to focus scientists on the right kinds of projects, or do they simply tend to induce inefficient fads or bandwagons into the research process?

Finally, one might try to enrich our framework by giving a more important role to academic openness than we have in this paper. As stressed above, the defining characteristic of academia in our model is the delegation of authority to scientists, much more than the free flow of ideas across academic institutions. In reality, it seems that openness is also a central attribute of academia, and it would be nice to better understand the complementarity between academic freedom and openness.

9 Appendix

Proof of Lemma 3: Let

$$f(x) = x - (w_p/\beta)(1 + \log(\beta x/w_p)) \quad (43)$$

and

$$g(x) = \alpha x - (w_a/\beta)(1 + \log(\alpha\beta x/w_a)). \quad (44)$$

Then for any $i \in \{1, \dots, k\}$, we know from (23) and (24) above that

$$E(\pi_i^{p*}) = f(\Pi_{i+1}) \quad (45)$$

and

$$E(\pi_i^{a*}) = g(\Pi_{i+1}). \quad (46)$$

Now, ignoring integer problems, if $(i-1)$ is the last transition stage between academia and the private sector (in fact we will show that there is only one transition stage), we must have:

$$f(\Pi_i) = g(\Pi_i), \quad (47)$$

where

$$\Pi_i = \max\{E(\pi_i^{p*}), E(\pi_i^{a*})\} = E(\pi_i^{p*}), \quad (48)$$

where the last equality follows from the fact that private sector management is optimal from stage i onwards.

We want to compare $f(x)$ and $g(x)$ over the range of $x = E(\pi_i^{p*})$'s which takes into account the fact that the optimal numbers of researchers in academia or in the private sector at all stages, and in particular at stage i , are non-negative. Now, using the analogs of (18) and (21) for stage i , we have:

$$n_i^p = (\log(\beta E(\pi_i^{p*})/w_p))/\beta, \quad (49)$$

and

$$n_i^a = (\log(\alpha\beta E(\pi_i^{p*})/w_a))/\beta, \quad (50)$$

so that we must have

$$E(\pi_i^{p*}) = \Pi_i > \max\left\{\frac{w_p}{\beta}, \frac{w_a}{\alpha\beta}\right\} \quad (51)$$

in order for both, n_i^p and n_i^a to be positive.

Next, consider the first derivatives of $f(x)$ and $g(x)$. We have:

$$f'(x) = 1 - \frac{w_p}{\beta} \frac{1}{x}; g'(x) = \alpha - \frac{w_a}{\beta} \frac{1}{x}, \quad (52)$$

which, together with (51), implies that the two derivatives $f'(x)$ and $g'(x)$ are strictly positive over the relevant range of $x = E(\pi_i^{p*})$.

Both curves $f(x)$ and $g(x)$ are convex. Moreover one can easily show that

$$f(x) > g(x) \quad (53)$$

for x sufficiently large. But this in turn implies that the two curves can cross at most once over the range where both are upward-sloping.

This, together with the fact that $E(\pi_i^{p*}) = \Pi_i$ is strictly increasing in i , allows us to establish the lemma. More specifically, if \bar{x} denotes the intersection between the two curves on their upward-sloping parts, we already know that

$$\bar{x} < V, \quad (54)$$

otherwise research would always be performed in academia at all stages, contrary to our assumption that it is optimally performed in the private sector at stage k . Now, either

$$\bar{x} < E(\pi_0^{p*}), \quad (55)$$

in which case it is socially optimal that research be performed at all stages in the private sector, or

$$\bar{x} \in (E(\pi_0^{p*}), V), \quad (56)$$

in which case it is socially optimal that research be run by academia for $i \leq i^*$ and in the private sector for $i > i^*$, where the (unique) transition stage is simply defined by:

$$\bar{x} = E(\pi_{i^*}^{p*}). \quad (57)$$

Proof of Proposition 5: From the perspective of a social planner, there is a unique transition point i_p^* , such that it is optimal for stage i_p^* to be the first stage of the parent line conducted in the private sector. This transition point i_p^* is the smallest value of i that satisfies:

$$\hat{x} = E(\pi_{i^*}^{p*}), \quad (58)$$

where \hat{x} is the unique intersection between the curves

$$f(x) = x - (w_p/\beta)(1 + \log(\beta x/w_p)) \quad (59)$$

and

$$h(x) = \alpha x + (1 - \alpha)p_r\gamma\Omega_o^* - (w_a/\beta)(1 + \log(\alpha\beta x/w_a)), \quad (60)$$

that lies on the upward sloping part of $f(x)$. (Note that $h(x)$ is equal to $g(x)$ plus the expected value $(1 - \alpha)p_r\gamma\Omega_o^*$ from creating new lines in academia).

Consider the three curves $f(x)$, $g(x)$ and $h(x)$ and the intersections \bar{x} between $f(x)$ and $g(x)$ and \hat{x} between $f(x)$ and $h(x)$, which correspond respectively to the basic model without offspring lines and to the extended model with offsprings. One can immediately see that

$$\hat{x} > \bar{x}, \quad (61)$$

which in turn immediately implies that the optimal transition from academia to private sector research occurs later than in the basic model without offsprings.

References

- [1] Acemoglu, Daron, Michael Kremer, and Atif Mian (2003), “Incentives in Markets, Firms and Governments”, NBER working paper 9802.
- [2] Aghion, Philippe, and Patrick Bolton (1992), “An Incomplete Contract Approach to Financial Contracting”, *Review of Economic Studies* 59, 473-494.
- [3] Aghion, Philippe, and Jean Tirole (1997), “Formal and Real Authority in Organizations”, *Journal of Political Economy* 105, 1-29.
- [4] Agrawal, Ajay, and Rebecca Henderson (2002), “Putting Patents in Context: Exploring Knowledge Transfer from MIT”, *Management Science* 48, 44-60.
- [5] Arrow, Kenneth (1962), “Economics of Welfare and the Allocation of Resources for Invention”, in R. Nelson ed., *The Rate and Direction of Inventive Activity*, Princeton: Princeton University Press.
- [6] Carmichael, H. Lorne (1988), “Incentives in Academics: Why Is There Tenure?”, *Journal of Political Economy* 96, 453-472.
- [7] Cockburn, Ian, and Rebecca Henderson (1998), “Absorptive Capacity, Coauthoring Behavior, and the Organization of Research in Drug Discovery”, *Journal of Industrial Economics* 46, 157-182
- [8] Dasgupta, Partha, and Paul David (1994), “Towards a New Economics of Science”, *Research Policy* 23, 487-521.
- [9] Grossman, Sanford, and Oliver Hart (1986), “The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration”, *Journal of Political Economy* 94, 691-719.
- [10] Guedj, Ilan (2004), “Ownership versus Contract: How Does Vertical Integration Affect Investment Decisions? Evidence from Pharmaceutical R&D Projects”, MIT working paper.
- [11] Hart, Oliver (1995), *Firms, Contracts and Financial Structure*, Oxford: Oxford University Press.
- [12] Hart, Oliver, and Bengt Holmstrom (1987), “The Theory of Contracts”, in T. Bewley ed., *Advances in Economic Theory*, Cambridge: Cambridge University Press.
- [13] Hart, Oliver, and Bengt Holmstrom (2002), “A Theory of Firm Scope”, Harvard working paper.
- [14] Hart, Oliver, and John Moore (1990), “Property Rights and the Nature of the Firm”, *Journal of Political Economy* 98, 1119-1158.

- [15] Hart, Oliver, Andrei Shleifer, and Robert W. Vishny (1997), “The Proper Scope of Government: Theory and an Application to Prisons”, *Quarterly Journal of Economics* 112, 1126-1161.
- [16] Heller, Michael, and Rebecca Eisenberg (1998), “Can Patents Deter Innovation? The Anticommons in Biomedical Research”, *Science* 280, 698-701.
- [17] Hellmann, Thomas, and Enrico Perotti (2004), “Circulation of Ideas: Firms versus Markets”, Stanford working paper.
- [18] Henderson, Rebecca, Adam Jaffe, and Manuel Trajtenberg (1998), “Universities as a Source of Commercial Technology: A Detailed Analysis of University Patenting, 1965-1988”, *Review of Economics and Statistics* 80, 119-127.
- [19] Holmstrom, Bengt (1999), “Managerial Incentive Problems: A Dynamic Perspective”, *Review of Economic Studies* 66, 169-182.
- [20] Howitt, Peter (2003), “The Economics of Science and the Future of Universities”, *The Timlin Lecture 2000*, University of Saskatchewan.
- [21] Kaplan, Steven, Berk Sensoy, and Per Stromberg (2005), “What Are Firms? Evolution From Birth to Public Companies”, University of Chicago working paper.
- [22] Lacetera, Nicola (2005), “Different Missions and Commitment Power: An Institutional View of Industry-University Relations”, MIT working paper.
- [23] Lach, Saul, and Mark Schankerman (2004), “Incentives and Invention in Universities”, London School of Economics working paper.
- [24] McPherson, Michael, and Morton Schapiro (1999), “Tenure Issues in Higher Education”, *Journal of Economic Perspectives* 13, 85-98.
- [25] Merton, Robert K. (1973), *The Sociology of Science: Theoretical and Empirical Investigations*, Chicago: University of Chicago Press.
- [26] Mishra, Raja (2005), “Cancer Research Has Many Healthy Side Effects”, *Boston Globe*, February 8, page F1.
- [27] Murray, Fiona, and Scott Stern (2004), “Do Formal Intellectual Property Rights Hinder the Free Flow of Scientific Knowledge? An Empirical Test of the Anti-Commons Hypothesis”, Northwestern working paper.
- [28] Nelson R. (1959), “The Simple Economics of Basic Scientific Research”, *Journal of Political Economy* 67, 297-306.
- [29] Stein, Jeremy C. (2002), “Information Production and Capital Allocation: Decentralized vs. Hierarchical Firms”, *Journal of Finance*, 57 1891-1921.

- [30] Stern, Scott, (2004), “Do Scientists Pay To Be Scientists?”, *Management Science*.
- [31] Wysocki, Bernard (2004), “Columbia’s Pursuit of Patent Riches Angers Companies”, *Wall Street Journal*, December 21, page A1.