

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

Philippe Aghion

Harvard University, NBER, CEPR and CIAR

Mathias Dewatripont

ECARES, Universite Libre de Bruxelles and CEPR

Jeremy C. Stein

Harvard University and NBER

January 9, 2008

Abstract

We develop a model that clarifies the respective advantages and disadvantages of academic and private-sector research. Rather than relying on lack of appropriability or spillovers to generate a rationale for academic research, we emphasize 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 numerous institutions, as well as to Estelle Cantillon, Micael Castanheira, Georg Kirchsteiger, Nicola Lacetera, Josh Lerner, Andreu Mas-Colell, Fiona Murray and Nicolas Sahuguet 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, as articulated by Nelson (1959) and Arrow (1962), 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.

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 1998.¹

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 a shift of research activity to the private sector should be efficiency-enhancing. However, the trend has been controversial, particularly in fields like 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.”²

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. Unlike the traditional literature, our model does not rely on imperfect appropriability to generate a role for academia. This is not to say that appropriability considerations are not important in some cases; certainly, there are many sorts of early-stage ideas where IPR protection is simply not feasible. 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. Moreover, while arguments based on imperfect appropriability make clear the need for basic research to be subsidized, they are less clear on why this

¹See also Henderson, Jaffee, and Trajtenberg (1998).

²See Murray and Stern (2004) for an empirical analysis of the anti-commons argument.

subsidy needs to happen in a different organizational form—i.e., in a university, as opposed to in a private corporation.

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 scientists retain the decision rights over what specific projects to take on, and what methods to use in tackling these projects. Indeed, in our model, academia effectively boils down to a commitment mechanism that ensures scientists that these decision rights will not be abrogated. 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.

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.

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 may end up working on projects that they find interesting, or prestige-enhancing, 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 primary 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, differences in IPR protection, or any of the other frictions that are usually invoked to rationalize a role for public funding of research. In so doing, we are able to offer a precise description of the potential costs associated with any anti-commons effect. If, for whatever reason, a research line transitions to the private sector sooner than is socially optimal, the inefficiency is manifested in the fact that—because of the higher wages—a private firm may employ too few scientists to work on the line, relative to what would happen in academia.

We then go on to explore several extensions of the basic model. The first of these considers the possibility of research lines “branching out”—i.e., giving rise to multiple potential directions for further work, some of which are nearer to commercial payoff than others. To be specific, suppose that there are two potentially legitimate research projects that make use of a firm’s patented idea. 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 has the decision rights, it will allocate all of its scientists to the applied project, and completely ignore the basic project.³

By contrast, if the private firm had never acquired the patent, and the ideas were left freely available to academic scientists, there would naturally tend to be some progress on both projects, as individual scientists followed their own interests. Moreover, the resources invested 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 deviation from the first best 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 basic project.

Another set of extensions looks at hybrid governance structures that lie in between the extremes of a totally authoritarian private sector and a completely laissez-faire version of academia. We analyze 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. These variations add some interesting nuances to our story, but do not alter the central message.

The remainder of the paper is organized as follows. In Sections 2 and 3,

³Note that this conclusion does not rest on any exogenously-imposed constraints regarding the scope of private firms: it is not that a firm cannot manage multiple projects at once, it simply finds it uneconomic to do so in this example, given its wage structure.

we introduce and then solve the most basic version of the model, in which at any given stage there is only one productive 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. As noted above, this implies that the primary welfare cost of going private too early is that too few scientists may end up working on the idea.

In Section 4, we introduce 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 promising. 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. Section 6 discusses the connection to related literature, and Section 7 concludes.

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. These scientists can only pursue research activities if they are hired either by an academic institution or a private-sector firm—they cannot be self-employed. Each scientist also 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.

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).⁴

We assume that the ex ante probability that a scientist prefers to follow the practical strategy is given by α . In addition, 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.⁵ 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. We suspect that the non-profit nature of academia plays a central role in this regard, a point that we develop more formally in Section 5.1 below. 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.⁶

In the baseline version of the model, we set aside the possibility of incentive schemes (either implicit or explicit) in academia. In Section 5.2, 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

⁴One can imagine other reasons why a scientist would prefer to work on the alternative project, in spite of its lower economic payoff. It may be more intellectually challenging, or it may hold the promise for greater professional prestige.

⁵More precisely, we need to avoid the limiting case where there is effectively a continuum of scientists with independent preferences. In this case, hiring n scientists in academia is functionally equivalent to hiring αn scientists in the private sector—with probability one, both yield the same amount of research effort devoted to the practical strategy—and the solution at all stages of the research line will be to go with whichever option involves lower total wages.

⁶See Hart, Shleifer and Vishny (1997) and Acemoglu, Kremer and Mian (2003) for related discussions about the role of non-profits.

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 scientists have no 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 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 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 buying laboratory equipment that is only 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.

3 The Case of a Single Research Line

The starting point for our analysis is the case of a single research line. We begin 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 tradeoff 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:

$$E(\pi_k^p) = pV - w_p. \quad (1)$$

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

Next, 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:

Proposition 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:

$$\Omega(\text{allprivate}) = p^k V - w_p \sum_{j=1}^k p^{j-1}, \quad (11)$$

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

Proposition 2: *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 .

More generally, it is easy to show that a necessary condition for academia to be viable one stage earlier than the private sector is that $\alpha z > R$. Intuitively, when αz is large relative to R , this tends to make academia relatively attractive 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.

For expositional simplicity, we have been assuming that the probability of success p associated with the practical strategy is a fixed constant, and does not vary with the research stage. This is obviously restrictive. However, the fundamental result in Proposition 1—namely, that it is always optimal for academia to precede the private sector—is much more general, and holds even if the success probabilities vary by stage in an arbitrary way. This can be proven by induction. Denote the probability of success at stage i by p_i . According to equation (10), if academia is preferred at stage i , it must be that $p_i \Pi_{i+1} < z$. But then the private sector cannot be preferred at stage $(i - 1)$, because this would require $p_{i-1} \Pi_i > z$, which is impossible because $\Pi_i < p_i \Pi_{i+1} < z$.

The intuition for this result is that, as research progresses from one stage to the next, the value function is always rising fast enough to increase the relative appeal of the private sector, even if the success probabilities are falling. One might at first think that if p_i is much lower than p_{i-1} , it might be optimal to be in the private sector at stage $(i - 1)$, and then revert back to academia at stage i . However, the flaw in this reasoning is that if p_i is low enough to make academia attractive at stage i , then the value Π_i of completing stage $(i - 1)$ —and thereby gaining entry to stage i —must be low as well. Indeed, Π_i must be so low in this case that the private sector cannot be optimal at stage $(i - 1)$.

3.1.2 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 1: *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:*

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

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

Proposition 3: *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; or ii) 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:

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

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.⁷ However, we will also briefly consider what happens when $w_a > 0$.

3.1.3 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 where ideas can be sold to the private sector at any point in their development, 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 a 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:

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

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

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

This condition is obviously more restrictive than the EAF condition corresponding to socially optimal transition. That is, it is harder for the research program to be ex ante positive-NPV if it is managed suboptimally than if it is managed optimally.

⁷ 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.

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 . Clearly this is an extreme case, and we do not mean to suggest that it is the most realistic one. However, it serves as a convenient benchmark that allows one to clearly see the potential benefits of academia for early-stage research.

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 is 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.⁸

3.1.4 Are wage differentials quantitatively important?

All of the above 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

⁸We 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 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 the timing of the transition to the private sector.

that a superficial glance at these sorts of numbers may lead one to underestimate the actual importance of academic vs. private-sector cost 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 to 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 willing to go to the private sector at a 25 percent wage premium, and others 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 the first type of scientists at a discount, it also represents the only way to ever hire the second type.

A further observation is that, as we have cast it, the basic version of the model understates the total wage bill associated with the private-sector form, because it assumes 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. We model the costs of supervision more explicitly in Section 5.1 below.

Finally, it should be noted that the “branching” 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 a lower-cost academic environment.

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

There is one effect which is conspicuously absent from the correlated-draws 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 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.

3.2.1 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 $\beta = -\log(1 - p) > 0$. When (18) is negative, i.e., when $\beta V/w_p < 1$, we are at a corner, with private-sector employment and profits both equal to zero. When (18) is positive, expected stage- k private sector profit is given by:

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

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. If (21) is positive, 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 , so long as $\beta\Pi_{i+1}/w_p > 1$, and $\alpha\beta\Pi_{i+1}/w_a > 1$, respectively, private-sector and academic-sector expected profits are given by:

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

and

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

Note that for academia to be viable at an earlier stage than the private sector, i.e., for $E(\pi_i^{p*})$ to be zero and $E(\pi_i^{a*})$ to be positive for some i , we require that $\alpha/w_a > 1/w_p$, which is equivalent to $\alpha z > R$. This is the same condition that we stated above for the version of the model with $n = 1$.

In the appendix, we prove the following analog to Lemma 1:

Lemma 2: *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. Suppose also that $\alpha z > R$, so that academia is viable at an earlier stage than 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 $n_{i^e-1}^p = 0$. 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 wind up in the private sector earlier, no firm would ever invest positive resources in it.

3.2.2 An additional benefit of academia

The following example 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: *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.*

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, 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

$i = 2$, 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.

4 Branching Out: The Potential For 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 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.

4.1 Adding offspring lines to the model

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 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:

$$p^{k_o}V - w_p \sum_{j=1}^{k_o} p^{j-1} < 0 \tag{25}$$

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*.⁹

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^* . To keep things simple, we focus on the correlated-draws case where $n = 1$. The independent-draws case where $\phi(n) = (1 - (1 - p)^n)$ does not add any further insight in this setting, so we omit it for the sake of brevity.

Lemma 3: *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, 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 4: *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 alternative strategy.*

Again, the noteworthy feature of the branching model is that it is now possible to have a substantially lower rate of innovation—defined in terms of the

⁹For the purposes of this section, we are implicitly assuming that the private sector cannot sell the rights to offspring lines back to academia, at least not at a price approaching their full value in an academic setting. This assumption seems to be empirically realistic, and can be defended based on: i) information asymmetries; ii) free-riding problems among universities; or iii) various bureaucratic constraints.

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 $n = 1$.

4.2 Empirical implications of the branching model

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*.

This implication seems to fit with the broad spirit of recent empirical work by Kaplan, Sensoy and Stromberg (2006). 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 out 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 (2006) are suggestive, they hardly represent a sharp test of our theory. There would clearly seem to be room for further work in this vein.

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.¹⁰ 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:

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

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

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

¹⁰We thank Rebecca Henderson for suggesting these examples.

where the wage w_{pi} is idetermined as:

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

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.

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.¹¹

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.¹² 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,

¹¹It should be emphasized that we are effectively making the best possible case for incentives in academia, by allowing bonuses to be tied directly to strategy choice. In reality, academic incentives are likely to be more general in nature, and much less directly linked to the ultimate commercial value of a research project. In this sense, our basic model, which omits academic incentives entirely, may actually be closer to capturing the truth.

¹²Note 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.

the model also suggests that to the extent that such incentives exist, it probably makes sense for them to be relatively low-powered.

6 Connection to the Literature

What is the role of academia in the innovation process? One answer is that because of knowledge spillovers and imperfect IPR protection, the 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)). While this story certainly has 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, 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 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.¹³

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

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

¹³In 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?

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.

7 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 where ideas can be sold to the private sector 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 research projects 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 and clarifies some of the concerns raised in the policy literature about the potential for an “anti-commons” effect due to too-early privatization.

In terms of directions for follow-up work, 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 introduce inefficient fads or bandwagons into the research process?

The model also has several empirical implications that might be worth investigating. Perhaps the most noteworthy comes from the branching version of the model, which predicts that a given idea will yield a more diverse array of offspring if it is located in academia, rather than the private sector. We hope to see this hypothesis tested in future work.

8 Appendix

Proof of Lemma 2: Let

$$\begin{aligned} f(x) &= x - (w_p/\beta)(1 + \log(\beta x/w_p)) \text{ if } \beta x/w_p > 1 \\ &= 0 \text{ otherwise.} \end{aligned} \tag{34}$$

Similarly, let

$$\begin{aligned} g(x) &= \alpha x - (w_a/\beta)(1 + \log(\alpha \beta x/w_a)) \text{ if } \alpha \beta x/w_a > 1 \\ &= 0 \text{ otherwise.} \end{aligned} \tag{35}$$

Since the value function Π_i is increasing in i , the proof of the lemma boils down to establishing that as x increases, there is a unique point where the $f(x)$ and $g(x)$ functions cross (other than where they both take on the value zero). Under the assumption in the lemma that $\alpha z > R$ (or alternatively, that $\alpha/w_a > 1/w_p$), it can be seen that $g(x)$ first becomes positive at a lower value of x than does $f(x)$. So if we define the differential $\psi(x) = f(x) - g(x)$, then for low values of x , we have that $\psi(x) < 0$. It is also clear that for sufficiently large values of x , $\psi(x) > 0$.

To see that there is a single crossing point, i.e., a single point where $\psi(x) = 0$, note that:

$$\psi''(x) = (w_p - w_a)/\beta x^2 > 0. \tag{36}$$

In other words, the function $\psi(x)$ is convex. This implies that if it takes on negative values for small x , and positive values for large x , it has exactly one zero. (It must be increasing when it first crosses zero, and from this point on, convexity implies that it keeps increasing, and therefore never crosses zero again.) This establishes the lemma.

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] 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.
- [5] Carmichael, H. Lorne (1988), “Incentives in Academics: Why Is There Tenure?”, *Journal of Political Economy* 96, 453-472.
- [6] Dasgupta, Partha, and Paul David (1994), “Towards a New Economics of Science”, *Research Policy* 23, 487-521.
- [7] 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.
- [8] Hart, Oliver (1995), *Firms, Contracts and Financial Structure*, Oxford: Oxford University Press.
- [9] Hart, Oliver, and Bengt Holmstrom (1987), “The Theory of Contracts”, in T. Bewley ed., *Advances in Economic Theory*, Cambridge: Cambridge University Press.
- [10] Hart, Oliver, and Bengt Holmstrom (2002), “A Theory of Firm Scope”, Harvard working paper.
- [11] Hart, Oliver, and John Moore (1990), “Property Rights and the Nature of the Firm”, *Journal of Political Economy* 98, 1119-1158.
- [12] 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.
- [13] Heller, Michael, and Rebecca Eisenberg (1998), “Can Patents Deter Innovation? The Anticommons in Biomedical Research”, *Science* 280, 698-701.
- [14] Hellmann, Thomas, and Enrico Perotti (2004), “Circulation of Ideas: Firms versus Markets”, Stanford working paper.

- [15] 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.
- [16] Holmstrom, Bengt (1999), “Managerial Incentive Problems: A Dynamic Perspective”, *Review of Economic Studies* 66, 169-182.
- [17] Howitt, Peter (2003), “The Economics of Science and the Future of Universities”, *The Timlin Lecture 2000*, University of Saskatchewan.
- [18] Kaplan, Steven, Berk Sensoy, and Per Stromberg (2006), “What Are Firms? Evolution From Birth to Public Companies”, University of Chicago working paper.
- [19] Lacetera, Nicola (2005), “Different Missions and Commitment Power: An Institutional View of Industry-University Relations”, MIT working paper.
- [20] Lach, Saul, and Mark Schankerman (2004), “Incentives and Invention in Universities”, London School of Economics working paper.
- [21] McPherson, Michael, and Morton Schapiro (1999), “Tenure Issues in Higher Education”, *Journal of Economic Perspectives* 13, 85-98.
- [22] 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.
- [23] Nelson R. (1959), “The Simple Economics of Basic Scientific Research”, *Journal of Political Economy* 67, 297-306.
- [24] Stern, Scott, (2004), “Do Scientists Pay To Be Scientists?”, *Management Science*.