

~

PROCURING KNOWLEDGE

Stephen M. Maurer
GSPP, U.C., Berkeley
2607 Hearst Ave, Berkeley, CA 94720-7320
maurer@econ.berkeley.edu

and

Suzanne Scotchmer
Dept Economics and GSPP, U.C., Berkeley
2607 Hearst Ave, Berkeley, CA 94720-7320
and NBER

July 28, 2003

ABSTRACT: There is growing public interest in alternatives to intellectual property including, but not limited to, prizes and government grants. We argue that there is no single best mechanism for supporting research. Rather, mechanisms can only be compared within specific creative environments. We collect various historical and contemporary examples of alternative incentives, and relate them to models of the creative process. We give an explanation for why federally funded R&D has moved from an intramural activity to largely a grant process. Finally, we observe that much research is supported by a hybrid system of public and private sponsorship, and explain why this makes sense in some circumstances.

JEL Classification: O310, O380

Keywords: patents, prizes, IP, government grants, contests, Bayh-Dole, CRADA

forthcoming, *Advances in the Study of Entrepreneurship, Innovation, and Economic Growth*, Vol. 15, Gary D. Libecap, editor, JAI Press (Elsevier Science) 2004.

1 Introduction

Patents and other intellectual property rights (IPRs) became stronger during the 1980s and 1990s than at any point in the preceding century. Nevertheless, other mechanisms for eliciting and investing in new ideas have remained important. This is particularly true in the public sector, where federal spending accounts for almost a third of the nation's total R&D. (The fraction is higher in most other countries.) Funding mechanisms used by the public sector routinely include in-house development, procurement through competitive bidding, and research grants to universities and promising scientists.

The public sector also uses hybrids that mix sponsorship with intellectual property. The 1980 Bayh-Dole and Stevenson-Wydler Acts authorized the patenting and licensing of federally funded research outputs, and the Federal Technology Transfer Act of 1986 authorized the formation of Cooperative Research and Development Agreements (CRADA's) between national laboratories and industry, granting intellectual property to the industrial partners.

Despite the prominence of intellectual property incentives, prize-based and other incentives remain strong, and may even be undergoing a renaissance. Beginning in the Clinton Administration, NASA and Congress have debated a prize-based approach to new space technology (Space Projects 2002). The Defense Advanced Research Projects Agency (DARPA) has offered a \$1 million prize designed to elicit a 40-fold improvement in robotic off-road vehicle design (DARPA 2003; Holden 2003). Even the private sector has begun to experiment with prizes, and in at least one offering, the prize mechanism has become institutionalized. Eli Lilly Corporation has established a company called Innocentive to elicit new ideas for solving problems in biology, chemistry, biochemistry, and materials science. "Seeker companies" post problems on Innocentive's web site and offer rewards ranging from \$5,000 to \$100,000. Scientists are invited to submit solutions, and the seeker companies pick the best ones (Innocentive 2003).

In this paper we catalog some of the many incentive schemes that have been suggested and used. We argue that the sensibleness of each scheme – indeed, the

attractiveness of intellectual property – cannot be debated without reference to a specific environment of knowledge creation. The diversity of models and incentive schemes discussed here suggest that no single incentive mechanism dominates, but that different models of knowledge creation call for different incentive schemes. There are also creative environments where the social value of an innovation is not appropriable by private firms or intellectual property rights are insufficient to cover costs.¹

The main objective of the incentive schemes we discuss is to elicit investments in new knowledge. Another objective is to do this at the lowest possible cost. The cost of investing in knowledge has two components: resource costs and the diminution of value that may arise through proprietary pricing and deadweight loss. The latter particularly afflicts intellectual property. Although raising funds for general revenue is not costless, the associated inefficiencies are usually thought to be less onerous than taxing a single market, as proprietary pricing does. This observation leads to the conjecture that, in many creative environments, some form of public procurement may be preferred.

We begin in Section 2 with a model of the creative process, and then discuss the incentive mechanisms that naturally flow from it. The efficacy of prize systems, discussed in Section 3, will depend largely on whether prizes can be made to reflect the values of delivered innovations. But even if that is possible, prize incentives, like patent incentives, are a crude mechanism for discriminating among ideas for R&D investment. In Section 4 we discuss contests that are targeted specifically to the problem of finding the innovators with the best ideas. In Section 5 we give an economic rationale for the federal grant process, in which ideas for funding originate with grantees, there is no enforcement mechanism to ensure that the grantee’s proposal is feasible, and no means to reclaim payment if the objective is not accomplished. In Section 6 we turn to the virtues of sponsorship for innovations such as nuclear fusion for energy, where it is *ex ante* unknown what the best approach is. In Section 7 we address what is probably the most controversial development in the R&D establishment, the

¹One important example is what Mokyr (2002) calls “propositional knowledge.” Even if the knowledge that “DNA is a double helix” were patentable, it would be hard to collect a royalty on every use of that knowledge.

fact that many innovations are both subsidized by public sponsors, and also receive intellectual property protection. We show that, provided industry is required to give matching funds, this system can solve the dual problems of generating enough funding for “big science”, while also tapping into industry’s expertise, to avoid ill-conceived ideas.

2 Ideas and Innovations

To compare incentive schemes from a normative point of view, we must have some notion of an efficient investment plan. How much money should be devoted to R&D? Where should R&D dollars be channeled? Which incentive system is mostly likely to achieve the desired goals?

From a business point of view, an “efficient” investment plan is the one that maximizes profit – investing in knowledge is only “efficient” if it leads to profit. However, the business point of view does not help society as a whole choose among incentive mechanisms. Profit is only a partial view of benefits, since it does not account for benefits that accrue directly to users. These are especially important for innovations funded by the public sector and put in the public domain.

To compare the efficiency of incentive mechanisms, we must first have a notion of which R&D investments themselves are efficient. The notion that underlies this paper is a standard economics notion, namely that an R&D investment is efficient if it provides higher (discounted) consumers’ surplus net of costs than any other feasible investment or investment path, at least in expectation. This is a definition that takes account of how investments in knowledge are funded. If investments are funded by intellectual property, then they generate less consumers’ surplus (or more deadweight loss) than if they are funded by public sponsors and put in the public domain.

Just as importantly, efficiency of an R&D investment, and therefore efficiency of incentive mechanisms, can only be understood within a model of the creative process. We will use the “ideas” model of O’Donoghue, Scotchmer and Thisse (OST, 1998; see also Scotchmer (2004)), which makes a clean separation between an exogenous process that generates ideas for innovations, and the decisions whether to invest in them. An

innovation requires both an idea and an investment in it.

Not all economic models of technological advance are set up to evaluate efficiency in investment, or to compare incentive mechanisms. For example, the evolutionary model of Nelson and Winter (1982) (see also Mokyr 1990, Ch. 11) is a positive model of technical change, rather than a normative model. Regardless of the incentive mechanisms in place, firms are assumed to invest in R&D when profit falls below a certain level. There is no notion of rationality that drives investment decisions.

Another class of models relies on an exogenous production function for knowledge that is commonly known to everyone (Wright 1983, Scherer 1984, Shavell and Van Ypersele (2001) and earlier models using the Poisson process, summarized by Reinganum 1989). These models do not recognize that ideas can be scarce. Indeed, an important intuition for why intellectual property may be a better mechanism than any form of government sponsorship is that it can harness the creativity that is widely dispersed among humans. Intellectual property can encourage investment in an idea that only a single person thinks of. But as we shall see, other mechanisms also have merit in such environments.

We describe an idea by a pair (v, c) (per-period social value and cost).² The variable v represents the per-period consumers' surplus with competitive supply. If the social value lasts forever, then, if the invention is in the public domain, the discounted social value is $\frac{1}{r}v$. If the invention is marketed by a proprietary firm, we will assume that the per-period profit is πv , where π is a fraction less than one. Then the proprietary profit available under a patent that lasts for discounted length T is $\pi v T$.³ We assume that the associated deadweight loss is dv per period, so that the associated deadweight loss is dvT .

Following the economics literature on contracting, we will make a distinction between the case that the per-period social value, v , or the cost of achieving it, c , is

²The interpretation of an idea will depend on how ideas are related to each other. In OST, ideas follows each other in a progression up a quality ladder; a later idea must be preceded by the earlier ones. That context is not discussed here.

³For tidier notation, we always use discounted time instead of "real time." If the real length is τ , then the discounted length is $T = \int_0^\tau e^{-rt} dt$. There is a one-to-one map between real time and discounted time, and the maximum discounted time is $\frac{1}{r}$.

observable to a third party, and the case that it is *verifiable* by an enforcing party such as a court. If it is not verifiable, then it is impossible to enforce a contract in which payments depend on that variable. This limits the kinds of incentive contracts that are available.

We begin with a preliminary comment that relates patents and prizes. If the social value of an invention is verifiable, then prizes clearly dominate intellectual property as an incentive system provided the efficiency concern is aggregate deadweight loss. Whatever the duration, T , of the proposed protection, a prize equal to $\rho(v) = \pi v T$ entails the same incentives to invent as the IP system has, but with less deadweight loss, provided the prize giver puts the innovation in the public domain. This observation remains true in other models of knowledge creation, such as where there is a knowledge production function.

If the value v is observable, an invention authority can generally do even better, by offering a prize $\rho(v)$ that is tailored to the distribution of costs. In fact, it might seem natural to base rewards directly on the R&D cost c , since that would enable a sponsor to minimize the amount of money that must be raised for rewards. In practice, however, this would not work. First, a sponsor cannot observe cost by consulting the researcher's accounting data. The economic definition of cost is the minimum cost required to achieve the result. Wasteful or inefficient efforts should not count. If the sponsor simply reimbursed the accounting costs, he might, for example, give the researcher an incentive to go to the beach under the guise of attending research conferences. Second, in most research endeavors, the laboratories and researchers' time are spread among many research projects. Overhead costs must be apportioned among the projects. No one except the researcher, and possibly not even him or her, can know how to apportion the costs.

Finally, and perhaps most importantly, some research efforts do not pay off with certainty. PhRMA, the trade organization of the pharmaceutical firms, estimates that fewer than one in five drug development efforts results in a successful drug. The failures obviously cannot be identified in advance, else the drug companies would avoid them. If only the successful drugs results in profitable IP rights or prizes, the cost

being covered by the IP rights or prizes would have to include the cost of failures as well as successes. If only the successful drugs are rewarded, they must be rewarded at least five times more than the average cost per drug discovery effort.

Conditional on v , the sponsor should have some subjective notion of the distribution of costs. An optimally chosen prize should reflect this conditional distribution. If costs for ideas with value v are lower on average for some types of innovations, then presumably the prize $\rho(v)$ should also be lower. For other types of innovations, the expected costs might be higher on average, and the prize should be higher. In making these judgments, the sponsor recognizes that, since the cost c is unobservable, a prize $\rho(v)$ smaller than the social value $\frac{v}{r}$ may not cover cost, and a prize larger than the patent value πvT may be larger than cost, and attract too many competitors. The benefit of giving lower prizes must be balanced against the possibility of discouraging some innovations.

A prize based on the value created will not be credible if the sponsor has opportunity to renege, or is thought likely to renege. This is where verifiability by a court may be important. For a prize system to work, something about the merits of the delivered invention must be observable to an enforcer, even if it is only a noisy proxy for v .⁴ It also helps if the prize giver can make himself credible in a repeat game, as was true to a large extent in 18th century and 19th century France.

The examples of incentive mechanisms that we give below address two types of creative environments: those where the creation of knowledge addresses a known need, and those where the need had not been identified, or at least articulated by a sponsor, prior to someone thinking of the idea. In the latter context, it is natural to call the idea *scarce*, in the sense that there are no substitute ideas that would address the same economic need.

The distinction between ideas that address well-known needs and ideas that are scarce has implications for what types of mechanisms can be used. If the need is known in advance, then it makes sense for a sponsor to post a prize to address it, or otherwise solicit solutions *ex ante*. The prizes would typically be contingent on

⁴Gallini and Scotchmer (2002) point out that the same problem of verifiability arises with patents, but only when they are litigated.

some performance standard dictated by the need. In contrast, an incentive system to reward scarce ideas could not very easily be established in advance, at least with a performance standard. The terms of the reward must be established *ex post*.

The distinction between well known needs and scarce ideas is, of course, fuzzy. A need can have varying degrees of specificity, from “demonstrate that Maxwell’s equations are correct” to “improve the efficiency of harnessing water power”, both discussed below.

We close this section with three examples that illustrate the scarcity of ideas, and also illustrate the growing realization that, before good R&D investments can be made, the individuals with good ideas must be identified.

The first example is NASA’s Institute for Advanced Concepts, which spends \$4 million per year funding ideas for space propulsion. The funding is in two phases. Initial grants are up to \$75,000, which is entirely for exploring an idea. The agency then funds Phase II grants of to \$500,000 to validate them. The two phases arguably correspond to the idea and the investment, although the second stage is still rather far from a usable innovation. Speculation is encouraged: the director urges audiences “Don’t let your preoccupation with reality stifle your imagination.” Not surprisingly, results have been mixed. Some grants (antimatter rockets, solar sails, space elevators) re-work ideas that have been widely discussed for decades. Other grants have elicited risky but physically plausible schemes. In fact, NASA has already picked up one idea – creating magnetic “sails” to surf the solar wind – for further development. Critics charge that another (“hydrino” propulsion) rests on pseudoscience ruled out by the known laws of physics (Reichardt 2002).

The second example concerns jet fighters. The U.S. Air Force has always recognized that an essential part of advancing aircraft capabilities is attracting new ideas. The current F-22 design effort began in 1981 with a formal request to industry for ideas. Suggestions ranged from an ultra-lightweight fighter to a 120,000 lb. “battle cruiser.” The Air Force then took the best ideas from each proposal and prepared detailed specifications. Even then, the Air Force resisted prototyping because that would reduce the number of participants and ideas. It was not until 1986 that Lock-

heed and Northrup were selected for the remaining competition (Sweetman 1991).

The third example comes from the private sector. In 2002, Google announced its first annual prize for a scalable computer program that “does something interesting” to the company’s internal library of pre-parsed web pages. The winning entry merges Google’s data with census bureau address data to create a database that can limit searches by geographic area. (Contestants granted Google a worldwide, non-exclusive license to make, sell, or use any technologies they developed (Google 2002).) Although the winning contestant was required to submit working code, the main benefit to Google resided more in the idea than in the implementation. Any programmer could have turned the idea into an innovation; it was the idea itself that had value, and that is what Google was looking for.

3 Prizes: Linking Payments to Value

An important class of mechanisms are prizes. These come in two types that correspond to the two innovative environments that arise from well-known needs and scarce ideas. *Targeted* prizes, which are posted *ex ante* and have clear performance standards, reward solutions to needs that originate with sponsors. The inventor’s idea is a solution to the sponsor’s stated need. *Blue-sky* prizes are given when no need has been stated in advance, but judges are allowed to “know it when they see it.” (In blue-sky prizes, “the sky is the limit.”) If the reward is specified in advance, it cannot be tailored to the idea. For example, Nobel prizes are a fixed award, regardless of the nature of the accomplishment. More often, as we shall see, blue-sky prizes are tailored *ex post* to the value of the innovation, but this is a difficult problem.

The problem with blue-sky prizes is how to tailor them to the value v . This same problem carries over to contests, discussed below. If a prize cannot depend on the value of the innovation that is delivered, then the sponsor cannot avoid rewards to useless innovations or rewards to the wrong innovators.

We illustrate two methods that have been used to ensure that a prize reflects the value v , namely, making the prize conditional on a verifiable performance standard, and giving the inventor the option to choose intellectual property protection instead.

In the latter case, the prize effectively becomes a patent *buy-out*. The inventor will not agree unless the prize is at least as large as the patent value πvT , and in this sense, the prize is constrained to reflect the value.

An example of how the buy-out works occurred when the French inventors of photography, Daguerre and Nahin, sold their rights in an *ex post* negotiation in 1839. The inventors received pensions totaling 10,000 francs per annum in exchange for revealing the secrets of the process at a joint meeting of the French Academies of Art and Science. Afterwards, the process was put in the public domain (Newhouse 1988). We can only presume that the French inventors received value commensurate with their invention, since they would not have accepted a prize less than the patent value.⁵

In contrast, John Hyatt chose patent protection over a prize for inventing celluloid in the 19th century (Porter 1994). Hyatt had originally invented celluloid in order to claim a prize posted by a manufacturer of billiard balls who wanted to replace ivory. However, when Hyatt realized that his invention had wider applicability, he chose patent protection instead, apparently judging the value of the patent to be greater than the prize.

Michael Kremer (1998) suggests a direct way to buy out the patent at a price that reflects the value. Since the value of the invention is likely to be observable by rival firms after the invention has been made, the patent authority can turn the patent into a prize as follows. The invention authority appropriates the patent, and auctions it to rival firms. He promises that there is a small probability that he will actually deliver the patent to the highest bidder in return for the bid price. The rivals will bid the same as if the patent would be transferred with probability one, thus revealing their valuations. The winning bid should be (close to) the private value πvT . The sponsor can divide the winning bid by πT to get an estimate of the value v . Kremer suggests that the sponsor pay the inventor a prize equal to his estimate of the social value, $\frac{v}{r}$, from general revenue. With low probability he transfers the patent to the winning bidder in return for the bid price. With high probability, he puts the

⁵Contemporary observers contrasted this procedure with what happened in England, where a better process was stifled by patent litigation (Newhouse 1988).

innovation in the public domain, thus turning the patent system into a prize system. This scheme is enforceable, provided the parameters πT are specified in advance, so that a court could enforce it.

The greatest virtue of patent buy-outs is that they allow the government to confiscate inventions with high public consequences, such as pharmaceuticals, without hurting the innovator. If anything, the Kremer scheme increases the incentive to invest, since the innovator's reward $\frac{v}{r}$ is larger than the value of the patent $\pi v T$. However, such a high prize also leads to social waste if it incites rival firms to duplicate each other's costs in a wasteful race. Under these circumstances, it may be better to choose a prize $\rho(v)$ below the social value, as discussed in the previous section.

An example of linking the prize to verifiable performance standards occurred in the Lyonnaise silk-weaving industry. Members of the Fabrique Lyonnaise could make improvements to weaving on their own initiative, and then petition a prize committee for remuneration. It is unclear whether most of these improvements would alternatively have been patentable, but the terms of the *ex post* reward evidently did not rely on that alternative. Instead, the prize committee set the terms of the rewards based on performance criteria, such as the number of weavers who adopted it (Foray and Perez 2000). Such terms allowed the committee to ensure that the prize was given only in return for value.

Targeted prizes were common in France along with blue-sky prizes. In 1795 a prize was offered for a means to preserve food to feed Napoleon's vast armies and navy. The prize was awarded in 1810 to Nicolas Appert on condition that he publish the technique and put it in the public domain (Porter 1994, p. 16). His technique, which involved heat-sterilization of food packed in bottles, is still in use. (If the solution seems obvious, we should remember that the causes of spoilage were not understood.) Other targeted prizes led to improvements of the steam engine (Porter 1994) and also water power, which led to the first water turbine (Strandh 1979).

One of the most famous targeted prizes was for discovering how to determine longitude (Sobel 1995). By the time the English prize of £20,000 was posted in 1714, similar prizes in other countries had gone unclaimed. The failure to determine longi-

tude was the principal cause of a large and growing number of naval and commercial catastrophes. No one disputed the high value of a solution, but mariners despaired of finding one. In the course of the English competition, two very different ideas for solving the problem emerged. John Harrison had the idea to engineer seaworthy clocks that would allow the mariner to keep London time, and thus compare local time to London time. Such a comparison would tell him how far East or West he had sailed. The other idea was to use lunar observations in comparison with the stars to ascertain longitude. Both were developed in parallel, but the prize eventually went to Harrison.

In the modern renaissance of prizes, the X Prize Foundation was established in 1996 with a \$10 million prize for the first private firm to carry three passengers to a suborbital height of 100 km twice within a single two week period.⁶ To date, 23 teams representing seven countries have entered (X Prize Foundation 2003). Observers report that at least three teams are likely to produce actual hardware. One effort has reported that it plans to spend \$5-7 million in the effort. Some commentators believe that several teams have a reasonable chance of claiming the prize before it expires in 2005 (Hoffman 2003).

Another modern targeted prize was \$30m announced in 1992 to develop a new, "Super Efficient Refrigerator" (Penn 1993, Zuckerman 2003). In order to receive payment the winner had to sell 250,000 copies of the new refrigerator by 1997. Although Whirlpool won the contest, they did not meet the performance standard, selling only 200,000 units, and the prize was not paid.

Finally, we mention two targeted prizes with performance standards that have been offered for inventions that have great value to the sponsor or to society, but no

⁶Long distance flight has traditionally been supported by prizes. Famous examples include the first flight across the English channel in 1909 (\$5,000), the first flight across the North Atlantic in 1919 (£10,000), and the first non-stop flight from New York to Paris in 1927 (\$50,000). The 1927 contest, which resulted in Lindbergh's flight is particularly interesting. The prize was originally offered in 1919 for the first flier to accomplish the feat before 1924. This incentive may have been insufficient given the technologies of the time; in any case, there were no takers. The prize was then renewed for an additional five year period drawing a total of nine teams into the race, although at least one of these efforts failed to obtain adequate financing. The remaining teams spent a total of \$400,000 in the effort. Eight of these aircraft either crashed on take off or disappeared. Lindbergh's successful effort cost \$25,000 (X Prize Foundation 2003, Jablonski 1972).

commercial value. The first was a German prize for demonstrating that Maxwell's equations of electromagnetism were correct. The prize was offered by the Berlin Academy in 1879, and the feat was accomplished by Henrich Hertz, albeit long after the prize expired (Bryant 1988).

A modern prize with similar character is the RSA Factoring Challenge, which offers a \$250,000 prize for the first team to factor a specified 2048-bit integer. Rewards for smaller specified numbers start at \$10,000 (RSA Factoring Challenge 2003). RSA sells a commercial implementation of public key encryption whose security rests on the computational difficulty of factoring large numbers. If the prize goes unclaimed, RSA has more confidence in the security of its system. This is a case where the sponsor is better off if it does not get what it is looking for. The successful claimant cannot profit from the discovery in any way other than the prize.

These examples show that prizes can be linked to value by (i) administering prizes in a legal environment where intellectual property protection is also available and (ii) making the prize contingent on a performance standard. These techniques can apply to either targeted prizes or blue-sky prizes, although patent protection is probably a more important backup for blue-sky prizes, as a way to avoid *ex post* hold-up.

So far these arguments have been entirely about the feasibility of linking prizes to the value of the innovation. They do not inform us about the optimal *size* of the prize. If costs are "expected" to be much lower than value, then it is sensible to set a prize much lower than the social value, even if the social value can be observed. The appropriate size of a prize should also be tailored to the number of participants that the sponsor wishes to attract.

In the simple ideas model given above, it is desirable to avoid duplication of costs. When ideas are scarce, this is not a problem – scarcity means that only one inventor can fill the market niche defined by his idea. In that case, the size of the prize should simply balance the possibility of giving up the innovation (when the prize turns out to be less than cost) against the social cost of over-rewarding the innovator. Such a balancing must be done by reference to a subjective distribution of costs, conditional

on value.

For targeted prizes, ideas are typically not scarce. There are typically several ideas that could fill the targeted need. In such an environment, simple prizes share a defect of intellectual property, even if the prize can depend on the value. There is no way to ensure that only the agent with the best idea invests, or even that a single agent invests. Those problems can be partly addressed by the contests we discuss in the next section.

4 Research Contests

A difficult problem for a public sponsor is when the prize cannot be made to reflect the value of the delivered invention. The sponsor is then in jeopardy of giving a prize for no value. Just as bad, potential inventors may refuse to take the bait, fearing that the sponsor will renege on promises. In this environment, intellectual property starts to look attractive, since the value of an intellectual property right is automatically linked to the social value of the invention.

However, we should not give up on prizes too soon. For one thing, there are many inventions for which intellectual property will be ineffectual because the social value cannot be appropriated. Examples that spring to mind are military wares and pure scientific knowledge. Fortunately, the sponsor still has options.

We discuss how the sponsor can elicit investments in innovations of high value even if the value is not verifiable *ex post* to an enforcement body. He will do this by setting up contests that involve enforceable contracts. In two of the contests we discuss, the contract terms will not involve the enforcement of any payment conditional on the observation of value. These two contests have the downside that the best contest may nevertheless lead to duplicated efforts of the contestants. The inability to make enforceable contracts in which payments depend on the delivered quality is costly.

A Simple Commitment to Pay The first commitment device is extremely simple, and also very common. The sponsor sets a prize, and commits only to give it away. He does not make the payment conditional on any specified performance, but

rather announces his objectives, and lets the contestants choose their strategies. Ex post he will choose the contestant who comes closest to meeting his objective. There is no reason for an enforcement body to be involved, except to make sure that the prize money is actually paid to someone. The sponsor cannot renege on paying the prize, and has no reason to pay it to any contestant other than the one who comes closest to his objective.

Much of the development of airplanes was motivated by such contests, especially before World War II. The Gordon Bennett, Schneider, Thompson, and Bendix prizes supported air races, typically offering purses between \$5000 to \$10,000.⁷ This promoted fierce competition from 1911 until 1939. While it is impossible to measure technological impact of these competitions with any certainty, prize winners routinely outperformed contemporary military aircraft in speed and endurance until the late 1930s. More specifically, the participation of two companies in pre-war racing – Rolls-Royce and Supermarine – pioneered technologies for advanced aircraft and engines that were eventually incorporated in World War II designs like the British Spitfire. Although much of this effort was subsidized by government loans of pilots and specially-built aircraft, other racing machines were built by private investors in search of prize money. The most famous example was the lethal (but very fast) “GeeBee” design that dominated racing in the early 1930s (Vorderman 1991, Porter 1994).

The next two contest-like mechanisms that we discuss are more complicated than a simple commitment to give away the money, but they have additional efficiency properties. As we have already mentioned, the simple commitment to pay a fixed prize has a defect that it shares with patent incentives, namely, that it provides no control over which contestants enter, or even how many. Suppose that a prize or patent is worth \$100. Suppose there are two potential contestants, one of which can enter for a cost of \$70, and the other of whom can enter for a cost of \$60. Only one will enter, since entry by both means that the expected winning of \$50 will not cover

⁷These amounts tend to understate the incentive, since a single aircraft could compete in multiple races and win multiple purses. In a few cases, aircraft were also able to compete in successive years. [Vordman (1991)]

the cost of either. But if the higher-cost firm enters, then the lower-cost firm will not. There is no guarantee that it is the lower-cost firm that enters in equilibrium. Similarly, if the two firms have different prospects for success, there is no guarantee that it is the better firm that enters. A further problem is that, depending on costs as compared to the prize, both firms may enter even if investment by a single firm would suffice. The incentive system then leads to a wasteful duplication of costs. A good incentive mechanism should avoid these forms of waste.

We now discuss two more refined contests that to some extent address these problems. Suppose that two firms have substitute ideas (v_i, c_i) , $i = 1, 2$, for how to serve the sponsor's need. We assume that, if the sponsor could observe these values, then he would choose the idea that provides the greatest social surplus, that is, the sponsor would choose firm 1's idea if $\frac{1}{r}v_1 - c_1 \geq \frac{1}{r}v_2 - c_2$, and would otherwise choose firm 2's idea. This is the firm with the lower cost if $v_1 = v_2$. The problem is that the sponsor cannot observe $(v_1, c_1), (v_2, c_2)$. The firms may have an incentive to lie about them if they think that lying will garner a profitable government contract.

A Vickrey Auction The Vickrey auction goes some distance toward eliciting efficient investment, but only works if the value can be verified *ex post*.⁸ So far as we know, this has not been used in practice. In the second-price (Vickrey) auction, the sponsor asks each prospective innovator i to report the social surplus $s_i = \frac{1}{r}v_i - c_i$ that it could provide. The firms report some values s_1, s_2 . Of course the sponsor cannot verify that either firm is reporting s_i honestly. In addition, he does not know whether a high value of the reported surplus s_i means that v_i is high or c_i is low. However this will not matter, as the promised payments will give the firms an incentive to be honest.

The sponsor chooses the firm that reports the highest net surplus. Suppose that this is firm 1. He promises a payment to the winning firm, firm 1, that is equal

⁸It is named after economist William S. Vickrey, who first expounded its remarkable properties in his (1961) paper. The simplest example is where an auctioneer wants to transfer an object to the agent with the highest valuation, but cannot observe the valuations. He asks each one to state his valuation, then gives the object to the highest bidder, but only charges a price equal to the second highest bid. In that situation, there is only one unobservable for each bidder, and nothing needs to be observed *ex post*. In the version given here, each agent has two unknown variables, cost and value. The value must be observed *ex post*, but not the cost.

to $\frac{1}{r}v_1 - s_2$. Notice that the payment depends on the *other* firm's bid, s_2 , not on the winning bidder's bid s_1 . The sponsor then asks firm 1 to invest. After paying its own costs c_1 , firm 1 ends up with the payment $\left[\frac{1}{r}v_1 - s_2 - c_1\right] = s_1 - s_2$.

The second-price auction has the following properties: Assuming that the firms report their respective surpluses honestly, then (1) the payment to the winning firm, say firm 1, will be close to (but no smaller than) the cost c_1 if s_1 is close to s_2 , and (2) the winning firm makes nonnegative profit by delivering the innovation in return for the specified payments.

We can now ask whether the premise is valid: Does either firm have anything to gain by misrepresenting its net surplus to the sponsor? In the case that $s_1 < s_2$, does firm 1 want to overstate s_1 ? Unless its lie causes firm 1 to win the bid instead of losing it, the lie has no effect, since firm 1's payment does not depend on its own bid s_1 . If the lie is large enough to change the outcome, then firm 1's profit is $s_1 - s_2$, which would be negative. Firm 1 is better off losing the bid and making zero profit.

In the case that $s_1 > s_2$, does firm 1 want to understate the surplus s_1 ? Unless its lie causes firm 1 to lose the bid instead of winning it, the lie again has no effect. If the lie is large enough to change the outcome, then firm 1 makes zero profit by losing the bid instead of earning $s_1 - s_2$, which is positive. Thus, firm 1 has no incentive to lie about s_1 . The same argument applies to firm 2.

Thus, the most important feature of the second-price auction is that each firm has an incentive to report faithfully on the net surplus it can deliver, and the sponsor can safely pick the firm that claims the highest surplus.

The second-price auction is particularly attractive if the surplus available from the two rivals is expected to be similar. In that case, the payment to the winning bidder will be close to the winning bidder's cost, and equal if $s_1 = s_2$. In general, however, a sponsor would care about the size of the transfer that he must pay to the winning bidder. The social cost of raising funds for general revenue is smaller than the social cost of taxing a single market, but is still not zero. The second-price auction yields an efficient outcome in the sense that the high-surplus firm is asked to invest, and there is no duplication of cost.

The second-price (Vickrey) auction assumes that the *ex post* payment can depend on the delivered value. But that is the difficulty we would like to avoid. The contest we now discuss, following Che and Gale (2003), is a hybrid between the commitment prize and the Vickrey auction.

A Prototype Contest The *ex ante* problem of the sponsor is to elicit investment in the best idea, given that he cannot condition his payments on the value of the delivered innovation. But in order to elicit investment, he needs a mechanism that commits him to pay off *ex post*, instead of trying to negotiate a low price once the prototypes are delivered. If the sponsor cannot commit to pay a price above production cost, the innovators will not invest.

The prototype contest solves this problem by allowing the firms and the sponsor to make contingent contracts before any investments are made. The contracts specify what price the sponsor will pay to each innovator, contingent only on buying that firm's innovation. The sponsor's observation of quality is reflected only in the decision to buy. The only enforcement problem is to make sure that *if* the inventor is chosen, the price will be as specified in the contract written before the innovator invested. Each firm would like to get a high price if its innovation is chosen, but a high price increases the chance of not being chosen *ex post*. This constrains the prices demanded by the firms in negotiating the contracts *ex ante*. On the other hand, since the contingent contracts are negotiated before costs are sunk, the firms will not offer prices so low that they do not cover costs.

To see this mechanism more explicitly, suppose the two potential innovators have ideas $(v_1, c_1), (v_2, c_2)$.⁹ Suppose, for concreteness, that $v_1 > v_2$. If v_1 is very close to v_2 , then in an *ex post* auction, the winning bid would be close to zero, and neither bidder would be likely to cover its cost in expectation. This is remedied in the contest we now describe by letting the firms bid *ex ante*, before they sink their costs. For simplicity, assume that the qualities v_1, v_2 are known to all parties, even if they

⁹The model presented here is simpler than that of Che and Gale. In their model, the firms can choose what qualities to develop, using a more standard model in which there is a known function that yields quality in return for investment effort. The firms announce their qualities in advance, coupled with a random distribution on prices as below.

cannot be verified in court. Thus, all the parties know what qualities of innovations will be delivered *ex post*. The firms will announce *ex ante* the prices (ρ_1, ρ_2) at which they are willing to sell their innovations *ex post*. The firm will receive this price, but only if chosen by the sponsor *ex post*.

Our objective is now to characterize the firms' equilibrium bids. These are bids such that neither firm has an incentive to revise its bid, assuming that the other firm's bid is fixed.

We first argue that the contingent bid prices cannot be deterministic. To see this, suppose that $v_1 = v_2$ and $c_1 = c_2 = c$, and consider what prices the firms will demand. A natural guess is $(\rho_1, \rho_2) = (2c, 2c)$. These are the minimum prices the firms could demand and still cover their costs in expectation, assuming that the tie-breaking rule is to randomize evenly between the firms. Since each firm would win the bid with probability $1/2$, the revenues would be $(1/2)\rho_1 = (1/2)\rho_2 = c$.

However, these prices are not an equilibrium. If firm 2 demands price $2c$ in the event that the sponsor chooses that firm, then firm 1 can improve profit by reducing its own demand to $\rho_1 = 2c - \varepsilon$, where ε is a small number. With prices $(\rho_1, \rho_2) = (2c - \varepsilon, 2c)$, the sponsor will choose firm 1, and firm 1 will make profit $\rho_1 - c = c - \varepsilon$ instead of 0. This shows that the zero-profit prices $(\rho_1, \rho_2) = (2c, 2c)$ are not an equilibrium. Of course there are no lower prices that are an equilibrium either, since at least one firm would then not cover cost in expectation.

The solution to this problem is mixed strategies. Instead of choosing a deterministic price, each firm chooses a probability distribution over prices, and the price actually offered to the sponsor is a random draw from this distribution.

An equilibrium will have the property that the firms randomize on whether they develop the innovation, as well as on the price. With some probability each firm will drop out, which means that it does not innovate and demands a zero price. If both firms drop out, the sponsor will fail in his objective of procuring the innovation by setting up a contest. We will denote the cumulative distributions on price by $F_1, F_2 : [0, \frac{v}{r}] \rightarrow [0, 1]$. For each $\rho \in [0, \frac{v}{r}]$, the probability that the firm chooses a price no larger than ρ is $F_i(\rho)$, for $i = 1, 2$. The firm will never choose a price larger

than (v/r) because the sponsor would never pay a price greater than the value of the innovation. The price $\rho = 0$ will imply that the firm does not innovate, and a positive price will imply that the firm does innovate. Thus $F_1(0), F_2(0)$ are the probabilities that the two firms do not innovate.

The firms' decisions whether to innovate and what price to demand will be made independently of each other, not observing each other's choices. Their choices are commitments in the sense that they cannot change the contract *ex post* if, for example, the other firm's price turns out to be lower.

For the example with symmetric ideas $(v_1, c_1) = (v_2, c_2) = (v, c)$, there is an equilibrium in which the firms drop out with probabilities $F_1(0) = F_2(0) = \frac{c}{(v/r)}$, and the distributions on prices are $F_1 = F_2 = F$, defined as follows:

$$F(\rho) = \begin{cases} \frac{c}{(v/r)} & \text{if } 0 \leq \rho < c \\ 1 - \frac{c}{\rho} + \frac{c}{(v/r)} & \text{if } c \leq \rho \leq \frac{v}{r} \end{cases}$$

This implies that probability $c/(v/r)$ the firms drops out and asks for price $\rho = 0$, that the firm never chooses any other price between 0 and c , and that the probability distribution has density function c/ρ^2 for prices between c and v/r . That is, conditional on innovating, the firm puts most of its probability weight on a price near the cost c , with the probability weight declining to the maximum (v/r) .

This is an equilibrium because each price in the support of the distribution yields the same expected profit as any other price, namely zero, and no other price would yield greater profit. If firm 1 drops out (chooses price $\rho = 0$), it makes zero profit. If firm 1 develops the innovation and demands any price ρ in the interval $[c, \frac{v}{r}]$, firm 1's expected profit is $\rho[F_2^0 + 1 - F_2(\rho)] - c = \rho \left[\frac{c}{\rho} \right] - c = 0$. The term $[F_2^0 + 1 - F_2(\rho)]$ represents the probability that firm 1 wins the bid *ex post*. With probability F_2^0 firm 2 drops out, and with probability $1 - F_2(\rho)$ firm 2 innovates but demands a higher price than firm 1's price.

These strategies hold the firms to the lowest possible expected profit that will induce them to invest, given that two of them are asked to innovate. However it is important to notice that, aside from the oddity of random prices, the investment decisions are inefficient. With some probability, the sponsor does not get the innovation, and even if he gets it, there is a large probability that the costs will be duplicated.

Allowing duplication is how the sponsor induces rivalry to keep the procurement price down.

The reader can work out how this mechanism must be modified if the firms' ideas $(v_1, c_1), (v_2, c_2)$ are different. In that case, the random prices, conditional on both firms innovating, may have the consequence that the sponsor does not always choose the highest-value innovation. Instead, the innovator chooses the innovation that generates the highest surplus, which will be the lower-value innovation if the innovator also demands a very low price.

The main example of prototype competition has occurred in military procurement. Prior to the 1970's, the U.S. Air Force usually acquired prototypes from a single vendor after a contest to choose the best written proposal, and then decided whether to accept it. That procedure led to widely criticized cost overruns in the 1960's. In the 1970's the Air Force moved to a system where two rival companies received contracts to build prototypes followed by a flight competition to demonstrate quality. This process led to the F-16 and F-18 fighter jets (Sweetman 1991, Sweetman et al, 1987).

While these contracts do not mirror the exact mechanism discussed above, they show that the sponsor can use prototype competition to keep quality high and costs low, at least relatively. Rogerson (1994) points out that the payoff to winning the prototype competition is a lucrative production contract. According to Rogerson, such procurement can be viewed as a three stage process consisting of (a) a design phase in which multiple firms pursue competing designs, (b) a selection phase in which a limited number of firms compete to produce prototypes and/or a final design, and (c) a production phase, typically involving sole source production by a single firm. He argues that the production phases allows the DoD to "award larger prizes to more important innovations, at least in a rough sense" as well as providing ongoing incentives to improve the produce after initial adoption. Firms that reach the production phase typically enjoy economic profits (above the normal return to capital) amounting to 4.4%. The super-normal rate of return can be viewed as a prize.

5 Government Grants

In this section we turn from the grand challenge prizes, such as those for longitude and food preservation, and focus on the more modest stuff of academic and industrial life, the small innovations that move knowledge forward in increments. Our premise here is that ideas are scarce, and that the agent with the idea should make an informed choice whether to invest in it. Under intellectual property incentives, or a prize system such as that in Lyon, the informed choice will follow the profit motive. But for most ordinary science, neither prizes nor patents have been the main mechanism for funding research. Instead, sponsors have employed scientists directly, or given them up-front grants in return for some unenforceable promise to implement a self-generated idea. How could that possibly work?

The grant process is a relative modern invention. For most of history, publicly sponsored research was in-house. This dates back at least to ancient Egypt, where the engineer Imhotep was hired to build bigger and better pyramids (DeCamp (1980)). European monarchs routinely promoted knowledge creation by attaching gifted individuals to their courts, for example, Archimedes, Kepler, Brahe, Euler, and Lagrange. Later examples of government employees who became famous innovators include Charles Darwin, science officer on the *Beagle*, and Herman Hollerith, the census bureau inventor of punch card technology (Porter 1994). In modern versions of intramural research, the emphasis has been on teams. During WWII, the U.S. government hired large groups of academics to produce weapon prototypes, e.g., radar at the MIT Radiation Laboratory, nuclear weapons at Los Alamos, torpedos at the Harvard Underwater Sound Laboratory, and rockets at the Jet Propulsion Laboratory. Additional advances were generated within the government's own laboratories, notably the invention of an advanced "laminar flow" wing that was used in the "Mustang" fighter plane of WWII (Sweetman 1984). This arrangement is continued by NASA and the national laboratories funded by the Department of Energy.

The paid researchers who are most familiar to readers of this article are university faculty members. This arrangement also has a long history. The preeminent example in the ancient world was Egypt's Library of Alexandria, which supported

domestic and foreign scholars who were free to work on their own ideas. Resident scholars included Archimedes, Hipparchos, Eratosthenes, Euclid, and Hero of Alexandria. These researchers made fundamental advances in astronomy and mathematics, as well as a variety of clever mechanical and clockwork devices. The mechanic Ktesibios, also employed there, was widely acclaimed as “The Edison of the Ancient World” (De Camp 1980; Casson 2001). Later, the famous inventors Galileo, Newton, Dalton, Volta, Ampere and Maxwell (Porter 1994) all worked as employees of universities, rather than in response to project-specific incentives.

The industrial laboratories that came into existence in the 20th century mimic the university model, especially with the establishment of campus-like settings for research in electronics and biotechnology. Probably the most famous of these is Xerox Parc in Palo Alto, which originated many of the computer technologies that can probably be found on the reader’s desktop.

Before WWII, R&D funding by the U.S. federal government was a small percentage of total R&D, and what existed was mostly intramural. The nature of public funding for R&D changed precipitously after the war.¹⁰ The federal government became a primary source of R&D funds, of which 70% is currently given out as extramural funding. The percentage of total R&D that is paid for by the federal government has fallen steadily between 1953 and the present, from about 2/3 to 1/3, but it is still much larger than the percentage in the 1930’s, which was between 12% and 20% (Mowery and Rosenberg 1998, p. 27).

To explain the virtues of the grant process, we shall retain our model of innovation in which agents have ideas for innovations, and must be given incentives to invest in those ideas. Even though the grant giver cannot monitor the grantee or withdraw payment in case of failure, and has limited ability to monitor the past record, the system can do a very good job of selecting the best researchers and making sure that they perform. The grantor’s main instrument of coercion is that he can cut off funding for future projects if a grantee fails to deliver on a past grant.

As we show, the system works by self selection. The funding agency evaluates

¹⁰For a concise account, see Mowery and Rosenberg (1998, Ch 2). For the basic data on R&D funding, see NSF, Science Indicators 2002.

ideas (proposals), but will only fund them, even if good, if the researcher has not failed to deliver in the past. In order to stay in the system, the researcher will have to be honest about his ideas for grants – else he won't succeed – and must actually bear the cost of implementing his idea. It is only the researchers with relatively fertile minds who will have an incentive to do this, and this supports the funding agency's objective. However, the granting agency cannot actually observe the fertility of the researcher's mind or keep track of how many innovations the researcher has made. We show that, nevertheless, the grant system works well to motivate effort and select the best researchers. (If even the best and most honest researchers can fail by mistake rather than malfeasance, the argument we give must be modified so that the grantor can also use information on the history of successes.)

Since the fertility of a researcher's mind can be captured either in the rate at which he has ideas or in the value of the ideas, we will assume for simplicity that all ideas have the same value and cost, (v, c) , but that researchers receive ideas at different rates $\lambda \geq 0$ per year, where λ represents the fertility of the researcher's mind. The objective of the grant agency is to reward the high- λ researchers.

When the researcher receives an idea, he can file a grant proposal with the sponsor, and the sponsor will decide whether or not to fund him. The researcher does not need to worry about getting paid, because he gets the money up front. However, whether or not the sponsor can observe the merits of the idea (v, c) , he has to worry that the researcher either cannot execute the idea or prefers to pocket the award and go to the beach.

In a given time period, a researcher with parameter λ has an average of λ ideas. (The parameter can be less than one, in which case the researcher will not, on average, have an idea in every period.) We can think of the parameter λ as measuring how creative he is. For a researcher with creativity λ , the expected present discounted value of investing in all the ideas he has in a given period at date t is the following, when r is the discount rate and the size of the grant per idea is ρ :

$$\frac{\lambda}{(1-r)^t}(\rho - c)$$

Suppose now that the researcher has received a grant of size ρ . He must decide whether

to perform the research or go to the beach. His net gain if he fails to perform is the saved cost c . His net loss from lost future grants is

$$\sum_{t=0}^{\infty} \frac{\lambda}{(1-r)^t} (\rho - c) = \lambda(\rho - c) \sum_{t=0}^{\infty} \frac{1}{(1-r)^t} = \frac{\lambda}{r} (\rho - c)$$

Thus, he will perform instead of pocketing the money if

$$c \leq \frac{\lambda}{r} (\rho - c) \tag{1}$$

or

$$c \leq \frac{\lambda}{\lambda + r} \rho$$

We can see that, for fixed awards ρ , only researchers who expect to have lots of ideas (high λ) will perform in return for future options on grants. If the inequality (1) holds for any λ , it also holds for any researcher with a higher value of λ . We can also see that, for a fixed rate of idea formation λ , researchers will only perform if the rewards ρ are high enough. And, of course, they will never perform if the award ρ is smaller than the cost c .

Let H be the distribution of λ in the population of researchers, so that $1 - H(\lambda)$ is the fraction of researchers with parameter greater than λ . For a fixed award size ρ , let $\hat{\lambda}(\rho)$ be the minimum λ for which (1) holds, that is, the value for which (1) holds with equality. The function $\hat{\lambda}$ is decreasing in ρ : For higher prospective awards, even less creative researchers are willing to perform in return for future options on grants. Then the number of funded researchers $1 - H(\hat{\lambda}(\rho))$ increases with the size of the award, ρ . The total budget of the sponsor per unit time is

$$B(\rho) = \int_{\hat{\lambda}(\rho)}^{\infty} \lambda \rho dH(\lambda)$$

In this system, if the sponsor wants to ensure that only the more creative researchers are funded, he must give awards without rationing. The only researchers who continue to apply for grants are those with creativity parameters larger than $\hat{\lambda}(\rho)$, and all of them are funded. The only way to increase the amount of research (number of researchers) is to increase ρ for all researchers, which means that the budget increases by more than the payments added for new researchers.

It is worth understanding what is lost by the grant giver if he cannot condition his awards on past success. Conditioning the size of awards on past success will reduce

the grant giver’s total budget, conditional on a fixed amount of innovation. Because the sponsor cannot observe λ , or condition grants on λ , sponsorship is more costly than it would be otherwise.

If the sponsor could observe the fertility λ of the researcher’s mind, he could make different payments for different researchers; the grant per idea, ρ , would be a function of λ . For each λ , he would choose a prize $\rho(\lambda)$ for which (1) holds as an equality. The function ρ would then be a decreasing function of λ ; more productive researchers would receive less money per funded idea. However, one can also see from (1) that, since $c(\lambda + r) = \rho(\lambda)\lambda$, the more creative researchers (those with higher λ) would receive more grant funding per unit time, $\lambda\rho(\lambda)$.

The fertility of each researcher’s mind will depend, among other things, on whether he is operating in an environment of “open science” (David 2003). If open science means that each discovery is published, hence shared, each researcher will be stimulated to have ideas at a higher rate. We can thus interpret open science to imply that each researcher’s λ is higher than otherwise. With higher λ ’s, a given total grant budget will support a higher rate of progress.

6 Public Sponsors and the Direction of Research

For some research objectives, there may be several paths to the result, each one risky in its own idiosyncratic way. That is, one of the things that is unknown is whether a particular approach will succeed. In this environment, an idea is a triple (v, c, p) , where v is the value of the specified objective, c is the expected cost of a particular research approach, and p is the probability that the approach will fail.

It is natural to assume that the successes and failures of different approaches are independent. Thus, if there are several ideas $\{(v_i, c_i, p_i)\}_{i \in S}$ for how to achieve the objective, and the sponsor elicits investment in all of them, the probability of success is $1 - \prod_{i \in S} p_i$. (This model follows Wright (1983), except that in Wright’s model all competitors have the same idea. Only the sponsor is uninformed.)

In this environment, the main efficiency question is how many approaches will be taken, and which ones. One of the difficulties of patent incentives is that they can

yield either too much or too little entry (too many or too few approaches). To see this, consider the case that $v_i = v$ for each i (the potential knowledge is the same for each firm), $c_i = c$ for each i (the approaches are equally costly) and $p_i = p$ for each firm (they have the same likelihood of failure), but the approaches are nevertheless different, reflected in the fact that the probabilities of failure are independent. Suppose that the value of the property right will be πvT , as above. If n approaches are taken, the probability of success is $1 - \prod_i p_i = 1 - p^n$. Then the total expected revenue of the firms collectively is given by a function f defined as

$$f(n) = (1 - p^n) \pi vT$$

The social value of the innovation can be written $kf(n)$ for some number $k > 1$. This is because the social value is $\left[(m + \pi)vT + \left(\frac{1}{r} - T \right) v \right] (1 - p^n) > (1 - p^n) \pi vT$, where mv is the per-period consumers' surplus that is available with proprietary pricing.

The number of approaches that maximizes social welfare, namely, the n^* that maximizes $kf(n) - cn$, will not be the same as the number of approaches \hat{n} that are taken by competitive firms in a patent race. Firms will enter will different approaches up to the point where $\frac{f(\hat{n})}{\hat{n}} = c$; that is, they dissipate the rent. To see that the optimal number n^* can be either larger or smaller than the equilibrium of the race, \hat{n} , notice first that, since f is a concave function originating at 0, it holds for all n that $f'(n) < \frac{f(n)}{n}$. The optimal n^* satisfies $kf'(n^*) = c$, which implies that $\frac{kf(n^*)}{n^*} > c$. If k is close to 1, then, since $\frac{f(n)}{n}$ is decreasing, $\hat{n} > n^*$. But as k becomes large, n^* must also become large, with the consequence that it may hold that $n^* > \hat{n}$.

It is difficult to remedy the incorrect patent incentives. If too few firms would enter, it is hard to entice more, since the additional entrants would make negative profit in expectation. If too many firms would enter, it would be hard to exclude them, especially if the ideas were common knowledge. If some competitor dropped out, another competitor would enter with the unused approach.

Thus, if the approaches to solving the problem are common knowledge, the public sector can do a better job of coordinating the different approaches, by adjusting their relative funding or funding the appropriate selection from them.

The advantages of public sponsorship become even more pronounced if we augment the model to account for learning. We will not give a formal treatment, but will summarize the idea. An R&D effort typically takes place over time. As portions of the cost c are invested, more is learned about the probability that each approach will be successful. If one approach seems particularly promising, then it is rational to shift resources away from the less promising toward the more promising. At some point the promising venture may play itself out, and then the resources should again be reorganized.

This rational shifting of research priorities can best take place if the intermediate knowledge learned is aggregated in a systematic way. A problem with private IP incentives is that firms have no incentive to share knowledge in a way that is socially efficient. Their incentives to share knowledge are polluted by a desire to win the race, and hence to get rivals to drop out, even if that is inefficient. The rivals know that each of them has such an incentive, and therefore the firms' signalling about their private states of knowledge will be suspect.¹¹

An example of different ideas embodying different (uncertain) approaches can be found in the attempt to develop nuclear fusion as an energy source. Since 1945, governments have understood that current fossil-fuel and nuclear-energy technologies cannot meet society's demand for electricity indefinitely. By contrast, a successful fusion technology would generate essentially limitless energy from an isotope, deuterium, found in sea water. Researchers in the US, Europe, Russia, and Japan have all been interested in this problem, using different approaches, at considerable expense. Over the past 25 years, the U.S. expenditure on fusion-related R&D has consistently ranged (in constant dollars) from \$100 million to \$500 million per year (G Kulcinski & J. Santaristu 1998). The current budget stands at \$257 million (Seife 2003).

There are at least twenty approaches to the problem of generating energy through nuclear fusion, depending on how one counts.¹² Almost all of these approaches

¹¹See Scotchmer and Green (1990) for a model that is similar in that firms may induce shakeouts by revealing their intermediate progress in pursuit of a goal.

¹²Approaches are conventionally divided between "magnetic confinement schemes" that use powerful magnets and/or electric fields to contain ultra-hot gas and "inertial confinement" schemes that extract energy from fuel pellets before they explode. Technologically distinct methods for magnetic

were understood at the very beginning, in 1951-52 (Herman, 1990, Lindl 1998). Despite years of public funding and experimentation, none of the approaches has been proved to achieve goals – and none has been definitively ruled out by experiment. The technical approach favored by the Russians – so-called tokamaks – leaped ahead in trials for a brief period in the 1960's, and worldwide funding shifted toward that approach. By 1972, at least seventeen new tokamaks were completed or under construction outside the USSR (Herman 1990). Despite the disappointing record, partisans in the global partnership remain confident, and in 1986 they hatched a plan to spend \$10 billion commercial-scale fusion reactor. The plan has been controversial due to the unproven nature of the technology. The commitment of funds was cut in half after the U.S. briefly dropped out in 1998 when detractors were advocating other approaches and more basic science (Seife 2003).

It is unclear whether private ventures would ever invest in an innovation as costly as nuclear fusion, regardless of the potential payoff, and that may be the single most compelling reason that it should be funded publicly, if at all. But the other important lesson embedded in the example is that researchers can and do learn from each other. The tokamak may or may not end up the survivor, but it is significant that researchers responded to experimental evidence in its favor. Private competitors might have suppressed such evidence in order to maintain a competitive advantage.

7 Mixed Private/Public Incentives

In contrast to many of the examples given above, and also in contrast to how the public/private question is usually posed by academic economists, incentives are generally not provided entirely through intellectual property or entirely through public sponsorship. The modern economy has evolved such that public and private funds

confinement include pinch effect machines, stellarators, tokamaks, spherical tokamaks, spherical torus machines, spheromaks, electric tokamaks, colliding beam machines, magnetic mirrors, floating multipole machines, reversed field pinch machines, magnetic targets, inertial electrostatic confinement machines, and Z pinch machines. Inertial confinement schemes include direct drive by lasers, direct drive by ion beams, and indirect drive by lasers. The foregoing categories are not exclusive. Dark horse candidates have included cold fusion, sonoluminescence, and dense plasma focus machines. For a detailed survey, see National Fusion Energy Science Web Site (1998).

are often mixed.

There are at least three pieces of legislation that underlie that phenomenon. The most important are probably the Bayh-Dole and Stevenson-Wydler Acts of 1980, which authorize the patenting of innovations developed with federal funds. They affect the national laboratories operated by the Department of Energy, such as Lawrence Berkeley Laboratory, Livermore Laboratory, Los Alamos, and many others, which collectively account for about 3% of total R&D spending in the U.S. They also affect most research done in universities, which collectively account for about 11% of total R&D spending in the U.S, but receive about 60% of their research funds from the federal government.¹³ Subsequent to the Bayh-Dole Act, most research universities have instituted Offices of Technology Licensing. Most such offices prefer exclusive licensing (Jensen and Thursby 2001), suggesting that they intend to exercise their rights as profit maximizers, just as a private entity would.

The other legal foundation is the Federal Technology Transfer Act of 1986. This is a tool that lets national laboratories (and other federal enterprises) create partnerships with private entities, called CRADAs, and thus spin off their innovations into the private sector. By the mid-1990's, about 1000 CRADA's were being formed each year, almost half in partnership with the Department of Energy.¹⁴

Finally, there are occasional examples of public/private incentives, resulting in intellectual property rights for the private sector, that fall outside these two broad categories. DARPA has offered a \$1 million Grand Challenge prize for winning a race among robotic ground vehicles (DARPA 2003). Entrants not only compete for this rather substantial prize, but will own any intellectual property that results.

Many commentators question the rationale for funding R&D directly with public funds, but at the same time offering intellectual property rights. The fear is that inventors get funded twice - once with grants and once with IP - thus compounding the burden to users and taxpayers. In addition, such lucrative funding can lead to wasteful duplication of efforts.

Many of these public/private ventures involve "big science", where one might

¹³NSF, Science Indicators 2002.

¹⁴See Maurer (2002) for a more complete discussion.

argue that neither the private sector nor public sector alone can “afford” to fund it. That argument is not very convincing, since large benefits should attract large investments, even if the funds have to be borrowed. In any case, no organization has better ability to raise money than the federal government.

Nevertheless, the public/private partnership resulting in intellectual property rights may have merit. Our argument is not focussed on the difficulty of raising money, but rather on how the system of dual funding can overcome other problems that would otherwise afflict both sponsors. For big science, industry has the problem that, although an innovation may have some commercial value, the commercial value under existing intellectual property laws may not be sufficient to cover costs. This is especially true for investments that have unappropriable social benefits, or where intellectual property rights are so narrow that the benefits will be eroded by competitors. The public sector could simply fund such projects, but then it faces the problem of choosing the ones that are likely to be fruitful, or making sure the funds are used as intended, especially when the expertise resides mostly in the private sector.

We have already shown how a grantor can overcome these problems for small, frequent innovations, simply by threatening to cut the innovator out of the grant process if he or she does not deliver enough innovations, or innovations of high enough value. But “big science” will not produce frequent deliveries of small innovations, regardless of how meritorious it is. In fact, for many big science projects such as energy through nuclear fusion, the real quality indicator is the likelihood of success, and even the best project may fail. The public sponsor therefore needs some mechanism to screen for the right investments, and to make sure that the researchers invest as directed.

We suggest that a system of matching funds, where it is industry matching the government subsidy rather than vice versa, together with partial payment to industry through intellectual property rights, can solve the dual problems of ensuring that industry covers its cost, and avoiding wrong investments by the public sector.

Each point in Figure 1 represents an idea (v, c) . We will suppose that industry is the repository of the best information about investments, so that the value and cost




Figure 1: Hybrid Incentives: Public Subsidies with Private Matching Funds

(v, c) is known to the researcher, but not to a government sponsor. The objective of the government sponsor is to invest in those ideas for which $\frac{1}{r}v - c > 0$. The value of the intellectual property available to a private firm if it invests on its own is πvT , and without a government subsidy, it will invest if $\pi vT - c > 0$. In Figure 1, the area under the lower diagonal line, labeled πvT , represents the ideas for which the private incentives in intellectual property would be sufficient.

Suppose, however, that there are high-value ideas that are more costly than the value of the intellectual property right (above the line πvT), but still worth doing ($\frac{1}{r}v > c > \pi vT$). The government might like to sponsor these.

Suppose that the government simply offers a subsidy s . Then all the ideas (v, c) under the higher diagonal line $\pi vT + s$ will be undertaken. The subsidy will increase research, but indiscriminately. In particular, there are likely to many low- v ideas, toward the origin in Figure 1, for which the subsidy of s is a waste of money. It is easy to imagine an endless series of subsidy claims for worthless innovations.

The government can solve this problem by insisting that the claimant make a matching commitment of funds in some amount, say m . Then in order to claim the

subsidy, the sponsor and the claimant will contribute s and m respectively to the research budget, and invest in ideas suggested by the claimant. If the claimant suggests an idea (v, c) such that $c > s + m$, then he must provide the required supplement. In that case, the industry partner pays $c - s$ rather than m . If the claimant suggests an idea (v, c) such that $c < s + m$, then the surplus goes to supporting graduate students or other research enterprises. For simplicity, we assume in Figure 1 that the claimant receives intellectual property in amount $\pi T v$ on the subsidized innovations, although in practice, the sponsor may require a reduction in intellectual property rights for some broader social purpose. For example, the NIH imposes guidelines under which research tools must be made freely available to other academic researchers, with an intent to protect other grantees of the NIH.

We can now see what happens under this incentive mechanism.

First, ideas (v, c) that satisfy $c < m$ will not be subsidized. The industry partner gets the intellectual property rights whether or not he is subsidized, and it is more profitable to get these property rights for the lower cost c rather than m . Thus, the dark shaded area that is below both the diagonal line $\pi v T$ and below the horizontal line m represents ideas that industry will invest in without claiming any subsidy.

It is only in ideas (v, c) that satisfy $\pi v T > m$ that the industry partner would claim a subsidy. These ideas are all to the right of \underline{v} in figure 1. This is what solves the “moral hazard” problem. Partners will not try to collect subsidies on worthless innovations because that would obligate them to commit funds in amount m .

We can further refine the ideas for which the industry partner would be willing to accept a subsidy, namely, the higher shaded area in Figure 1. These are the ideas to the right of \underline{v} , above the horizontal line m , and below the line $\pi v T + s$. Ideas such that $c > \pi v T + s$ will not be attractive to an industry partner because the partner must pay the surfeit $c - s$.

We thus see that the partnership with mandatory matching funds will allow the public sector to subsidize ideas that would otherwise not elicit investment (those between the lines $\pi v T + s$ and $\pi v T$, to the right of \underline{v}), without causing the sponsor to

hemorrhage money in subsidies to worthless innovations.

In recent years, industry matching requirements have become increasingly important in government aircraft procurement. During the 1980's, industry contributed roughly one-half the cost of developing prototypes for the F-22 and F-23 fighters (Sweetman 1991). During the 1990's, NASA agreed to spend \$912 million on the X-33 spacecraft, provided that Lockheed contributed \$212 million. Lockheed, which had agreed to pay for any cost overruns, eventually spent \$357 million before the project was cancelled in 2001 (NASA 2001a,b).

The rationale contained in the Bayh-Dole Act is not the one given here. The rationale is not about *ex ante* incentives to create innovations, but focusses instead on the *ex post* incentive of universities and national laboratories to disseminate them. The basic assertion is that licenses on university patent rights, especially exclusive licenses, give industry licensees enough protection to make the collateral investments required to commercialize those innovations.

But there is an equally compelling argument that publicly sponsored research should be put in the public domain. This would arguably maximize the number of users, and therefore maximize the number of follow-on innovations. Whether the *ex post* rationale for the Bayh-Dole Act is convincing depends on whether one believes that intellectual property rights are constructed sensibly to begin with. If the intellectual property for follow-on innovations is sufficient, then exclusive licenses on the underlying publicly sponsored inventions are not necessary. And if insufficient, then the Bayh-Dole Act remedies a defect of intellectual property law. Wouldn't it be better to fix the thing that's broken?

8 Conclusion

Public procurement provides the option to pay for knowledge out of general revenue, which is generally thought to impose less deadweight loss than taxing a single market, and then to put the knowledge in the public domain. Thus, public sponsorship of knowledge that is destined for the public domain has a natural advantage over intellectual property: it avoids the inefficient exclusion of users whose willingness to

pay is smaller than the proprietary price, but larger than the marginal cost of supply. For such exclusion, we use the term “deadweight loss”.

However, intellectual property also has a natural advantage over public sponsorship: it ensures that only the users pay the costs of inventions, and since these payments are voluntary, no citizen can be made worse off by an invention. The latter advantage is particularly persuasive for knowledge that is narrowly targeted to a special audience, but it is not determinative even there. Users of knowledge in two disjoint targeted audiences could all be better off by making reciprocal subsidies, each contributing their tax dollars toward the knowledge that is useful to the other targeted audience, under the condition that no one is excluded from use.

On grounds that citizens can provide reciprocal subsidies, we have implicitly taken the view in this paper that it is best to avoid deadweight loss (proprietary pricing). This pre-judges the choice between public and private sponsorship, and in doing so, allows us to focus on other important aspects of the creative environment, in particular, which schemes are best at harnessing the creative genius that is widely dispersed among humans.

It is worth noting that sometimes there is no alternative to public sponsorship funded through taxation, since it may be impossible to exclude beneficiaries from using new knowledge, and therefore impossible to charge each user a price.

Conversely, invention can sometimes create its own reward without any formal incentive mechanism. In the early days of airplane development, manufacturers often subsidized the aircraft that were used in competitions for long-distance flights, since winning could be expected to create commercial advantages.¹⁵ Benefits presumably included publicity and gains to reputation, the opportunity to develop trade secrets and tacit knowledge, and the firm’s normal share of externalities to aviation as a whole. In the case of air racing, prizes were also supported by gate receipts (Vorderman 1992). The X-Prize competitors have similar business plans. Contestants hope to tap a small but lucrative market for passenger flights into space (Hoffman 2003).

¹⁵In one striking example, Bellanca told the Spirit of St. Louis team that it was willing to cut the price of its aircraft from \$25,000 to \$15,000 – but only if they let the company name its own pilot. Lindbergh declined. [Jablonski (1972).]

But in addition, many racing aircraft received massive subsidies from governments and, in the 1930's, also from charitable or patriotic individuals (Vorderman 1992, Sweetman 1984). The pattern again continues in the X-Prize competition, where at least two competitors are reportedly bankrolled by dot.com millionaires. Most of the X-Prize itself was raised from such donors (Hoffman 2003). Motivations include civic pride, patriotism, and an ideological commitment to spaceflight as an end in itself (X-Prize Foundation 2003.)

References

- [1] Bryant, J. H., 1988. *Henrich Hertz: The Beginning of Microwaves*. Institute of Electrical and Electronics Engineers, Inc., Piscataway, NJ.
- [2] Casson, Lionel, *Libraries in the Ancient World*. Yale University Press, New Haven, CT.
- [3] Conant, J., 2003. *Tuxedo Park: A Wall Street Tycoon and the Secret Palace of Science That Changed The Course of World War II*. Simon & Schuster, New York.
- [4] David, P., 2003. The Economic Logic of “Open Science” and the Balance between Private Property Rights and the Public Domain in Scientific Data and Information: A Primer. In National Research Council, *The Role of the Public Domain in Scientific and Technical Data and Information*, National Academies Press, Washington, D.C.
- [5] Defense Advanced Research Projects Agency, DARPA Grand Challenge, <http://www.darpa.mil/grandchallenge/>.
- [6] Foray, D. and Hilaire-Perez, L., 2000. *The Economics of Open Technology: Collective Organization and Individual Claims in the “Fabrique Lyonnaise” during the Old Regime*. mimeograph.
- [7] Gallini, N. and Scotchmer, S., 2001. Intellectual Property: When is it the Best Incentive Mechanism? , *Innovation Policy and the Economy*, Vol. 2. MIT Press, Cambridge, MA, pp. 51-78.
- [8] Google, Google Programming Contest. <http://www.google.com/programming-contest/>.
- [9] Herman, R., 1990. *Fusion: The Search for Endless Energy*. Cambridge University Press, New York.
- [10] Hoffman, C., 2003. The Right Stuff. *Wired*, July 2003, 134-145.
- [11] Holden, C., 2003. DARPA’s Desert Dash. *Science* 300, 733.

- [12] Innocentive, 2003. About Innocentive. <http://www.innocentive.com/about/>.
- [13] Jensen, R. and Thursby, M., 2001. Proofs and Prototypes for Sale: The Licensing of University Inventions, *American Economic Review* 91, 240-259.
- [14] Jablonski, E., 1972. *Atlantic Fever*. Macmillan, New York.
- [15] Kulcinski, G. and Santarisu, J., 1998. Reducing the Barriers to Fusion Electric Power. *Journal of Fusion Energy* 17, 17- 23.
- [16] Kremer, M., 1998. Patent Buyouts: A mechanism for encouraging innovation, *Quarterly Journal of Economics* CXIII, 1137-1168.
- [17] Lindl, J. D., 1998. *Inertial Confinement Fusion*. American Institute of Physics Press, New York.
- [18] MacArthur Foundation, The MacArthur Fellows Program, http://www.macfound.org/programs/fel/fel_overview.htm.
- [19] Maurer, S. M., 2002. Promoting and Disseminating Knowledge: The Public/Private Interface. National Academies of Sciences, http://www7.nationalacademies.org/biso/Maurer_background_paper.html.
- [20] Mokyr, J., (1990), *The Lever of Riches*, Oxford University Press, Oxford.
- [21] Mokyr, J., 2002. *The Gifts of Athena*. Princeton University Press, Princeton, NJ.
- [22] Mowery, D. and Rosenberg, N., 1998. *Paths of Innovation*. Cambridge University Press, New York.
- [23] National Fusion Science Energy Web Site, 1998. *Fusion Innovations: A Report from ICC98*. http://www.fusionscience.org/ICC98/ICC98_Sum.pdf.
- [24] National Science Foundation, 2002. *Science Indicators 2000*. www.nsf.gov/sbe/srs/seind02/pdf_v2.htm.
- [25] Newhouse, E. L. ed., 1988. *Inventors and Discoverers: Changing Our World*. National Geographic Society, Washington, D.C.
- [26] NASA, March 1 2001. Press release. <ftp://ftp.hq.nasa.gov/pub/pao/pressrel/2001/01-031.txt>.

- [27] NASA, 2001. Ames X-33 Public Web Site. <http://asm.arc.nasa.gov/x33/x33home.shtml>.
- [28] Nelson, R. R. and Winter, S.D., 1982. An Evolutionary Theory of Economic Change. Belknap Press of Harvard University Press, Cambridge, MA.
- [29] O'Donoghue, T., Scotchmer, S. and Thisse, J.-F., 1998. Patent Breadth, Patent Life and the Pace of Technological Progress. *Journal of Economics and Management Strategy* 7, 1-32.
- [30] Penn, C., 1993. Super Efficient Refrigerator Finalists. *Home Energy Magazine*. <http://hem.dis.anl.gov/eehem/93/930305.html>.
- [31] Perez, L.H. and Foray, D., (2001). Inventing in a World of Guilds: The Case of Lyon in the XVIIIth Century. mimeograph, Université Paris-Dauphine. <http://emlab.berkeley.edu/users/bhhall/ipconf/forayperez01.pdf>.
- [32] Porter, R., 1994. *The Biographical Dictionary of Scientists*. Oxford University Press, New York.
- [33] Reichhardt, T., 2002. Out of This World. *Nature* 420, 10-11.
- [34] Reinganum, J., 1989. The Timing of Innovation: Research, Development and Diffusion, in: Schmalensee, R. and Willig, R.D. (Eds.), *Handbook of Industrial Organization*. Elsevier, Amsterdam, pp. 849-908.
- [35] Rogerson, W., 1994. Economic Incentives and the Defense Procurement Process. *Journal of Economic Perspectives* 8, 65-90.
- [36] RSA Security, 2003. Cryptographic Challenges. <http://www.rsasecurity.com/rsalabs/challenges/index.html>.
- [37] Scherer, F. M., 1984. *Innovation and Growth*. MIT Press, Cambridge, MA.
- [38] Scotchmer, S., 2004. *Innovation and Incentives*. forthcoming. MIT Press. Cambridge, MA.
- [39] Scotchmer, S., 1999. On the Optimality of the Patent System. *The RAND Journal of Economics* 30, 181-196.

- [40] Scotchmer, S. and Green, J., 1990. Novelty and Disclosure in Patent Law. *The RAND Journal of Economics* 21, 131-146.
- [41] Seife, C., 2003. United States Rejoins International Fusion-Research Project. *Science* 299, 801.
- [42] SpaceProjects.com, 2003. Competitive Prices Regarding Space as a Waste-Reducing Means of NASA Procurement. <http://www.spaceprojects.com/prizes/>.
- [43] Shavell, S. and van Ypersele, T., (2001). Rewards versus Intellectual Property Rights. *44 Journal of Law and Economics* 525
- [44] Strandh, O., 1979. *A History of the Machine*. A&W Publishers, New York.
- [45] Sweetman, W., 1984. *High Speed Flight*. Janes, London, UK.
- [46] Sweetman, W., 1991. *YF-22 and YF-23 Advanced Tactical Fighters: Stealth, Speed and Agility for Air Superiority*. Motorbooks, Osceola WI.
- [47] Sweetman, W., Gething, M., Richardson, D., Spick, M. and Gunston, W., 1987. *The Great Book of Modern Warplanes*. Salamander, New York..
- [48] Vickrey, W., 1961. Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance* 16, 8-37.
- [49] Wright, B., 1983. The Economics of Invention Incentives: Patents, Prizes and Research Contracts. *American Economic Review* 73, 691-707.
- [50] Vorderman, D., 1993. *The Great Air Races*. Bantam Books, New York.
- [51] X Prize Foundation Web Site, 2003. <http://www.xprize.org/>.
- [52] Zuckerman, S., 2003. The New Wave of Energy Efficient Refrigerators. <http://www.ecomall.com/greenshopping/icebox2.htm>.