Economics Working Paper Series

2015/011

Push or Pull? Grants, Prizes and Information

David Rietzke

The Department of Economics
Lancaster University Management School
Lancaster LA1 4YX
UK

Push or Pull? Grants, Prizes and Information *

David Rietzke†

May 18, 2015

Abstract

In the funding of R&D, push mechanisms, such as research grants, subsidize research input, while pull mechanisms, such as innovation prizes, reward research output. By rewarding research output, pull mechanisms create strong incentives for researchers to devote non observable inputs to R&D. Push mechanisms, in contrast, may reward a researcher independently of her output. In the presence of moral hazard, it might seem that push mechanisms generate weak incentives for non observable inputs from the researcher and, absent risk-sharing considerations, would be inferior to pull mechanisms; it is the aim of this paper to critically assess this hypothesis. I analyze a principal-agent model in which a funder encourages R&D activity through a push incentive (a grant) and/or a pull incentive (a prize); R&D input consists of both an observable and non observable component. In contrast to the stated hypothesis, it is shown that a grant may emerge as an optimal means of funding as a result of the interaction between adverse selection and moral hazard. The model also helps to explain the use of matching grants, it is shown that such grants serve as an effective sorting device in the presence of adverse selection.

KEYWORDS: Grants, Prizes, Moral Hazard, Adverse Selection, Innovation, Principal-Agent Problem
JEL Classifications: D82, D86, O31

*I thank Stan Reynolds, for his valuable input throughout this project. I also thank Asaf Plan, John Wooders, Andreas Blume, Martin Dufwenberg, Rabah Amir, Derek Lemoine, Brian Roberson,Charles Moul, and Tim Flannery for helpful comments and suggestions.
†Department of Economics, Lancaster University, d.rietzke@lancaster.ac.uk

1 Introduction

Innovation drives long-run economic growth, but in many instances the social value of an innovation exceeds the private value to the innovator. As a result, investment in R&D may fall well below the socially optimal level.[1] To encourage R&D spending, research funders have a number of tools at their disposal. Broadly, these tools may be categorized as either push or pull mechanisms. Push mechanisms, such as research grants, tax credits on R&D, or direct inputs from a funder, operate upstream and subsidize research inputs. Pull mechanisms, such as innovation prizes, tax credits on sales, or patent buyouts, operate downstream and reward successful research output. Research grants and innovation prizes are common examples of push and pull mechanisms, respectively, and are used by both public and private agencies to encourage R&D activity. The Bill and Melinda Gates Foundation, for example, issued more than $2.6 billion in grants in 2012,[2] much of which was intended to promote the development of new pharmaceuticals. Innovation prizes are a key component of the Obama Administration's attempt to stimulate American innovation.[3] From 2010 to 2012, 200 new innovation prizes were offered by federal agencies in areas ranging from national defense to education.[4]

One inherent challenge faced by research funders is that some relevant inputs may not be verifiable, and therefore, not contractible (moral hazard). By rewarding research output, pull mechanisms create the natural incentive for a researcher to devote non-contractible inputs to R&D, and are well suited to deal with this issue.[5] Push mechanisms, on the other hand, may reward a researcher independently of her success or failure in a research project. In the presence of moral hazard, it might seem that push mechanisms generate weak incentives for non observable inputs from the researcher and, absent risk-sharing considerations, would be inferior to pull mechanisms; it is the aim of this paper to critically assess this hypothesis.

I analyze a principal-agent model in which a researcher (the agent) of uncertain ability expends costly research inputs to increase the likelihood that she develops a new technology; success or failure in the project is verifiable. All else equal, a researcher of higher ability is more likely to innovate; the researcher knows her own ability (type), while the funder

[1] Jones and Williams (1998) estimate that the socially optimal level of R&D spending is 2 to 4 times greater than observed investment. For a recent survey of the empirical literature on returns to R&D see Hall et al. (2009)

[2] http://www.gatesfoundation.org/Who-We-Are/General-Information/Financials

[3] http://challenge.gov/about

[4] http://www.whitehouse.gov/blog/2012/09/05/challengegov-two-years-and-200-prizes-later

[5] That compensation be tied to observable outcomes in settings with moral hazard dates back to the work of Stiglitz (1974) Jensen and Meckling (1976), Myers (1977), Innes (1990), among others. Hölmstrom (1979) made the observation that any observable signal, which provides information about the level of unobserved effort could improve the payoff to the funder

knows only the distribution over types. Success generates a small profit to the researcher, but this incentive is insufficient to warrant R&D activity. The funder values the innovation more than the researcher, and encourages R&D through a push and/or pull incentive. The push incentive, which I call a grant, is a reward received by the researcher, independently of whether she successfully innovates. The pull incentive, which I call a prize, is any reward received by the researcher only if she successfully innovates. In order to focus solely on the incentive properties of push and pull mechanisms, I abstract from risk-sharing considerations and assume the researcher and the funder are risk neutral.

In the canonical model of moral hazard, (i) output is observable, but (ii) input is unobservable.[6] My model follows in this tradition with respect to (i), but differs critically from (ii). In the R&D process, some relevant inputs may be readily verifiable by the funder. The purchase of large-scale capital such as a new research laboratory, or equipment necessary for R&D, could be agreed upon ex ante, and verified ex post. Still, it may be difficult or impossible to monitor and confirm the time, effort, energy, etc. that the researcher devotes to the project. In this spirit, I depart from (ii) and consider an environment with two research inputs. One input is verifiable (I refer to the observable input as "investment"), while the other input is not (I refer to the non-observable input as "effort"). The two inputs are strategic complements, so that greater investments increase the marginal returns to effort (and vice versa). The model admits both a moral hazard problem - since some relevant research inputs are not observable, as well as an adverse selection problem - since the researcher is privately informed of her own ability.

It is shown that a grant may emerge as the optimal means of funding as a result of the interaction between adverse selection and moral hazard. To convey the intuition, first consider a pure adverse selection environment in which all relevant research input is contractible, but the researcher's ability is unknown to the funder. In this setting, the optimal means of funding uses only a grant. The reason is that the information rents captured by the researcher are proportional to the reward for innovation. Prizes, therefore, generate more rent for the researcher, and are a more costly means of funding from the funder's perspective. In a pure moral hazard environment (i.e. some relevant research inputs are non observable, but the researcher's ability is known), the optimal funding mechanism, in general, uses only a prize. This result follows the aforementioned rationale for pull mechanisms under moral hazard; the prize creates a strong incentive for the researcher to exert effort since it is only received if the researcher successfully innovates.

It is worth pointing out, however, that a grant can act as an incentive to encourage effort in this pure moral hazard environment. The grant may elicit greater investment from the

[6] See, for example, Bolton and Dewatripont (2005) (Ch. 4)

researcher, which increases the marginal returns to effort. As there is some profit incentive associated with innovation, this leads the researcher to exert greater effort, provided the moral hazard problem is not too severe.[7] Indeed, it is shown that if the researcher's effort supply is limited, and the moral hazard problem is not too severe, the optimal means of funding can involve either a pure grant, a pure prize, or a combination of the two. The channel through which the grant encourages effort depends critically on the presence of both contractible and non contractible inputs. There do not appear to be many studies that have considered such an environment, but one exception is Choi (1992). In the context of a research joint venture, Choi shows shows that inefficiency, resulting from under-provision of a non-contractible input, can be alleviated if there is another complementary and contractible input.

When adverse selection and moral hazard interact, the optimal means of funding depends on the relative strength of the two informational asymmetries. The funder faces a familiar trade-off between providing strong incentives for effort, and limiting the information rents that accrue to the researcher. Generally, a prize is more efficient in dealing with the former, while the grant is more efficient in dealing with the latter. If the moral hazard problem is not too severe, a pure grant emerges as the optimal means of funding. Interestingly, the combination of adverse selection and moral hazard may render prizes as a *less* attractive means of funding as compared to a setting with pure adverse selection. In settings where the moral hazard problem is relatively strong, a grant still may emerge as part of an optimal funding mechanism, used in conjunction with a prize.

The effectiveness of the grant in my model under moral hazard contrasts similar contracting models with only a single, unobserved input.[8] Lewis and Sappington (2000a, 2000b and 2001), for example, study a similar principal-agent problem under risk neutrality, allowing for both adverse selection and moral hazard. In these studies there is only a single, unobservable, input from the agent; a grant is an ineffective means of encouraging effort in such a setting. In my setting, the observable input, combined with the presence of the agent's profit incentive, opens the door for grants as a relevant instrument. Maurer and Scotchmer (Scotchmer, 2004, ch. 8) put forth an alternative explanation for how grants might overcome moral hazard. The authors posit that the repeated nature of the interaction between researchers and funders instills a discipline in the researcher; researchers that fail to deliver in the past are denied future grants. This insight is useful for understanding the

[7]The severity of the moral hazard problem is measured by the size of the per-unit effort cost, relative to the researcher's profit incentive. The larger (smaller) the effort cost (profit incentive) the more severe is the moral hazard problem.

[8]Although, output-independent rewards may arise in the presence of moral hazard, as a way to provide insurance for a risk-averse agent. Risk sharing considerations are discussed in more detail in the conclusion.

overall structure of the grant-awarding process, but it does not explain why a funder might use a grant over some other pull mechanism.

Prizes have received considerable attention in the literature, with particular attention paid to informational trade-offs between prizes and patents.[9] Both prizes and patents are pull mechanisms, however, and fewer papers have assessed the trade-offs between push and pull incentives. One notable exception, is Fu et al. (2012) (henceforth, FLL). FLL consider a budget-constrained funder who allocates her budget between a prize and a subsidy to each of two agents engaged in a patent race. The prize is a reward received by the first researcher to innovate, while the subsidy is a direct input that increases the productivity of effort. The authors identify trade offs between prizes and subsidies, which are driven, in part, by their effect on the nature of competition.[10] In this paper, I abstract from competition, and instead focus on the informational trade offs between push and pull incentives.

Both FLL and Lewis and Sappington (2000b) consider push mechanisms that involve direct inputs from the funder, which increase the marginal returns to effort. On their own, these inputs have no value to the researcher, and in consequence, the concern that she might abscond with the funds is mitigated. But grants are monetary payments, not direct inputs to R&D. Funders recognize the potential for inappropriate usage of these funds, and structure contracts that clearly outline the expectations of how the funds are to be used (the terms of the grant). My model attempts to capture this aspect, as both the size and terms of the grant (in my setting, the level of investment) are endogenously determined. By incorporating this feature, a new insight is obtained on the use of matching grants - grants that require an investment from the recipient, above and beyond the value of the grant. In the model under pure adverse selection, it is shown that the optimal grant contract induces greater investments from higher types, in exchange for larger grants. But to deter lower ability researchers from over reporting their type, the funder requires investments to be made in excess of the grant. Thus, matching grants serve as an effective sorting device in the presence of adverse selection.

2 The Model

There are two players: the funder and the researcher. The funder may be thought of as a private NGO like the Bill and Melinda Gates Foundation, or a government agency such as

[9]See, for example, Wright (1983), De Laat (1997), Scotchmer (1999), Hopenhayn et al. (2006), and Chari et al. (2012)

[10]FLL show, for instance, that as the degree of asymmetry between contestants increases, the funder allocates more of her budget towards a subsidy, which is given to the less-able contestant in order to "level the playing field". This fosters greater competition and promotes faster innovation.

the NSF. The researcher could be a for-profit firm, an individual, or a nonprofit research foundation. The funder values the development of new technology and offers incentives to the researcher to encourage her to devote her own resources towards R&D. Successful innovation requires both observable and non-observable inputs. The observable input, which I refer to as investment, may be thought of as a capital investment such as the purchase of new lab equipment, or the construction of a new research facility. The unobservable input, which I refer to as effort, may be thought of as the time or energy devoted by the researcher to the success of the project.

Denote by $x \geq 0$, the researcher's investment, and $y \geq 0$, the researcher's effort. The probability of successful innovation is $\theta\alpha(x, y)$. The function, $\alpha : \mathbb{R}_+^2 \to [0, 1]$, is twice continuously differentiable, strictly increasing in both arguments and strictly concave in (x, y). The two research inputs are strategic complements; higher levels of investment increase the marginal returns from effort (and vice versa). The parameter $\theta \in (0, 1]$ captures the researcher's ability; all else equal, a researcher of greater ability is more likely to innovate.[11] The assumptions on α are summarized as follows:

Assumption 1.

(i) $\alpha(0, y) = 0$ *and* $\alpha(x, 0) = 0$ *for all* x, y

(ii) α *is twice continuously differentiable with:* $\alpha_1 > 0$, $\alpha_{11} < 0$, $\alpha_2 > 0$, $\alpha_{22} < 0$, $\alpha_{12} > 0$, *and* $\alpha_{11}\alpha_{22} - \alpha_{12}^2 > 0$.[12]

The marginal cost of investment is normalized to unity, while the marginal cost of effort is $c \geq 0$. I allow for an upper bound on the researcher's effort $\overline{y} \in (0, \infty]$. A finite upper bound is natural if, for example, one interprets effort as the fraction of the researcher's time devoted to the project. Innovation generates a small profit, $\pi > 0$, to the researcher, but this profit is not large enough to warrant R&D activity from the researcher. This profit may be monetary or non monetary. The magnitude of the ratio, $\frac{c}{\pi}$, gives one measure of the severity of the moral hazard problem; the larger this ratio, the more severe is the moral hazard problem.

Success or failure in the project, and the researcher's investment are observable and verifiable, while the researcher's effort and ability are not. The parameter θ is drawn from a continuous distribution with support, $\Theta \equiv [\underline{\theta}, \overline{\theta}] \subset [0, 1]$, according to CDF, F, and corresponding PDF f. The researcher knows the true θ while the funder knows only its

[11] θ may also capture the difficulty of the project; a lower value of θ corresponds to a more challenging project.

[12] For a function, $H(x_1, \cdots, x_m)$, H_k denotes $\frac{\partial H(x_1, \cdots, x_m)}{\partial x_k}$

distribution. To avoid two simultaneous adverse selection problems, we assume that π is known by the funder. Successful innovation carries a benefit, $W > \pi$, to the funder, which could be thought of as the benefit to society, which is not captured by the firm.[13]

Assumption 2.

(i) $\overline{\theta}\alpha_1(0, y)\pi - 1 < 0$ *for all* y

(ii) $\underline{\theta}\alpha_1(0, 0)(W + \pi) - 1 > 0$

(iii) $\underline{\theta}\alpha_2(0, 0)(W + \pi) - c > 0$

Assumption 2(i) ensures that in the absence of any other incentives the researcher would choose to invest nothing (and hence exert no effort). Assumption 2(ii) and (iii) respectively ensure that the benefit to society from innovation is large enough to warrant a strictly positive level of investment and effort from a researcher of any type.

I rule out transfers from the researcher to the funder. There are two types of researcher-to-funder transfers that could potentially arise. First are entry fees, upfront, nonrefundable payments made by the researcher to qualify for a prize. Such payments commonly arise as part of an optimal mechanism in environments with moral hazard (see, for instance, Lewis and Sappington, 2000b, 2001). Substantial entry fees are not common in practice, and moreover, such a payment may be infeasible if the researcher faces a liquidity constraint (Che and Gale, 2003).[14] The second type of researcher-to-funder transfer that could potentially arise, is a payment made to the funder in the event of innovation success. This would amount to a payment of all, or part, of the researcher's profit, π. However, if π is not verifiable by a third party, such contracts would be difficult to enforce. We rule out both of these possibilities. This leaves the funder with two potential tools at her disposal: a grant, which is a reward received by the researcher independent of whether she succeeds or fails, and a prize, which is a reward received by the researcher only if she succeeds.

Since the researcher's investment level is contractible, the model allows the size of the grant/prize to depend on the level of investment. In practice, such an arrangement is common

[13]In some cases W may represent the net present value of consumer surplus associated with innovation. However, in many cases innovation may carry large benefits to society even though consumer surplus is small. Consumer surplus, as typically measured in the literature, is calculated by examining the willingness-to-pay of consumers. But, the willingness-to-pay for a new malaria vaccine, for example, may be quite low since those that need the treatment typically are unable to pay very much. Still, such an innovation could drastically improve the quality of life in many nations and lay the foundation for future economic growth. Thus, the benefit to society may be far greater than just consumer surplus.

[14]Such a liquidity constraint is not inconsistent with the lack of a capital constraint in the model. The researcher may have a line of credit to finance capital investments, while a substantial entry fee would require an alternative source of funding.

with the use of push mechanisms. The size of a research grant typically depends on the costs of a given project, and R&D tax credits are generally tied to investment costs incurred.[15] Examples of pull mechanisms with this feature are less common, but one such example can be found in the U.S. "Orphan Drug Act" of 1983. The act was passed with the intent to provide incentives to pharmaceutical firms to develop treatments for rare diseases within the U.S. One of the provisions in the act is a tax credit, earned by a firm if it successfully develops a new treatment, equal to half the cost of clinical development (Reider, 2000).

2.1 The Funder's Problem

The funder makes a take-it-or-leave-it contract offer to the researcher. It suffices to consider only the use of direct mechanisms. Thus, the researcher announces her type, θ, and the funder specifies an investment, $x(\theta) \geq 0$, a grant, $g(\theta) \geq 0$, and a prize, $V(\theta) \geq 0$. After the contract is formed, the researcher chooses her effort level, $y \geq 0$. If a type θ researcher invests x, and the reward for successful innovation is z, she chooses the effort level:

$$y(\theta, x, z) \equiv \arg\max_{y \in [0, \overline{y}]} \{\theta \alpha(x, y) z - cy\}$$

The payoff to a researcher of type θ who reports her type as $\hat{\theta}$, is:

$$U(\theta, \hat{\theta}) = \theta \alpha \Big(x(\hat{\theta}), y(\theta, x(\hat{\theta}), V(\hat{\theta}) + \pi) \Big) \Big[\pi + V(\hat{\theta}) \Big] - x(\hat{\theta}) - cy(\theta, x(\hat{\theta}), V(\hat{\theta}) + \pi) + g(\hat{\theta})$$

The funder's problem is then:[16]

$$\max_{\{x(\theta), V(\theta), g(\theta)\}} \int_{\Theta} \Big[\theta \alpha \big(x(\theta), y(\theta, x(\theta), V(\theta) + \pi) \big) \left[W - V(\theta) \right] - g(\theta) \Big] dF(\theta) \tag{1}$$

Subject to individual rationality (IR), incentive compatibility (IC), and the non-negativity constraints: $x(\theta) \geq 0$, $g(\theta) \geq 0$, and $V(\theta) \geq 0$. IR and IC require that for all θ, $\hat{\theta} \in \Theta$:

$$\text{(IR)} \quad U(\theta) \geq 0 \tag{2}$$

and

[15] For a recent overview of the U.S. R&D tax credit system see Hemel and Ouellette (2013)

[16] Note that the funder is not a social-welfare maximizer, as the researcher's profit does not enter her objective function. My results generalize immediately to a specification in which the researcher's profit enters the funder's objective function, but weighted by a constant, $a < 1$. This would capture, for example, an environment in which the funder's resources are obtained through socially costly taxation, or through costly fundraising activities.

$$\text{(IC)} \quad U(\theta,\theta) \geq U(\theta,\hat{\theta}) \tag{3}$$

Where $U(\theta) \equiv U(\theta,\theta)$ is the researcher's payoff when she reports her type truthfully. Note that the IR constraint given in (2) follows, as Assumption 2(i) implies the researcher's outside option is zero. The IC constraint requires that truth-telling is optimal for the researcher.

2.2 Efficiency and Perfect Information

It is useful to have a notion of input levels that are socially efficient, and those which are efficient from the perspective of the funder. This leads to the following two definitions

Definition 1.
Let $\mathcal{S}(\theta,x,y) \equiv \theta\alpha(x,y)(W+\pi)-x-cy$. The investment schedule, $x^e(\theta)$, and effort schedule, $y^e(\theta)$, are efficient if

$$\big(x^e(\theta), y^e(\theta)\big) \in \arg \max_{x(\theta)\geq 0, y(\theta)\in[0,\overline{y}]} E_\theta[\mathcal{S}(\theta, x(\theta), y(\theta))]$$

$E_\theta[\mathcal{S}(\theta, x(\theta), y(\theta))]$ is the ex-ante expected total surplus that is generated when the researcher invests according to $x(\theta)$ and chooses effort according to $y(\theta)$. The efficient investment/effort schedules maximize expected total surplus. Note that x^e is uniquely characterized by the first-order condition:

$$\theta\alpha_1(x^e(\theta), y^e(\theta))(W+\pi) - 1 = 0$$

At an interior solution, $y^e(\theta)$ satisfies:

$$\theta\alpha_2(x^e(\theta), y^e(\theta))(W+\pi) - c = 0$$

Note that if $\theta\alpha_2(x^e(\theta), \overline{y})(W+\pi) - c > 0$ then $y^e(\theta) = \overline{y}$. Our next definition concerns mechanisms that are efficient from the perspective of the funder.

Definition 2.
A mechanism, $\mathcal{M}$, attains the first-best for the funder if his equilibrium payoff is equal to $E[\mathcal{S}(\theta, x^e(\theta), y^e(\theta))]$

Definition 2 says that a a mechanism is first-best for for the funder if his payoff under the mechanism is equal to the maximum total surplus. Note that any payoff to the funder greater than $\max_{x(\theta),y(\theta)} E_\theta[\mathcal{S}(\theta, x(\theta), y(\theta))]$ would require that the researcher's ex-ante expected payoff is less than zero. Such a mechanism would violate voluntary participation.

In a world with perfect information (i.e. θ is known, and the probability of success depends only on the observable input, x), the principal could achieve the first-best outcome using either a prize or a grant. Under such conditions, the funder can calculate the socially efficient level of investment, $x^e \equiv \arg\max_x\{\theta\alpha(x)(W+\pi)-x\}$. She then requires the researcher to invest x^e to be eligible for the grant and/or prize. Finally, the funder specifies the grant/prize combination in such a way that the researcher is indifferent between investing x^e or not:

$$g + \alpha(x^e)V = x - \alpha(x^e)\pi$$

Hence, the funder is indifferent between the use of a prize only, a grant only, or any combination thereof that satisfies IC. With fully contractible research inputs and no uncertainty over the researcher's ability, one can simply think of the expected value of a prize as a a transfer which the researcher receives with certainty. The socially efficient investment level is achieved and the funder attains the first-best. To gain some traction on the contracting problem at hand, our analysis will first consider adverse selection and moral hazard in isolation, before combining the two.

3 Adverse Selection

In this section, we study adverse selection in isolation. To abstract from the moral hazard problem, we assume that the researcher's effort is limited and can be supplied at no cost: $\overline{y} < \infty$ and $c = 0$. Thus, WLOG, we may assume that the researcher always exerts full effort so that $y = \overline{y}$. For convenience, we suppress the second argument of α and simply write $\theta\alpha(x)$ as the probability of success. Let $h(\theta) \equiv \frac{1-F(\theta)}{f(\theta)}$ denote the inverse hazard rate. As is standard in the contracting literature, we assume $h'(\theta) < 0$, and $f(\theta) > 0$ for all θ.

If $\pi = 0$, the funder could achieve the first-best through the use of a grant only (Bolton and Dewatripont, 2005, pp. 231).[17] To achieve this outcome, the funder chooses the efficient investment schedule, and simply reimburses the agent with a grant equal to the investment cost. A researcher of any type is indifferent between following the funder's recommendation or not. In the current setting with $\pi > 0$, this contract is not incentive compatible. As the efficient investment level is increasing in type, the researcher would always report being the highest type in order to maximize the probability of innovating and receiving π. Our first result demonstrates, however, the optimal means of funding does still involve only a grant. Our second result establishes that the optimal grant contract elicits larger investments from

[17] This point is also made in Lewis and Sappington (2000b)

researchers of greater ability, offers only partial reimbursement of investment costs, and the funder's payoff falls short of the first-best. Before stating the first result, we re-write the maximand in (1) to reflect the fact that the funder's payoff is the total surplus generated less the rent that accrues to the researcher. The funder's problem may be reformulated:

$$\max_{x(\theta),U(\theta)} \int_{\Theta} [\theta\alpha(x(\theta))(W+\pi) - x(\theta) - U(\theta)]\, dF(\theta) \quad \text{subject to IR and IC} \tag{4}$$

Proposition 1. *When the funder faces adverse selection only, the optimal funding mechanism uses only a grant:* $V^*(\theta) = 0$ *for all* $\theta \in \Theta$

To understand why the funder offers no prize, consider a simple two-type version of the model: $\Theta = \{\underline{\theta}, \overline{\theta}\}$, where $\overline{\theta} > \underline{\theta}$. From (4) it is clear that the funder's payoff is decreasing in the rent captured by the researcher. Under the optimal mechanism, the low type will therefore capture zero rent. To maintain IC, the high type must be permitted to capture at least as much rent as she could by imitating the low type. If $\underline{z}$ is the reward for innovation offered to the low type, the amount of rent that the high type can capture by misreporting is:

$$(\overline{\theta} - \underline{\theta})\alpha(x(\underline{\theta}))\underline{z} \tag{5}$$

Thus, the rent that accrues to the high type is proportional to the low type's reward for innovation. Hence, a prize creates greater information rent for the higher type than does a grant. As these rents are costly to the funder, the optimal funding mechanism uses only a grant. We now explore the properties of the optimal grant contract. Before doing so, note that by using Proposition 1 and standard techniques, it is shown in Appendix A that we may write the funder's problem as:

$$\max_{x(\theta)\geq 0} \int_{\underline{\theta}}^{\overline{\theta}} \Big(\theta\alpha(x(\theta))[W+\pi] - x(\theta) - h(\theta)\alpha(x(\theta))\pi\Big)\, dF(\theta) \tag{6}$$

Subject to the non-negativity constraint: $g(\theta) \geq 0$ and the constraint $x'(\theta) \geq 0$. Let

$$\underline{\theta}^{AS} \equiv \max\{\underline{\theta}, \theta | \theta\alpha'(0)(W+\pi) - 1 - h(\theta)\alpha'(0)\pi = 0\}$$

Note that assumption 2 ensures that $\underline{\theta}^{AS} < \overline{\theta}$.

Proposition 2. *When the funder faces adverse selection only:*

(i) For all $\theta \in [\underline{\theta}^{AS}, \overline{\theta}]$ $x^*(\theta)$ *satisfies:*

$$\theta\alpha'(x^*(\theta))(W+\pi) - 1 - h(\theta)\alpha'(x^*(\theta))\pi = 0 \quad (7)$$

and for all $\theta \in [\underline{\theta}, \underline{\theta}^{AS}]$ $g^*(\theta) = x^*(\theta) = 0$

(ii) *For all* $\theta \in [\underline{\theta}^{AS}, \overline{\theta}]$*, the investment schedule,* $x^*(\theta)$*, and the grant schedule,* $g^*(\theta)$*, are strictly increasing in* θ *with* $g^*(\theta) < x^*(\theta)$

(iii) $x^*(\theta) < x^e(\theta)$ *for all* $\theta < \overline{\theta}$ *and* $x^*(\overline{\theta}) = x^e(\overline{\theta})$

(iv) $\frac{\partial x^*(\theta)}{\partial \pi} > 0$ *if and only if* $\theta > h(\theta)$

Part (i) of Proposition 2 characterizes the optimal investment schedule chosen by the funder. To provide some intuition for the optimal schedule (at an interior solution), note that by re-arranging equation (7) we obtain:

$$\theta\alpha'(x^*(\theta))W = 1 - \theta\alpha'(x^*(\theta))\pi + h(\theta)\alpha'(x^*(\theta))\pi$$

When an researcher of type θ invests x, the expected benefit to the funder is $\theta\alpha(x)W$. Hence, the left-hand side of the above equation is the marginal benefit of investment to the funder. The cost incurred by the funder, when the researcher invests x may be broken up into two terms. First, the funder must compensate the researcher to maintain the individual rationality constraint. When a type θ researcher invests $x > 0$, the loss incurred by the researcher is $\theta\alpha(x)\pi - x < 0$. The funder must compensate the researcher to offset this loss and ensure participation. The first term on the right-hand side of the equation above, $1-\theta\alpha'(x^*(\theta))\pi$, is thus the marginal cost of maintaining the individual rationality constraint. The term $h(\theta)\alpha(x)\pi$ is the cost to the funder of maintaining the incentive compatibility constraint. Researchers of type $\theta > \underline{\theta}^{AS}$ capture strictly positive information rents since they may always profit by imitating lower types. To obtain higher levels of investment from higher types, the funder must offer these high types larger grants. The term $h(\theta)\alpha'(x^*)\pi$ is the marginal cost of maintaining IC.

Parts (ii) and (iii) of Proposition 2 show that in the pure adverse selection environment, the funder only partially reimburses the researcher for investment costs, and that the second-best investment path lies below the efficient path of investment. Both of these facts contrast results in Bolton and Dewatripont (2005) and Lewis and Sappington (2000b). Driving the difference is the presence of the researcher's profit incentive, $\pi > 0$. Since the funder is unable to extract this profit from the researcher, this leaves information rents for all types (except the lowest type). Going back to the two-type model, equation (5) shows that the amount of rent the high type can capture is positively related to the low type's investment level. To limit these rents, the investment level is below the efficient level for all types less than $\overline{\theta}$. The availability of information rents means that the funder cannot offer full reimbursement of investment and simultaneously sort over types. Full investment reimbursement would always create the incentive for the agent to invest as if she were the highest type (the type receiving the largest investment recommendation), in order to maximize the probability of receiving π. Thus, offering only partial reimbursement of investment is crucial to prevent lower types from over reporting. This observation helps to explain the use of matching grants, which require investments above and beyond the value of the grant.

One interesting feature is the relationship between the investment path characterized in (7) and π. Part (iv) shows that $x^*(\theta)$ may either increase or decrease in π. This contrasts the full information environment where, it can be shown, the optimal level of investment is always increasing in π. Intuitively, with full information, an increase in π simply reduces the cost to the funder of maintaining the individual rationality constraint. When the researcher's ability is unknown, $x^*(\theta)$ increases in π if and only if $\theta > h(\theta)$. As previously outlined, the total cost to the funder of incentivizing a type θ researcher to invest $x(\theta)$ may be expressed:

$$x(\theta) - \theta\alpha(x(\theta))\pi + h(\theta)\alpha(x(\theta))\pi$$

The first two terms above give the cost of maintaining IR and the third term is the cost of maintaining IC. As in the full information environment, when π increases, this decreases the cost to the funder of maintaining IR. At the same time, recall from equation (5) that the greater the reward for innovation, the more rent higher types are able to capture. Thus, an increase in π creates more rent for higher types and increases the cost of maintaining IC. An increase in π has two competing effects on the funder's cost function, and the effect on the optimal investment path depends on the balance of these two forces.

4 Moral Hazard

In this section we study the moral hazard problem in isolation, so we assume $c > 0$, but to abstract from adverse selection, we assume F places unit mass on some $\tilde{\theta} \in \Theta$. Let $y(x, z) \equiv \arg\max_{y\in[0,\overline{y}]}\{\alpha(x, y)z - cy\}$ denote the researcher's optimal effort choice when she invests x and the reward for innovation is z. At an interior solution, $y(x, z)$ satisfies the first-order condition $\alpha_2(x, y(x, z))z - c = 0$.

Note that the researcher's effort constraint binds if and only if $\alpha_2(x, \overline{y})z - c \geq 0$. For the moment, we ignore this possibility and assume $\overline{y}$ is arbitrarily large. For simplicity, we also assume $\alpha_2(0, 0)\pi - c > 0$, which ensures the optimal effort level is strictly positive for any level of investment and any prize.[18] The funder's problem may then be written:

$$\max_{x,y,g,V\geq 0} \left\{\alpha(x,y)\big[W-V\big]-g\right\}$$

subject to

$$\alpha(x,y)\big[\pi+V\big]-x-cy+g\geq 0 \tag{8}$$

and

$$\alpha_2(x,y)(\pi+V)-c=0 \tag{9}$$

Equation (8) gives the individual rationality constraint; note that this constraint binds under the optimal contract. If not, the funder could increase the level of investment slightly and increase her payoff. Equation (9) is the first-order condition for the researcher's maximization problem. The strict concavity of alpha in y ensures that this first-order condition is sufficient for a global optimum and that this solution is unique for any x and z, so that this first-order approach is valid.

Proposition 3.
In the model with moral hazard only, when effort is unconstrained the optimal mechanism uses only a prize ($g^* = 0$).

The intuition for Proposition 3 is straightforward; by only rewarding success, a prize creates a greater incentive for unobservable effort than does a grant. This result is in line with standard theory on moral hazard. It is worth noting, however, that in our environment with pure moral hazard, a grant may be used to elicit higher levels of unobservable effort. A grant is effective in inducing higher levels of contractible investment; the complementarity between investment and effort, in conjunction with the presence of the private benefit, π, then leads the researcher to exert higher levels of effort. In this way, grants can be relevant policy tools even when some relevant inputs are unobservable to the research funder. The next result shows that the optimal mechanism may use a grant when the moral hazard problem is weak enough.

Proposition 4.
Suppose effort is constrained and $\frac{c}{\pi} < \alpha_2(x^e, \overline{y})$, then the funder may achieve the first-best through the use of a pure grant ($V = 0$).

[18]This assumption is not necessary for the qualitative nature of our results, but simplifies exposition.

Recall that $\frac{c}{\pi}$ gives one measure of the strength or severity of the moral hazard problem. As the agent's effort cost (profit incentive) increases (decreases), the incentive to exert effort diminishes, and the principal faces a more severe moral hazard problem. Thus, the condition given in Proposition 4 restricts the severity of the moral hazard problem. The key insight of Proposition 4 is that push mechanisms are relevant instruments under moral hazard. This result contrasts many of the classic insights on moral hazard where output-based rewards are key to providing incentives for unobservable effort.

5 Combined Moral Hazard and Adverse Selection

In this section, we consider a version of the model with both moral hazard and adverse selection. That is, θ is unknown to the funder and the cost of effort is non zero: $c > 0$. In general, when both adverse selection and moral hazard problems are present, the optimal mechanism will involve some combination of a grant and prize. As the first result shows, however, if the moral hazard problem is not too strong then the optimal means of funding uses only a grant. Let $x^{AS}(\theta)$ and $g^{AS}(\theta)$ respectively denote the optimal investment and grant schedules characterized in Proposition 2 when the funder faces a pure adverse selection problem. For simplicity, to avoid corner solutions in this section we take W to be arbitrarily large.

Proposition 5.
Suppose $\overline{y} < \infty$ and $\frac{c}{\pi} < \underline{\theta}\alpha_2(x^{AS}(\underline{\theta}), \overline{y})$. Then the optimal funding mechanism uses a pure grant ($V^* = 0$). Moreover, for all $\theta \in [\underline{\theta}, \overline{\theta}]$ the optimal investment and grant schedules are $x^*(\theta) = x^{AS}(\theta)$, $g^*(\theta) = g^{AS}(\theta) + c\overline{y}$, where $x^*(\theta) < g^*(\theta)$ for all θ.

The above result shows that when the investment path $x^{AS}(\theta)$ is sufficient to induce maximal effort from the researcher, the optimal contract uses only a grant and the optimal investment path is unchanged from the version of the model with only adverse selection. Under the supposition of the proposition above, the moral hazard problem is relatively weak; the researcher has enough of an incentive to exert effort without the need for the funder to alter the contract offering from the version of the model with adverse selection only. The funder must offset the additional cost of effort; the grant offered by the funder increases by exactly $c\overline{y}$, relative to the case with adverse selection only.

It is interesting to note that, in some sense, the use of a prize actually become *less* attractive than the pure grant when there exists a weak moral hazard problem, relative to the pure AS environment. For convenience, suppose $\overline{y} = 1$ so that total effort cost when the agent chooses maximal effort is simply c. As compared to the pure AS environment, the funder's payoff under the optimal contract decreases by c - the effort cost. The researcher captures exactly as much rent. But suppose the funder, recognizing the moral hazard problem, sub optimally uses a pure prize. As compared to the pure AS environment, the prize offered to the lowest type must increase by just enough to offset the additional effort cost, in order to maintain IR. But recall from Section 3 that the rent captured by the researcher increases with the reward for innovation. To maintain IC, the funder must therefore allow higher types to capture more rent than in the version of the model with adverse selection only. This means the expected value of the prizes offered to higher types must increase by more than c. The expected payoff to the funder thus decreases by more than c when she uses a prize. Thus, prizes may actually become *less* attractive than grants when the moral hazard problem is present but fairly weak.

Our next result formalizes this intuition. To keep the analysis as simple as possible, suppose $\alpha(x,y) = y\alpha(x)$ and $y \in [0,1]$. Under this specification, the researcher's payoff is linear in effort, and hence, we can simply think of the researcher's effort choice as a binary decision: $y \in \{0,1\}$. If the researcher does not exert effort, $y = 0$, then innovation fails with certainty. If the researcher exerts effort, $y = 1$, then she incurs cost, c, and the probability of success is $\theta\alpha(x)$. If investment is x and the reward for successfully innovating is z then the researcher exerts effort only if: $\theta\alpha(x)z - c \geq 0$.

Before stating this result, let ϕ^{AS} and ϕ^{c} respectively denote the funder's payoff under the optimal contract in the pure AS environment and in the combined AS/MH environment. Also let ϕ_p^{AS} and ϕ_p^{c} respectively denote the funder's maximal payoff when she (sub-optimally) uses a pure prize in the pure AS environment and the combined AS/MH environment. Finally, we define $\tilde{h}(\theta)$ such that:

$$\tilde{h}(\theta) \equiv \frac{\int_{\theta}^{\overline{\theta}} \tilde{\theta} f(\tilde{\theta})\, d\tilde{\theta}}{\theta^2 f(\theta))} \tag{10}$$

Our next result assumes $\tilde{h}$ to be strictly decreasing. This assumption is qualitatively similar to the decreasing inverse hazard rate assumption. Both conditions rule out distributions with long, thin tails, and are satisfied for distributions such as the uniform, the triangular distribution, and many parameterizations of the beta distribution.

Proposition 6.
Suppose $\tilde{h}'(\theta) < 0$. Under the hypotheses of Proposition 5: $\phi^{c} - \phi_p^{c} > \phi^{AS} - \phi_p^{AS}$.

Proposition 6 says that, when a weak moral hazard problem is added on top of an adverse selection problem, the use of a pure prize results in a greater loss to the funder than in the pure AS environment. Our final result characterizes the optimal contract when the moral hazard problem is relatively severe. In this case, a prize is essential to induce effort, but a grant still plays an important role in the mechanism, as a means of limiting the rent that accrues to the researcher. We maintain the binary effort specification.

Proposition 7.
Suppose $\frac{c}{\pi} > \underline{\theta}\alpha(x^e(\theta))$ *for all* θ*. Then the optimal investment path is given by:* $x^*(\theta) = x^e(\theta)$ *for all* θ*. The grant schedule offers full reimbursement of investment (i.e.* $g^*(\theta) = x^*(\theta)$*). Finally, the prize schedule is strictly decreasing and is given by:* $V^*(\theta) = \frac{c}{\underline{\theta}\alpha(x^*(\theta))} - \pi$*. The mechanism is efficient but the funder does not achieve the first best.*

Under the hypotheses of Proposition 7, the moral hazard problem is more severe. This corresponds to an environment in which effort is particularly costly and/or the private benefit of innovation is particularly small. In such settings, a prize is necessary to elicit effort from the researcher. However, higher types do not need large prizes to induce effort; to limit the rents captured by the researcher, the prize schedule is thus decreasing in type. It is also interesting to note that the severe moral hazard problem leads to a socially more efficient outcome than environments with pure adverse selection. In the pure adverse selection environment, the optimal contract specifies an investment level below the socially efficient level. The funder does this in order to reduce the amount of rent that accrues to the researcher. When the moral hazard problem is severe, the funder requires greater investments in order to reduce the size of the prize that must be offered to induce effort. It is more cost effective from the funder's perspective to do this, because large prizes create significant rents for the researcher. It is worthwhile to note that the mechanism characterized in Proposition 7 is weakly implementable, in the sense that an researcher of any type is indifferent between reporting truthfully and not. If a high type underreports, she receives a larger prize, but also invests less, meaning she is less likely to receive the prize than if she reported truthfully. Under the optimal contract, these two effects offset one another, leading the researcher to be indifferent between truthful reporting and not.

The form of the optimal contracts in this combined AS/MH environment contrast existing results in the literature. Driving the difference is the presence of a productive observable input that increases the returns to the non observable input, and the private benefit, π. Lewis and Sappington (2000b) (Henceforth, LS) investigate a similar model in the presence of adverse selection and moral hazard in the context of the sale of a productive asset. The objective in LS is to understand the economic impact of wealth constraints. In our model, we rule out up front payments from the researcher to the funder; this is exactly the case in LS where the researcher has zero wealth. In the absence of wealth constraints, a researcher of high ability can credibly signal this to the funder by making a larger upfront payment to the funder, in exchange for a larger prize. When the researcher has zero wealth, she loses the ability to signal her type, and the funder's only relevant tool is the prize for success. Armed with this one policy tool, the funder in LS is forced to offer the same contract to all types.[19] In my model, the observable investment serves a similar role as an upfront payment to the funder. Higher types are able to distinguish themselves through their willingness to make larger investments. Moreover, because this input is contractible, grants become a relevant policy tool. Finally, because the researcher derives some private benefit for success, no prize may even be necessary under the optimal contract, as shown in Proposition 5. Even when moral hazard is a significant problem, as in Proposition 7, grants are still an important part of the optimal contract. This is because large prizes create significant informational rents for the researcher. The use of a grant reduces the size of the prize that must be offered, and limits the amount of rent the researcher captures.

6 Conclusion and Discussion

In this paper, the incentive properties of push and pull mechanisms are explored in the presence of both adverse selection and moral hazard. I depart from the mainstream literature on moral hazard and examine an environment with both observable and non observable inputs. Grants are shown to be a more efficient means of dealing with adverse selection, and matching grants are shown to be crucial to the funder for sorting over types. In the presence of moral hazard, grants are still a relevant policy instrument. In such an environment, the grant can be used to encourage the contractible input, which increases the marginal returns to the non contractible input and induces greater effort. In a pure moral hazard environment, if the moral hazard problem is weak, the optimal means of funding may involve a grant, a prize, or a combination of the two. When adverse selection and moral hazard interact, the optimal means of funding, in general, depends on the relative strength of the two informational asymmetries. When the moral hazard problem is relatively weak, a grant emerges as the optimal means of funding. Even if the moral hazard problem is strong, grants may play an important role in funding, as a means of limiting the researcher's information rents due to adverse selection.

The model is kept as simple as possible to convey the main insight, but there are a few potentially important considerations, from which I abstract. First, is the role of push and pull mechanisms in allocating risk. Push mechanisms may serve the role of providing insurance to a risk averse agent against unlucky outcomes. In this paper, I abstract from

[19] When there are at least two researchers, LS show that the funder re-gains the ability to tailor the contract to different types by awarding the contract to higher types more often.

such risk-sharing considerations in order to focus solely on the incentive properties of push and pull mechanisms. My model helps bring to light the positive incentive effects of research grants and shows that, even in the absence of such risk-sharing considerations, grants might be preferred over prizes in the presence of moral hazard. Failing to take into account the positive incentive effects of grants may greatly understate their importance in encouraging R&D. Second, I abstract from the role that some push mechanisms might play in helping to overcome constraints on capital. Frictions in financial markets could render researchers incapable of securing traditional sources of funding; research grants are one way for a funder to overcome this issue. Although capital constraints do not explicitly play a role in my model, it's assumed that, absent any other incentive, the researcher invests nothing in R&D. One plausible interpretation is that the agent has access to financial capital, but frictions in the financial sector make it prohibitively costly to acquire. Still, an explanation for the use of grants that rests only on capital constraints is not satisfactory, as it does not address the moral hazard issue. Moreover, a grant is not the only means through which capital constraints may be overcome. A prize, for example, should have the ability to attract financiers to share in the rewards. This is exactly the logic behind the U.S. Department of Labor's (DOL), "Pay for Success" model. According to the DOL:

> Under the Pay for Success model, a government agency commits funds to pay for a specific outcome that is achieved within a given timeframe. The financial capital to cover the operating costs of achieving the outcome is provided by independent investors. In return for accepting the risks of funding the project, the investors may expect a return on their investment if the project is successful; however, payment of the committed funds by the government agency is contingent on the validated achievement of results.[20]

Finally, push mechanisms, in general, do not share with grants the ability to alleviate constraints on capital. R&D tax credits, for example, require up-front investment by a researcher before the benefit is received, and therefore play no role in dealing with this particular issue. Third, in this analysis I abstract from grant/prize competition. My model could be interpreted as the second stage of a two stage game where, in the first stage, some signaling game is played and the funder selects one researcher with whom to contract. In the second stage, the funder makes the contract offering to this researcher. However, there are many different ways in which a competition could be set up, and comparing and contrasting the optimal structure of a grant/prize competition is well beyond the scope of the current analysis.

[20] http://www.doleta.gov/workforce_innovation/success.cfm

References

Bolton, P. and Dewatripont, M. (2005). *Contract theory*. MIT press.

Chari, V. V., Golosov, M., and Tsyvinski, A. (2012). Prizes and patents: Using market signals to provide incentives for innovations. *Journal of Economic Theory*, 147(2):781–801.

Che, Y.-K. and Gale, I. (2003). Optimal design of research contests. *The American Economic Review*, 93(3):646–671.

Choi, J. P. (1992). Cooperative R&D with moral hazard. *Economics Letters*, 39(4):485–491.

De Laat, E. A. (1997). Patents or prizes: Monopolistic r&d and asymmetric information. *International Journal of Industrial Organization*, 15(3):369–390.

Fu, Q., Lu, J., and Lu, Y. (2012). Incentivizing r&d: Prize or subsidies? *International Journal of Industrial Organization*, 30(1):67–79.

Hall, B. H., Mairesse, J., and Mohnen, P. (2009). Measuring the returns to R&D. National Bureau of Economic Research w15622.

Hemel, D. J. and Ouellette, L. L. (2013). Beyond the Patents-Prizes Debate. *SSRN Electronic Journal*.

Hölmstrom, B. (1979). Moral hazard and observability. *The Bell Journal of Economics*, pages 74–91.

Hopenhayn, H., Llobet, G., and Mitchell, M. (2006). Rewarding sequential innovators: Prizes, patents, and buyouts. *Journal of Political Economy*, 114(6):1041–1068.

Innes, R. D. (1990). Limited liability and incentive contracting with ex-ante action choices. *Journal of Economic Theory*, 52(1):45–67.

Jensen, M. C. and Meckling, W. H. (1976). Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure.

Jones, C. I. and Williams, J. C. (1998). Measuring the Social Return to R&D. *The Quarterly Journal of Economics*, 113(4):1119–1135.

Laffont, J.-J. and Tirole, J. (1993). *A theory of Incentives in Procurement and Regulation*. Cambridge, Mass: MIT Press.

Lewis, T. R. and Sappington, D. E. (2000a). Contracting With WealthConstrained Agents. *International Economic Review*, 41(3):743–767.

Lewis, T. R. and Sappington, D. E. (2000b). Motivating wealth-constrained actors. *American Economic Review*, pages 944–960.

Lewis, T. R. and Sappington, D. E. (2001). Optimal contracting with private knowledge of wealth and ability. *The Review of Economic Studies*, 68(1):21–44.

Myers, S. C. (1977). Determinants of corporate borrowing. *Journal of financial economics*, 5(2):147–175.

Reider, C. R. (2000). The orphan drug act: provisions and considerations. *Drug information journal*, 34(1):295–300.

Scotchmer, S. (1999). On the optimality of the patent renewal system. *The RAND Journal of Economics*, pages 181–196.

Scotchmer, S. (2004). *Innovation and incentives*. MIT press.

Stiglitz, J. E. (1974). Incentives and risk sharing in sharecropping. *The Review of Economic Studies*, pages 219–255.

Wright, B.-D. (1983). The economics of invention incentives: Patents, prizes, and research contracts. *The American Economic Review*.

Appendix A: The Pure Grant Problem

In this section we first show that the funder's grant problem given in equation (1) may be expressed as in equation (6). We adopt the binary effort specification, outlined in section 5. Let $y(\theta, x)$ be an indicator function equal to 1 if the agent exerts effort and equal to zero otherwise. The case of $c = 0$ corresponds to the pure AS environment. When the funder uses a pure grant, the payoff to a type θ researcher who reports $\hat{\theta}$, is:

$$U(\theta, \hat{\theta}) = \theta y(\theta, x(\hat{\theta}))\alpha(x(\hat{\theta}))\pi - x(\hat{\theta}) - cy(\theta, x(\hat{\theta})) + g(\hat{\theta})$$

By assumption 1 we have $\overline{\theta}\alpha'(0)\pi - 1 < 0$. This means that in the absence of any other incentives, the researcher would invest nothing, and earn zero profit. Thus, the individual rationality (IR) constraint requires $U(\theta) \geq 0$ for all θ. The incentive compatibility constraint (IC) requires $U(\theta, \theta) \geq U(\theta, \hat{\theta})$ for all $\theta, \hat{\theta} \in \Theta$. The IC constraint is satisfied if and only if

$U_2(\theta,\theta) = 0$ and $U_{12}(\theta,\hat{\theta}) \geq 0$ for all $\theta, \hat{\theta} \in \Theta$ (a proof may be found in Laffont and Tirole (1993) pp. 64 and 121). $U_2(\theta,\theta) = 0$ implies $U'(\theta) = y(\theta, x(\theta))\alpha(x(\theta))\pi$. Since $U'(\theta) \geq 0$, the IR constraint is satisfied so long as $U(\underline{\theta}) \geq 0$.

It may then be verified that $U_{12}(\theta,\hat{\theta}) \geq 0$ if and only if $x'(\theta) \geq 0$ for all θ. So, the IR/IC constraints may be summarized:

$$U(\underline{\theta}) \geq 0; \quad U'(\theta) = y(\theta, x(\theta))\alpha(x(\theta))\pi; \quad x'(\theta) \geq 0$$

Using the definition of U, the funder's problem from (1) may be expressed:

$$\max_{x(\theta),U(\theta)} \int_{\underline{\theta}}^{\overline{\theta}} [\theta y(\cdot)\alpha(x(\theta))[W+\pi] - x(\theta) - cy(\cdot) - U(\theta)] \, dF(\theta)$$

Subject to the IC/IR constraints. The funder's payoff is strictly decreasing in $U(\cdot)$. Clearly, the optimal solution must yield $U(\underline{\theta}) = 0$. This fact, combined with the IC condition on $U'(\cdot)$ imply:

$$U(\theta) = \int_{\underline{\theta}}^{\theta} y(\theta, x(\theta))\alpha(x(\theta))\pi \, d\theta$$

Integrating by parts:

$$\int_{\underline{\theta}}^{\overline{\theta}} U(\theta)\, dF(\theta) = \int_{\underline{\theta}}^{\overline{\theta}} h(\theta) y(\theta, x(\theta))\alpha(x(\theta))\pi \, dF(\theta)$$

Where $h(\theta) \equiv \frac{1-F(\theta)}{f(\theta)}$. Thus, we may re-express the funder's problem as

$$\max_{x(\theta)\geq 0} \int_{\underline{\theta}}^{\overline{\theta}} \Big(\theta y(\cdot)\alpha(x(\theta))[W+\pi] - x(\theta) - cy(\cdot) - y(\cdot)h(\theta)\alpha(x(\theta))\pi\Big) \, dF(\theta) \qquad (11)$$

Subject to the non negativity constraint $x(\theta) \geq 0$ and the IC constraint $x'(\theta) \geq 0$.

Appendix B: The Pure Prize Problem

We assume the binary effort specification as outlined in section 5. The case with adverse selection only corresponds to an environment in which $c = 0$ and $y = 1$.

First note that in this setting, the payoff to a type θ researcher who reports $\hat{\theta}$ is:

$$U(\theta,\hat{\theta}) = y(\theta, x(\hat{\theta}), V(\hat{\theta}) + \pi)(\theta\alpha(x(\hat{\theta}))(V(\hat{\theta}) + \pi) - c) - x(\hat{\theta})$$

As outlined in Appendix A, the IC/IR constraints then require for all θ and $\hat{\theta}$: $U'(\theta) =$

$U_2(\theta,\theta), U_{12}(\theta,\hat{\theta}) \geq 0, U(\theta) \geq 0$. These conditions may be restated: $U'(\theta) = y(\theta, x(\theta), V(\theta) + \pi)\alpha(x(\theta))(V(\theta)+\pi)$, $U(\underline{\theta}) \geq 0$ and $x'(\theta) \geq 0$. Note that $U(\theta) \geq 0$ means

$$y(\cdot)(\theta\alpha(x(\theta))(V(\theta)+\pi) - c) \geq x(\theta)$$

Since $x(\theta) \geq 0$ the IR constraint therefore requires $\theta\alpha(x(\theta))(V(\theta)+\pi) - c \geq 0$. That is, the IR constraint requires that any researcher given a positive level of investment will indeed exert effort. So, without loss of generality we will set $y = 1$ for all types given a positive investment recommendation. Then, a simple substitution into the IC constraint yields:

$$U'(\theta) = \frac{U(\theta) + x(\theta) + c}{\theta}$$

Using the fact that $U(\underline{\theta}) = 0$ and solving this differential equation we obtain

$$U(\theta) = \theta \int_{\underline{\theta}}^{\theta} \frac{x(\tilde{\theta}) + c}{\tilde{\theta}^2}\, dF(\theta)$$

Integrating by parts:

$$\int_{\underline{\theta}}^{\overline{\theta}} U(\theta)\, dF(\theta) = \int_{\underline{\theta}}^{\overline{\theta}} \tilde{h}(\theta)(x(\theta) + c)\, dF(\theta)$$

Where $\tilde{h}$ is defined in equation (10). Substituting $U(\theta)$ into the funder's problem for $V(\theta)$, the problem faced by the funder may be expressed:

$$\max_{x(\theta)} \int_{\underline{\theta}}^{\overline{\theta}} [\theta\alpha(x(\theta))(W+\pi) - x(\theta) - \tilde{h}(\theta)x(\theta) - c(1+\tilde{h}(\theta))]\, dF(\theta)$$

Subject to the non negativity constraint $x(\theta) \geq 0$ and the IC constraint $x'(\theta) \geq 0$. Note that concavity of α implies that the maximand above is strictly concave in x at each θ. Maximization over this expression yields the necessary and sufficient first-order condition:

$$\theta\alpha'(x^p(\theta))(W+\pi) - 1 - \tilde{h}(\theta) = 0$$

Differentiating the first-order condition with respect to θ we see that $\tilde{h}'(\theta) < 0$ is sufficient to ensure $x^p(\theta)$ is strictly decreasing, as required by IC. It is worth noting that, at an interior solution, the investment path $x^p(\theta)$ is chosen independent of the effort cost, c.

Appendix C: Proofs

Proof of Proposition 1. Incentive compatibility requires:

$$U'(\theta) = \theta\alpha(x(\theta))(V(\theta) + \pi)$$

Clearly, the rents captured by the researcher are increasing in the prize. Since the funder's payoff is decreasing in U, optimality requires $V^*(\theta) = 0$. □

Proof of Proposition 2. First note that equation (7) is the (point wise) first-order necessary condition for the problem given in (6). For $\theta > \underline{\theta}^{AS}$ it holds that $\theta(W + \pi) - h(\theta)\pi > 0$. It is straightforward to show that this means the maximand in (6) is strictly concave in $x(\cdot)$ for each θ. This ensures that the first-order conditions indeed characterize a global maximum point wise for each θ. Since $h(\cdot)$ is strictly decreasing, then for all $\theta \geq \underline{\theta}^{AS}$, we must have $\theta\alpha'(0)\pi - 1 - h(\theta)\alpha'(0)\pi \geq 0$. This ensures that the solution given by (7) is indeed non-negative.

Then, differentiating the first-order condition with respect to θ, it is easy to see that since $h'(\theta) < 0$ and $\theta(W + \pi) - h(\theta)\pi > 0$, we have $x^{*'}(\theta) > 0$. Hence, (7) characterizes the optimal level of investment.

Then see that that $U_2(\theta, \theta) = 0$ implies

$$x^{*'}(\theta)(\theta\alpha'(x^*_g(\theta))\pi - 1) + g^{*'}(\theta) = 0$$

By assumption, we have $\theta\alpha'(x)\pi - 1 < 0$ for all $x > 0$. Since $x^{*'}_g(\theta) > 0$ we must have $g^{*'}(\theta) > 0$. Now, note that the efficient investment schedule, $x^e(\theta)$, satisfies:

$$\theta\alpha'(x^e(\theta))(W + \pi) - 1 = 0$$

Since $h(\theta) > 0$ for all $\theta < \overline{\theta}$, equation (7) clearly implies $x^*(\theta) < x^e(\theta)$ for all $\theta < \overline{\theta}$, and since $h(\overline{\theta}) = 0$ we have $x^*(\overline{\theta}) = x^e(\overline{\theta})$. Finally, (iv) may be verified by differentiating (7) with respect to π.

□

Proof of Proposition 3. Let $\mathcal{C}^* = \{x^*, y^*, V^*, g^*\}$ denote the optimal contract and suppose $g^* > 0$. We will proceed by contradiction. Since optimality requires the IR constraint to bind, we must have: $g^* = x^* + cy^* - \alpha(x^*, y^*)(\pi + V)$. Define $\tilde{V} > 0$ such that $\alpha(x^*, y^*)\tilde{V} = g^*$. Finally, let $V^{**} = V^* + \tilde{V}$. Consider the new contract offering: $\mathcal{C}^{**} = \{x, y, V, g\} = \{x^*, y^{**}, V^{**}, 0\}$ where y^{**} is the researcher's optimal effort choice when investment is x^* and the prize is V^{**}. Let $U(\mathcal{C}^*)$ and $U(\mathcal{C}^{**})$ respectively denote the researcher's payoff under the contracts $\mathcal{C}^*$ and $\mathcal{C}^{**}$. We have:

$$
\begin{aligned}
U(\mathcal{C}^{**}) &= \alpha(x^*, y^{**})(\pi + V^{**}) - x^* - cy^{**} \\
&\geq \alpha(x^*, y^*)(\pi + V^{**}) - x^* - cy^* \\
&= \alpha(x^*, y^*)(\pi + V^*) - x^* - cy^* + \alpha(x^*, y^*)\tilde{V} \\
&= \alpha(x^*, y^*)(\pi + V^*) - x^* - cy^* + g^* \\
&= U(\mathcal{C}^*) \\
&= 0
\end{aligned}
$$

The first equality is definitional. The inequality follows from the optimality of y^{**} when investment is x^* and the prize is V^{**}. The second and third equalities follow from the definitions of V^{**} and $\tilde{V}$, respectively . The fourth equality is definitional, and the final equality is true since the payoff to the researcher under $\mathcal{C}^*$ must equal zero. Thus, the contract, $\mathcal{C}^{**}$, satisfies the researcher's IR constraint. From equation (9) it is clear that, at an interior solution, the researcher's optimal effort choice is strictly increasing in the reward for innovation. Since the funder offers a strictly positive grant under the contract $\mathcal{C}^*$, optimality of this contract requires $x^* > 0$. By assumption, $\alpha_2(0,0)\pi - c > 0$, which means $y^* > 0$. Since effort is unconstrained, y^* is interior, and so we must have: $y^{**} > y^*$.

Let $\phi(\mathcal{C})$ denote the funder's payoff under the contract $\mathcal{C}$. It is easily confirmed that, under our assumptions, the optimal contract must generate a strictly positive level of surplus. Since the researcher captures no rent, the funder captures the full amount of surplus generated: $\phi(\mathcal{C}^*) > 0$. This implies $W > V^{**}$. Thus:

$$
\begin{aligned}
\phi(\mathcal{C}^{**}) &= \alpha(x^*, y^{**})(W - V^{**}) \\
&> \alpha(x^*, y^*)(W - V^{**}) \\
&= \alpha(x^*, y^*)(W - V^*) - g^* \\
&= \phi(\mathcal{C}^*)
\end{aligned}
$$

The first inequality holds since $y^{**} > y^*$ and $W > V^{**}$. The last two equalities are definitional. □

Proof of Proposition 4. Our assumptions imply $\alpha_2(x^e, \overline{y})(W + \pi) - c > 0$ and hence $y^e = \overline{y}$. This means x^e solves: $\alpha_1(x^e, \overline{y})(W + \pi) - 1 = 0$. We will show that the funder can achieve the first-best without the use of a prize. Thus, we set $V = 0$. Let x^* be the optimal investment level chosen by the funder, and suppose that the researcher's effort constraint binds: $y(x^*, \pi) = \overline{y}$; we will then verify that this is the case. The funder solves:

$$
\max_{x,g} \{\alpha(x, \overline{y})W - g\} \text{ such that } \alpha(x, \overline{y})\pi - x - c\overline{y} + g = 0
$$

Substituting the constraint into the funder's problem and taking the first-order condition yields:

$$\alpha_1(x^*, \overline{y})(W + \pi) - 1 = 0$$

So $x^* = x^e$ and hence $y(x^*, \pi) = y(x^e, \pi) = \overline{y}$, which verifies that the researcher's effort constraint must bind. Thus, the funder can achieve the first-best through a grant only.

□

Proof of Proposition 5. First, consider the auxiliary problem where both x and y are observable and contractible. This auxiliary problem is a strict relaxation of the principal's problem under moral hazard, and so this provides an upper bound on the principal's payoff. We will show that the principal may indeed achieve this payoff with the grant/prize schedules given in the proposition. Following similar arguments as in the proof of proposition 2, when the principal may choose both x and y, her problem can be expressed:

$$\max_{x(\cdot),y(\cdot),V(\theta)} \int_{\underline{\theta}}^{\overline{\theta}} [\theta\alpha(x(\theta), y(\theta))(W + \pi) - x(\theta) - cy(\theta) - h(\theta)\alpha(x(\theta), y(\theta))(V(\theta) + \pi)]dF(\theta)$$

Subject to non negativity constraints, and the constraint $y \leq \overline{y}$. Examining the problem above, it is clear that the principal will choose $V^*(\theta) = 0$ for all θ. Let $\mu(\theta)$ be the multiplier associated with the constraint on effort. The Lagrangian is then:

$$L = [\theta\alpha(x(\theta), y(\theta))(W + \pi) - x(\theta) - cy(\theta) - h(\theta)\alpha(x(\theta), y(\theta))\pi]f(\theta) + \mu(\theta)(\overline{y} - y(\theta))$$

First suppose that $y^*(\theta) = \overline{y}$ for all θ. We will show that this is indeed the case. At an interior solution, the first-order condition with respect to x gives:

$$\theta\alpha_1(x^*(\theta), \overline{y})(W + \pi) - 1 - h(\theta)\alpha_1(x^*(\theta), \overline{y})\pi = 0 \tag{12}$$

Note that (12) is exactly the same FOC as in the pure AS environment. Thus, $x^*(\theta) = x^{AS}(\theta)$. Now, see that (12), and concavity of α implies:

$$\alpha_1(0, \overline{y})\pi - 1 + \alpha_1(0, \overline{y})(\theta W - h(\theta)\pi) > 0$$

By assumption 2: $\alpha_1(0, \overline{y})\pi - 1 < 0$, which means $\theta W - h(\theta)\pi > 0$. Thus far, we have assumed that the principal chooses $y^*(\theta) = \overline{y}$ for all θ. We now show that this is in fact the case. Note that when $x = x^{AS}$ then $y^*(\theta) = \overline{y}$ iff:

$$\theta\alpha_2(x^{AS}(\theta),\overline{y})(W+\pi)-c-h(\theta)\alpha_2(x^{AS}(\theta),\overline{y})\pi \geq 0$$

Rearranging:

$$\theta\alpha_2(x^{AS}(\theta),\overline{y})\pi - c + \alpha_2(x^{AS}(\theta),\overline{y})(\theta W - h(\theta)\pi) \geq 0 \tag{13}$$

By assumption, $\theta\alpha_2(x^{AS}(\theta),\overline{y})\pi - c > 0$, and as already shown, $\theta W - h(\theta)\pi > 0$. Thus, (13) holds with strict inequality, and the principal does indeed choose $y^*(\theta) = \overline{y}$ for all θ. Then, note that $U(\underline{\theta}) = 0$ implies:

$$g(\underline{\theta}) = -\underline{\theta}\alpha(x^{AS}(\theta),\overline{y})\pi + x^{AS}(\theta) + c\overline{y} = g^{AS}(\underline{\theta}) + c\overline{y}$$

Note that $-\underline{\theta}\alpha(x^{AS}(\underline{\theta}),\overline{y})\pi + c\overline{y} < 0$, which implies $g^*(\underline{\theta}) < x^{AS}(\theta)$. From, the agent's first-order condition it follows that:

$$\theta\alpha_1(x^{AS}(\theta),\overline{y})x'^{AS}(\theta)\pi - x'^{AS}(\theta) + g'^*(\theta) = 0$$

Which gives $g'^*(\theta) = g'^{AS}(\theta) < x'^{AS}(\theta)$. Thus, $g^*(\theta) = g^{AS}(\theta) + c\overline{y} < x^{AS}(\theta)$. Finally, the analysis so far has assumed that the principal was able to choose both x and y; the solution to this problem yields an upper bound on the principal's payoff under moral hazard. But note that under the hypotheses of the proposition: $\theta\alpha_2(x^{AS}(\theta),\overline{y}) - c > 0$, which implies that when the investment schedule is $x^{AS}(\theta)$ an agent of any type will indeed choose to exert maximal effort. This means that even when y is not contractible, the principal can achieve the upper bound on her payoff.

□

Proof of Proposition 6. We first show that, under the optimal contract, the funder's payoff decreases by exactly c in the combined AS/MH environment, as compared to the pure AS case. When W is sufficiently large, we may, without loss of generality, suppose that the optimal investment path is strictly positive for all θ. Under the hypotheses of Proposition 5, the optimal contract in the combined AS/MH environment induces the same investment schedule as in the pure AS environment. It is easily checked that:

$$\phi^{AS} - \phi^c = \int_{\underline{\theta}}^{\overline{\theta}} c\, dF(\theta) = c$$

Appendix B provides the solution to the funder's pure prize problem for any $c \geq 0$. Using this characterization, it may also be verified that

$$\phi_p^{AS} - \phi_p^c = \int_{\underline{\theta}}^{\overline{\theta}} [c + \tilde{h}(\theta)c]\, dF(\theta) > c$$

Hence,

$$\phi_p^{AS} - \phi_p^c > \phi^{AS} - \phi^c \iff \phi^c - \phi_p^c > \phi^{AS} - \phi_p^{AS}$$

□

Proof of Proposition 7. We first establish the following lemma.

Lemma 1. *Suppose $x^*(\theta) > 0$ for all θ. Then $U^{'}(\theta) \geq \frac{c}{\underline{\theta}}$ for all $\theta \in [\underline{\theta}, \overline{\theta}]$.*

Proof. First see that IC requires $U'(\theta) = \alpha(x(\theta))(V(\theta) + \pi)$. Then, since $x^*(\theta) > 0$, it must be that all types exert effort. This means $\theta\alpha(x(\theta))(V(\theta)+\pi) \geq c$ for all θ. In particular, we must have $\underline{\theta}\alpha(x(\underline{\theta}))(V(\underline{\theta}) + \pi) \geq c$, or $U'(\underline{\theta}) \geq \frac{c}{\underline{\theta}}$. IC requires:

$$U_{12}(\theta, \hat{\theta}) = x'(\hat{\theta})\alpha'(x(\hat{\theta}))(V(\hat{\theta}) + \pi) + \alpha(x(\hat{\theta}))V'(\hat{\theta}) \geq 0 \text{ for all } \theta, \hat{\theta} \in \Theta \tag{14}$$

and

$$U_2(\hat{\theta}, \hat{\theta}) = \hat{\theta}\Big[x'(\hat{\theta})\alpha'(x(\hat{\theta}))(V(\hat{\theta})+\pi)+\alpha(x(\hat{\theta}))V'(\hat{\theta})\Big] - x'(\hat{\theta}) + g'(\hat{\theta}) = 0 \text{ for all } \hat{\theta} \in \Theta \tag{15}$$

Equations (14) and (15) together imply: $x'(\theta) \geq g'(\theta)$ for all $\theta \in \Theta$. Now, note that $U'(\theta) = \alpha(x(\theta))(V(\theta)+\pi)$ may be equivalently expressed: $U'(\theta) = \frac{U(\theta)+c+x(\theta)-g(\theta)}{\theta}$. See that:

$$\begin{aligned}
U''(\theta) &= \frac{\Big[U'(\theta)+x'(\theta)-g'(\theta)\Big]\theta - \Big[U(\theta)+c+x(\theta)-g(\theta)\Big]}{\theta^2} \\
&= \frac{1}{\theta}\Big[U'(\theta) + x'(\theta) - g'(\theta) - \frac{U(\theta)+c+x(\theta)-g(\theta)}{\theta}\Big] \\
&= \frac{1}{\theta}\big[U'(\theta) + x'(\theta) - g'(\theta) - U'(\theta)\big] \\
&= \frac{1}{\theta}\big[x'(\theta) - g'(\theta)\big] \\
&\geq 0
\end{aligned}$$

So, $U'(\underline{\theta}) \geq \frac{c}{\underline{\theta}}$ and $U''(\theta) \geq 0$, which implies $U'(\theta) \geq \frac{c}{\underline{\theta}}$.

□

To minimize the rent captured by the researcher, the funder sets $U(\theta^m) = 0$ and $U'(\theta) = \frac{c}{\theta^m}$, so long as it is feasible to do so. This would imply: $U(\theta) = c\left(\frac{\theta-\underline{\theta}}{\underline{\theta}}\right)$. The payoff to the funder is then:

$$\int_{\underline{\theta}}^{\overline{\theta}} \left[\theta\alpha(x(\theta))(W+\pi) - x(\theta) - c - c\left(\frac{\theta-\underline{\theta}}{\underline{\theta}}\right)\right] dF(\theta) \tag{16}$$

Pointwise maximization of the expression above yields the first-order condition:

$$\theta\alpha'(x^*(\theta))(W+\pi) - 1 = 0$$

Hence, $x^*(\theta) = x^e(\theta)$. To ensure feasibility, we now establish that the grant and prize schedules are nonnegative. $U'(\theta) = \frac{c}{\underline{\theta}}$ implies $\alpha(x^*(\theta))(V^*(\theta)+\pi) = \frac{c}{\underline{\theta}}$ or $V^*(\theta) = \frac{c}{\underline{\theta}\alpha(x^*(\theta))} - \pi$. Under our assumptions, V^* is in fact nonnegative. This then gives:

$$U(\theta) = c\left(\frac{\theta-\underline{\theta}}{\underline{\theta}}\right) - x^*(\theta) + g^*(\theta)$$

Which implies $x^*(\theta) = g^*(\theta) \geq 0$. Finally, we establish that the integrand in (16) evaluated at $x^*(\cdot)$ is non-negative. Let $\Gamma(\theta)$ denote this integrand when investment is $x^*(\cdot)$. Note that $\Gamma(\underline{\theta})$ is maximum total surplus at $\underline{\theta}$, which is assumed to be non-negative. Also note that

$$\Gamma'(\theta) = \alpha(x^*(\theta))(W+\pi) - \frac{c}{\underline{\theta}}$$

Which is again positive by our assumptions. Hence, $\Gamma(\theta) \geq 0$ for all θ.

□