Push and Pull Funding for Social Innovation

Hassan Sayed* Christopher Snyder†

PRELIMINARY DRAFT/DO NOT CIRCULATE

January 9, 2026

Click here for latest version.

Abstract

If a mechanism designer would like to encourage firms to develop a socially beneficial innovation, is it cheaper to reward them ex-ante with push funding — such as research grants paid unconditionally — or ex-post with pull funding — such as prizes only paid upon successful innovation? We study this question in a multifirm setting where firms' probabilities of developing an innovation and costs of research and development are private information. We show that the most cost-efficient contract always involves a positive amount of pull funding, and derive a simple distributional condition under which a designer should also utilize a positive amount of push funding. While push funding screens firms based on their costs of research and development, because pull funding is paid out only upon the successful development of an innovation, it screens firms on both their costs and likelihoods of success. On the one hand, this selects in firms who are more likely to successfully develop the innovation. On the other, it allows firms to collect rents on two dimensions instead of one, generating a tradeoff. Additionally, we provide a simple condition characterizing when a pure pull contract is cheaper than a pure push contract, show that the form of pull funding that entails the lowest fundraising cost for the designer is a shared pot of money, and discuss how our results can be microfounded through a social planner's problem.

*Market Shaping Accelerator, Center for Global Development: hsayed@cgdev.org.
†Department of Economics, Dartmouth College: christopher.snyder@dartmouth.edu.

1 Introduction

The development of socially beneficial innovations often rests on external transfers from governments, nonprofits, and philanthropists to university researchers, government contractors, and private firms. [1] These transfers incentivize research into innovations that may be underserved by traditional markets, particularly when market failures such as lack of property rights preclude private sector investment. Consider, for example, a pharmaceutical company that researches a novel use for a generic drug. The company must spend millions of dollars on clinical trials to discover a new use. However, because the drug has a generic formulation and lacks patent protection, the firm will be forced to sell at a competitive price, meaning profits will be slim. Thus, the limited research that does take places in generic drug repurposing is often financed by external grants from sources like the National Institutes of Health (van der Pol et al., 2023).[2] Similarly, inventions that are primarily purchased by governments may be subject to buyer-side bargaining, forcing firms to sell near marginal cost and generating a hold-up problem that makes them unwilling to undertake R&D investments.

Transfers for stimulating socially beneficial innovations traditionally take the form of "push" funding: a funder provides up-front payments to firms to research an innovation. Federal grants, for instance, constitute push funding; a university lab may receive funding to carry out a clinical trial to repurpose a generic drug, but receives transfers regardless of the trial's conclusions. "Pull" funding, by contrast, pays firms only upon the successful discovery of innovations. These may include advance market commitments (AMCs), volume guarantees, or lump-sum prizes.[3]

If a mechanism designer wants to ensure a specific innovation is developed with a sufficient likelihood, is push or pull funding more cost-efficient? This paper provides a theoretical answer to this question by showing that designers should use a *combination* of both and, in some cases, only pull funding. We consider a market of firms that vary in both their costs of research and development and their ex-ante probability of succeeding at developing an innovation. These costs and probabilities are private information to the firms, and are not known to the mechanism designer, who holds a prior belief over their distribution. The designer can provide ex-ante push payments to firms; pull payments that are only paid upon the successful development of the innovation; or a combination of both. The designer is interested in achieving the development of the innovation with a certain probability — meaning she would like multiple firms to attempt the innovation in

[1] In the 2025 fiscal year alone, the United States government disbursed nearly $200 billion in federal funding for research and development, with close to $50 billion going to the National Institutes for Health and $7 billion to the National Science Foundation (Blevins et al., 2025). The Gates Foundation announced in 2022 an initiative to accelerate billions of dollars in investments to develop innovations that mitigate climate change, improve food security, and eradicate diseases (Gates, 2022).

[2] By contrast, *new* drugs provide firms 20 years of patent exclusivity, allowing them to reap far greater profits (U.S. Food and Drug Administration, 2020). In the case of generic drug repurposing, (Budish et al., 2025) estimate the value of missing research and development in this market at $2-10 trillion.

[3] An AMC resulted in the development of a vaccine for pneumococcal vaccine in the early 2010s; this included a funding pool that donor countries promised to use towards the purchase of a pneumococcal vaccine, conditional on its development (Kremer et al., 2020). Innovation prizes include those such as the Ansari XPrize for space travel (XPRIZE Foundation, 2004).

parallel. Conditional on meeting this probability, she would like to choose the payment scheme with the lowest expected cost.

The first result of the paper is to show that, except in boundary cases, the cheapest mechanism that achieves a given target probability of success always involves a positive amount of pull funding. This result revolves around the following tradeoff. Because push funding involves ex-ante payment, firms select in if their costs of research and development are below the value of that payment. Firms collect rents — driving up expected costs — but only on a single dimension. By contrast, since pull funding is only paid out upon success, firms will only participate if they have sufficiently low costs *and* sufficiently high likelihoods of success. Thus, while pull funding is more efficient at selecting in firms who are likely to succeed, firms are able to collect rents on *two dimensions*. When the initial amount of pull funding is zero, the rents firms reap from the two mechanisms is similar, but pull's ability to *screen* on probabilities of success makes it more efficient.

The result also shows that, in many cases, contracts involve both push funding and pull funding. In particular, it provides a sufficient condition on the distribution of costs that characterizes when the optimal mechanism involves either a strictly positive amount of push or *no* push. This condition is related to the monotone hazard rate, and characterizes the ratio of the moments of the cost distribution, which in turn represent the *relative* costs of push and pull.

The paper's second result compares *pure* push and *pure* pull contracts — i.e. those in which the designer *only* pays firms via push or pull — representing a simplification of the problem that real-life policymakers often utilize. It shows that when the designer's target probability of success is sufficiently high *or* the number of firms in the market is small, pure push is cheaper than pure pull, and that when the target probability is low *or* the number of firms in the market is large, pure pull is cheaper. The intuition is that when the target probability is sufficiently high, the designer needs most of the firms in the market to participate. If the designer offers a push payment equal to the *largest possible cost* a firm could have, she is able to screen all the firms in with probability 1. By contrast, even for arbitrarily large pull payments, there are always a small number of firms with very low probabilities of success who will be unwilling to participate. Thus, as the target probability grows large, the cost of pull approaches infinity, while the cost of push is bounded above. For small target probabilities, the intuition is similar to the case of combination payments; the designer needs to screen in only a small number of firms to achieve their target probability of success. Because the excess rents firms reap with a small amount of pull is counterbalanced by their greater likelihoods of success, pure pull is optimal. In both cases, as the number of firms grows large, the designer is able to shrink the set of firms she needs to screen in, since the probability that at least some of those large number of firms are efficient increases.

Crucially, our results do not rely on failures in contract enforcement. For example, push funds may be susceptible to embezzlement and moral hazard,[4] while pull payments may preclude par-

[4]For example, Kremer (2002) documents how USAID's attempt to use. $60 million in push funding to develop a malaria vaccine was plagued by fund diversion and theft.

ticipation from firms who lack up-front cash to finance research and development. Instead, we assume contracts are enforceable to highlight a fundamental difference in how these funding mechanisms *screen* firms. Firms promised push payments will attempt innovation as long as their costs of research and development are below the value of a push payment. Pull payments incentivize participation only if firms both have sufficiently low costs *and* they are sufficiently likely to develop an innovation. The combination of these two effects with the nature of rents that firms collect drives the results of our model.

Because pull funding often requires donors and governments to commit to funds in advance, there is also a question of what form pull funding should take to minimize *fundraising costs*. For example, should all firms be offered an individual lump-sum prize? Should a designer offer large lump-sum prizes to the first k firms that succeed and no others? We show that the form of pull funding that minimizes the amount a designer has to fundraise is given by a large pot of money that is equally divided between successful firms. The intuition is that the amount a designer has to fundraise is pinned down by the *maximum* amount she has to pay out. If the designer offered lump-sum prizes to everyone who participated, and up to N firms participated, she would have to raise N times the lump-sum prize. However, by raising a single pot of money, firms receive an expected utility greater than the lump-sum if few other firms succeed and lower than the lump-sum if lots of firms succeed; because the latter case occurs with low probability, the designer is able to utilize this feature to save on the amount she has to fundraise.

Finally, we connect the designer's problem to a social optimization problem, where the designer achieves a value V if an innovation is developed (by any firm). The designer trades off this value with the expected cost of developing the innovation. We provide a simple characterization of the designer's optimal choice, and using a duality result, show that any target probability of success can be rationalized by some social value V using a simple expression.

Our paper is broadly connected to the literature on optimal incentives for innovation, and is most closely related to Weyl & Tirole (2012). The key difference is that while their paper is concerned with how providing market power allows regulators to *screen* projects with higher social surplus, our paper is concerned with settings where firms cannot as-is receive *any* surplus, for example, because providing market power may be unrealistic, and where the designer knows in advance what the exact innovation will be. That is, our paper applies to settings where market failures preclude market power's ability to screen, as in their paper — thus differentiating our work from a long line of literature on the use of market signals to provide innovation incentives, including Acemoglu & Linn (2004), Chari et al. (2012), Veiga & Weyl (2016)[5], and Ahlvik & van den Bijgaart (2024).

Moreover, our framework is built to answer the question of what policies — such as prizes or grants — incur the lowest expected cost for a designer, i.e. what optimal industrial policy should be in settings where governments or other funders are choosing *how* they should disburse funds.

[5]Our condition on the ratio of distributions in determining when a pure pull contract can be improved with some push is somewhat related to that in Veiga & Weyl (2016).

To this end, our paper is broadly related to the literature on the utilization of pull mechanisms to stimulate innovation in areas where the private sector does not traditionally, including studies of AMCs (Kremer et al., 2022), the direction of innovation (Bryan & Lemus, 2017), and the use of follow-on contracts to motivate procurement of efficient innovation (Che et al., 2021). Our result on prize-sharing also connects us to the literature on information disclosure in innovation contests (Halac et al., 2017).

Because pull funding screens firms on both costs and probabilities of success, our paper is also related to the mechanism design literature on multidimensional screening. A review of the earlier literature is provided by Rochet & Stole (2003), who show how screening on multiple dimensions can link incentive contracts, in contrast to single-dimensional screening, but are oftentimes difficult to solve. Papers that solve for optimal mechanisms often rely on robustness or other techniques (see Carroll (2017), Deb & Roesler (2024), Yang (2025), and literature within). In our model, however, the two variables consisting firms' types — costs and probabilities of success — have particular interpretations. Firms willingness to take on an innovation project via push is only determined by their costs, while their willingness to take on pull funding is determined by a relative *ratio* of their costs to probabilities of success. At the same time, the designer directly benefits from firms with higher probabilities of success, while firms pose no (direct) externalities on each other through their types. While this draft considers a simpler class of mechanisms, instead of the optimal mechanism, we are able to gain tractable results in large part due to the specific form our multidimensional screening problem takes. This form stands in contrast to many other problems in the literature, such as those where firms face negative externalities if competitors win a contract (e.g. thorugh downstream competition) or consumers have multidimensional preferences over goods. Future drafts will study our problem in a more general mechanism design framework with an eye for the optimal mechanism.

Section 2 overviews our model. Section 3 goes over our main results, beginning with a characterization of firm entry before moving to Theorem 1, which shows that the optimal payment scheme always utilizes some pull funding, and characterizes when the optimal scheme involves some push. Next, we move to compare and contrast pure push and pull mechanisms; costs of fundraising for the mechanism designer; and how the problem we solve in Theorem 1 can be microfounded through a social planner's problem. We conclude with a brief discussion of the results.

2 Model

Setup Consider a market of firms indexed $i = 1, \dots, N$ who are overseen by a mechanism designer d. There are two time periods. In the first time period, each of these firms can choose to enter a project to develop a *pre-specified innovation* at cost $c_i \geq 0$. Let the variable e_i be equal to 1 if firm i undertakes the project and 0 otherwise. Conditional on entering, in the second time period, they succeed at developing that innovation with probability $p_i \in [0, 1]$. Similarly, let the variable

s_i be equal to 1 if firm i successfully develops the innovation and 0 otherwise.

The costs c_i represent firms' ex-ante, expected research and development costs from attempting to develop an innovation. The probabilities p_i represent firms' capacity to actually develop the innovation, conditional on incurring costs. [6] The number of prospective firms N can be thought of as representing the potential market size of innovators. Markets with high N may involve dozens of labs and private sector entities involved in the development of an innovation, while those with low N may have just a few small players developing a specific technology.

In the absence of any intervention from the mechanism designer d, we assume that firms receive no reward from successful development of the project: their utility is $-c_i$ regardless of whether the project succeeds or fails. We use bold letters throughout to indicate *profiles* of variables — e.g. $\mathbf{c} = \{c_j\}_{j=1}^n$ denotes the profile of firms' costs. We use the notation $-i$ to refer to the profile of a variable for all firms, excluding firm i.

We will refer to the tuple (p_i, c_i) as a firm's *type*, and assume it is private information for firm i. It is known to (and observed only) by firm i, and unknown to other firms and the mechanism designer. However, we assume firms' types are drawn *independently* by Nature from a continuous distribution $F(p, c)$. That is, firms have (symmetric) priors about other firms' types, which is the same as the mechanism designer's. We denote $f(p, c)$ the density of F, which has support on $p_i \in [0, 1]$ and $c_i \in [0, \overline{c}]$ for some $\overline{c} > 0$.

Designer's Mechanisms Without the mechanism designer's intervention, firms receive no utility upon successfully completing the innovation. This generates scope for the designer d to shape firm's entry decisions by choosing a mechanism that provides transfers to firms in the first or second time periods, potentially conditional on eliciting information on firms' types.

Denote $(\hat{p}_i, \hat{c}_i)$ firm i's reported type. A *mechanism* $\mathbf{M}$ is a function that specifies transfers to each of the firms $i = 1, \ldots, n$ across both time periods conditional on the vector of reported types $(\hat{\mathbf{p}}, \hat{\mathbf{c}})$, decisions to undertake the project $\mathbf{e}$, and the profile of firms who succeed at the project $\mathbf{s}$. That is, a mechanism specifies transfers $t_i^1(\hat{\mathbf{p}}, \hat{\mathbf{c}}, \mathbf{e})$ in the first period and $t_i^2(\hat{\mathbf{p}}, \hat{\mathbf{c}}, \mathbf{e}, \mathbf{s})$ in the second period. Fixing the profile of reports $(\hat{\mathbf{p}}, \hat{\mathbf{c}})$, firm i's expected utility is given by:

$$e_i\Bigg(t_i^1(\hat{\mathbf{p}}, \hat{\mathbf{c}}, 1, e_{-i}) - c_i + \mathbb{E}_{s_{-i}}\big[p_i \cdot t_i^2(\hat{\mathbf{p}}, \hat{\mathbf{c}}, 1, e_{-i}, 1, s_{-i}) + (1 - p_i) \cdot t_i^2(\hat{\mathbf{p}}, \hat{\mathbf{c}}, 1, e_{-i}, 0, s_{-i})\Bigg) \tag{1}$$

Equation (1) represents firm i's utility conditional on entering ($e_i = 1$). Conditional on entering, firm i receives a transfer $t_i^1(\hat{\mathbf{p}}, \hat{\mathbf{c}}, 1, e_{-i})$, which is also a function of others' entry decisions, less the cost c_i of undertaking the project. In the next time period, i succeeds at the project with probability p_i, in which case it receives a transfer $t_i^2(\hat{\mathbf{p}}, \hat{\mathbf{c}}, 1, e_{-i}, 1, s_{-i})$, which it then integrates across all the

[6] The costs and probabilities can represent the aggregation of multiple stages of project development. For example, consider a drug that must pass phases A and B of a clinical tiral. p_i represents the ex-ante likelihood a drug makes it through clinical trials: i.e. the probability the drug moves from A to B multiplied by the (conditional) probability the drug succeeds in phase B. c_i would be the cost of phase A plus the conditional probability of moving from phase A to B multiplied by the cost of phase B.

different profiles of other firms s_{-i} that might succeed. If it fails, it may still receive a transfer $t_i^2(\hat{\mathbf{p}}, \hat{\mathbf{c}}, 1, e_{-i}, 0, s_{-i})$.

While equation (1) documents a very general set of mechanisms, this paper's goal is to compare when a mechanism designer should emphasize ex-ante "push" funding — i.e. focusing on compensating firms via transfers t_i^1, without regard for whether firms develop an innovation — as opposed to ex-post "pull" funding — i.e. focusing on compensating firms via transfers t_i^2 only after they have successfully developed an innovation. To highlight the key intuitions driving our results, it is sufficient to focus on a very simple class of mechanisms that satisfy the following two assumptions.

Assumption 1. $t_i^1(\hat{\mathbf{p}}, \hat{\mathbf{c}}, 1, e_{-i}) = R \geq 0$. *Conditional on entering, any firm receives* R*, regardless of their report.*

Assumption 2. $t_i^2(\hat{\mathbf{p}}, \hat{\mathbf{c}}, 1, e_{-i}, 1, s_{-i}) = x$ *and* $t_i^2(\hat{\mathbf{p}}, \hat{\mathbf{c}}, 1, e_{-i}, 0, s_{-i}) = 0$. *Conditional on succeeding, a firm receives* x*, regardless of their report, and otherwise no additional payment in the second period.*

Assumption 1 says that conditional on particpating, all firms receive an up front payment of R, regardless of the type they report. Assumption 2 says that any firm that successfully develops the innovation receives an (expected) reward of x, and no additional reward if they do not.[7] Note that both R and x can embed a variety of prize-sharing rules; for example, if n firms succeed at developing the innovation, they may equally split a large pot of money, or may each get an individual prize of value x.

The set of mechanisms we consider can thus be characterized by the tuple (R, x). We refer to R as a *push payment*. R should be thought of as an advance payment to firms to research a problem and engage in research and development. For example, federal grant funding in the United States often takes this form. If $R > 0$ and $x = 0$, we refer to this as a *pure push* contract.

Relatedly, we refer to x as a *pull payment*. x should be thought of as a prize, the expected payment from an advanced market commitment, or the total value of a subsidy or top-up payment that can be reaped *only when* a firm develops an innovation. If $R = 0$ and $x > 0$, we refer to this as a *pure pull* contract.

The key question this paper asks is: to achieve a social objective, should the mechanism designer d utilize pure push mechanisms, pure pull mechanisms, or a combination? In future drafts, we return in a later section to the case of more general mechanisms, which primarily rely on eliciting information about firms' types, and shows how the insights of these simple mechanisms generalize.

Social Objective The mechanism designer d is concerned primarily with the development of the innovation: she wants to ensures the innovation is developed with an exogenous *target probability*

[7] The latter part of the assumption can be microfounded by rewriting rewards received at time 2 into a 'baseline" reward plus the additional reward received if the project is successfully completed, and shifting the "baseline" reward into the transfer in the first period.

of success θ. Given this target, she is interested in minimizing the *expected cost* of the mechanism — i.e. the expected cost of transfers she pays across the two time periods. The main theoretical results generalize as long as d would like a sufficient number of firms to participate, and one of our results shows how this preference can be microfoundeded via a social planner's problem.

For example, a government may want to incentivize the development of a vaccine to halt a pandemic. As long as firms can produce at scale, the designer simply wants to ensure that a vaccine is developed with *some* probability. Or, consider a nonprofit that wants agricultural actors to design climate resilient maize. While there may be benefits from having multiple climate resilient varieties, the nonprofit's primary objective is to ensure that farmers can benefit from at least one variety.

Utilizing the simple mechanisms laid out in Assumptions 1 and 2, the *expected cost* to the designer can be written as follows. Suppose a firm i enters if and only if $(p_i, c_i) \in E$ for some set $E \subseteq [0,1] \times \mathbb{R}^+$. Fixing this subset E, d's expected cost is given by:

$$N \cdot Pr\big((p_i, c_i) \in E\big) \cdot [R + x \cdot \mathbb{E}[p_i | (p_i, c_i) \in E] = N \cdot \int_E R + x \cdot p_i dF(p_i, c_i) \tag{2}$$

The first term on the left hand side represents the *expected number* of firms that will enter, $N \cdot Pr\big((p_i, c_i) \in E\big)$. Each of these participating firms is compensated R in push payments. Then, x is paid out only when firms succeed at developing the innovation; in expectation, since firms' types are distributed independently, the expected payout for a representative firm is given by $x \cdot \mathbb{E}[p_i | (p_i, c_i) \in E]$. The expression on the left hand side can also be written more succintly via the integral on the right hand side.

Given a subset E such that firms enter if and only if $(p_i, c_i) \in E$, the expected probability that at least one firm succeeds at the innovation is given by

$$1 - \prod_{i=1}^{N} (1 - Pr((p_i, c_i) \in E) \cdot \mathbb{E}[p_i | (p_i, c_i) \in E]) = 1 - \left(1 - \int_E p_i dF(p_i, c_i)\right)^N \tag{3}$$

which makes use of the fact that types (p_i, c_i) are drawn *independently*. Given a target probability of success θ, we have that

$$1 - \left(1 - \int_E p_i dF(p_i, c_i)\right)^N \geq \theta \iff \int_E p_i dF(p_i, c_i) \geq 1 - (1 - \theta)^{\frac{1}{N}}$$

Finally, note that this set of entering firms E is an equilibrium object, and additionally depends on R and x. Thus, the designer's problem can be given by

$$\min_{R,x} N \cdot \int_{E(R,x)} R + x \cdot p_i dF(p_i, c_i) \text{ s.t. } \int_{E(R,x)} p_i dF(p_i, c_i) \geq 1 - (1 - \theta)^{\frac{1}{N}} \tag{4}$$

The next section characterizes the set of entering firms E as a function of R and x and in turn provides a solution to equation (4).

Discussion Our framework is meant to capture situations where firms cannot reap profits upon the successful development of an innovation, thereby precluding innovation entirely. This may be due to lack of property rights, hold-up problems in innovation procurement, political economy frictions, or other market failures. That is, our model is meant to capture precisely innovations that *depend* on industrial policy or other external transfers to come to fruition. Relatedly, this also means that there is little scope for utilizing market power or other institutional tools as a means of screening the potential value of an innovation.

To this end, a key feature of our environment is that the innovation is pre-specified by the designer. For example, the designer may want firms to develop a vaccine for a specific disease that achieves a certain efficacy and safety profile. While technological agnosticism is permitted — for example, the designer may allow both mRNA and non-mRNA vaccines — our framework should *not* be thought of as a general model for incentivizing any socially beneficial innovations. Indeed, the benefits of many innovations may be unknown until they are discovered. Instead, our goal is to capture settings where the designer can articulate in advance the sort of product she would like firms to develop.

Many of our technical assumptions are meant to focus our analysis on the role of *screening* in generating differences in expected costs between push and pull mechanisms. Firms in the model, for example, cannot embezzle a share of funds (thus creating moral hazard). Firms also do not possess limited liability (i.e. they lose the entirety of c_i if they are unsuccessful). If we were to add these effects to the model, this would increase the effective cost of push, thus making pull funding more attractive. Similarly, if we were to allow firms to be capital-constrained — e.g., because they do not have cash on hand to finance all of c_i without external funding — this would create a lower bound on the size of push funding necessary to guarantee firm participation, making pull funding less attractive. Finally, we assume below that firms' costs and probabilities are drawn independently. While our results go through if we allow for correlation between the two, these mechanically shift expected costs in the directions of push or pull.[8] By abstracting from all these different effects and assuming contracts are enforceable as-is, we are able to focus solely on the role of how R and x generate changes in the *composition* of firms participating and, thus, the designer's expected costs.

3 Analysis

We first study how firms respond to d's incentives, which pins down the set $E(R, x)$ of firms that enter and attempt the innovation in equilibrium. Subsection 3.2 shows that, in general, any pure push mechanism can be made cheaper in expectation by replacing some push funding with pull. It also derives distributional conditions under which a pure pull mechanism can be improved by mixing in push funding.

[8]For example, if an increase in c_i caused a first order stochastic shift in the distribution of p_i, our results would skew in favor of pull, while the opposite would skew in favor of push.

We then study what mechanism minimizes expected costs when the designer must choose only between pure push and pure pull mechanisms — as is common in many public policy settings. We show that when θ is small, pure pull is cheaper in expectation than pure push, and that when θ is large, pure push is cheaper in expectation. Subsection 3.4 shows that, among all mechanisms that involve pull, a mechanism where firms share a single pot of money achieves the lowest *maximal* payout — and thus involves the lowest "fundraising cost" for the designer d.

Finally, subsection 3.5 studies a slightly more general problem where a designer trades off the probability of achieving some social surplus with its expected cost, shows its connections to the problems studied earlier in the section, characterizes its solution, and shows how this problem microfounds the designer's pursuit of a target probability of success.

3.1 Firm Behavior

Given a "push" payment of R and a pull reward of x, firm i has an incentive to enter if and only if

$$R + p_i \cdot x - c_i \geq 0$$

Any firm with $c_i \leq R$ has an incentive to enter. Firms with $c_i \geq R$ have an incentive to enter if and only if

$$x \geq \frac{c_i - R}{p_i}.$$

These two equations jointly characterize the set $E(R, x)$ of firms that enter for a given R and x. Figure 1 also illustrates this visually for a mechanism where R and x are both > 0.

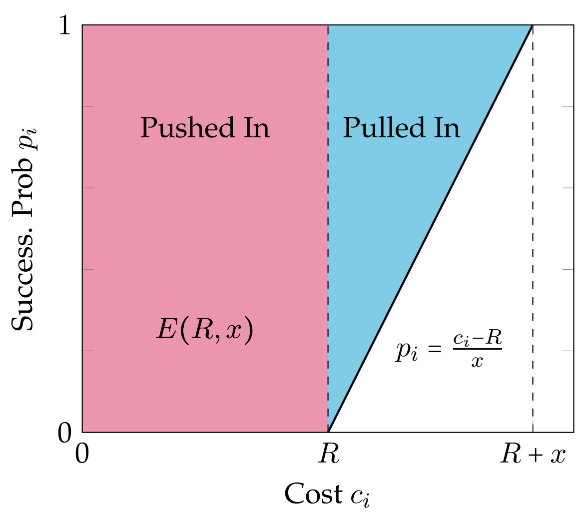

Figure 1: Set of Entering Firms $E(R, x)$

$E(R, x)$ is given by the trapezoid in the figure, illustrating the types (p_i, c_i) that enter into the mechanism. The red rectangle constituting the left side of the trapezoid represents firms whose participation is guaranteed by the push payment R. The blue triangle constituting the right of the trapezoid represents firms who require not only the push payment R but also the pull payment x.

When the size of the push payment R increases, the size of the red rectangle increases, while the size of the blue triangle stays the same; increasing push payments screens in more firms based on costs, but does not screen in firms based on their success probabilities p_i. When the size of the pull payment x increases, the red rectangle remains where it is, but the line $p_i = \frac{c_i - R}{x}$ capturing the edge of the blue triangle swings outwards. This screens in additional firms based on the atio of their costs to probabilities of success — in particular, those with $c_i \geq R$ and lower ratios $\frac{c_i - R}{p_i}$. These two methods of screening will play a crucial role in our results.

We define $\sigma(R, x)$ as follows:

$$\sigma(R, x) = \int_{E(R,x)} p_i dF(p_i, c_i)$$

σ represents the effective probability of success that emerges from a contract (R, x); it is the expected value of p_i for firms in $E(R, x)$ multiplied by the probability that a firm is in $E(R, x)$ to begin with. $N \cdot \sigma(R, x)$ gives the expected *number* of successes; for example, if $N = 1$, σ is the (unconditional) probability that the designer observes a success.

Next, we make the following assumption for analytical simplicity, as well as to show that our results do not mechanically depend on distributional correlations between p_i and c_i.

Assumption 3. *p_i and c_i are drawn independently from a continuous distribution. That is, there exist cdfs g and h such that $p_i \sim g$ and $c_i \sim h$.*

With the assumption of continuous, independent distributions for p_i and c_i, we can write σ as follows:

$$\sigma(R, x) = \int_0^1 \int_0^{R+p_i x} p_i g(p_i) h(c_i) dc_i dp_i \tag{5}$$

which in turn allows us to characterize the solution to (4) via the following Lagrangian equation:

$$\min_{R,x} R \cdot \int_0^1 g(p_i) H(R + p_i x) dp_i + x \cdot \sigma(R, x) + \lambda(1 - (1 - \theta)^{\frac{1}{N}} - \sigma(R, x)) \tag{6}$$

Finally, note that $\sigma(R, x)$ is bounded above by $\mathbb{E}[p_i]$ (which is its value if all N firms enter). Thus, we have that the probability of at least one success is bounded above by $\overline{\theta} \equiv 1 - (1 - \mathbb{E}[p_i])^N$.

3.2 The Optimal Mix of Push and Pull

We now turn to the first main result of the paper, which is to show that it is always optimal for the designer to utilize some "pull" funding and, under a regularity condition, also utilize some push funding. This suggests that policymakers who may entertain flexible funding models should look to hybrid mechanisms that mix both push and pull forms of funding.

Theorem 1. *Fixing the probability of success $\theta < \overline{\theta}$, the mechanism that minimizes the expected cost always involves a positive amount of pull funding ($x > 0$). If $\frac{H(c_i)}{h(c_i)c_i}$ is strictly increasing in c_i, then the optimal mechanism involves a positive amount of push funding ($R > 0$). If $\frac{H(c_i)}{h(c_i)c_i}$ is weakly decreasing in c_i and $x \geq \overline{c}$, then the optimal mechanism involves no push funding ($R = 0$).*

The proof of the result, as well as all results, is contained in Appendix A, and proceeds via three lemmata. Lemma 1 shows that any pure push contract can be made more cost-efficient by mixing in some pull funding. Lemma 2 shows that any pure pull contract can be made more cost-efficient by mixing in some push funding if the distribution of costs H obeys a certain constraint; sufficient conditions for these constraints are established in the consequent two lemmata.

Pull Improves Pure Push Contracts First, we show how any pure push contract can be improved with a small amount of pull funding. At a high level, utilizing pull funding selects in firms with smaller ratios of costs-to-probabilities of success, i.e. with lower ratios $\frac{c_i}{p_i}$. By contrast, push funding only selects firms with lower costs c_i, and does not inherently select on success probabilities p_i. However, by selecting in on $\frac{c_i}{p_i}$, pull allows firms to collect rents on two dimensions: both costs *and* probabilities of success. Push only allows firms to collect rents on a single dimension: costs c_i.

Suppose d were to achieve a target probability of success θ using a *pure push* mechanism. Pull's ability to screen on p_i generates expected cost-savings for the designer d. Recall that $\sigma = Pr((p_i, c_i) \in E(R, x))\mathbb{E}[P_i|(p_i, c_i) \in E(R, x)]$ represents the expected *firm-level* probability of success, integrating over the firm's probability of entry.

If d added one extra dollar of push funding at R, this would generate a marginal increase in σ of $\mathbb{E}[p_i]$ per dollar. However, since pull funding is only paid out (conditional on entry) with expected probability $\mathbb{E}[p_i]$, and in addition *selects* on higher probabilities of success if d instead added $\mathbb{E}[p_i]$ dollars to R, it would generate a marginal increase in σ of $\mathbb{E}[p_i^2]$. This logic is also borne out in Figure 2, below.

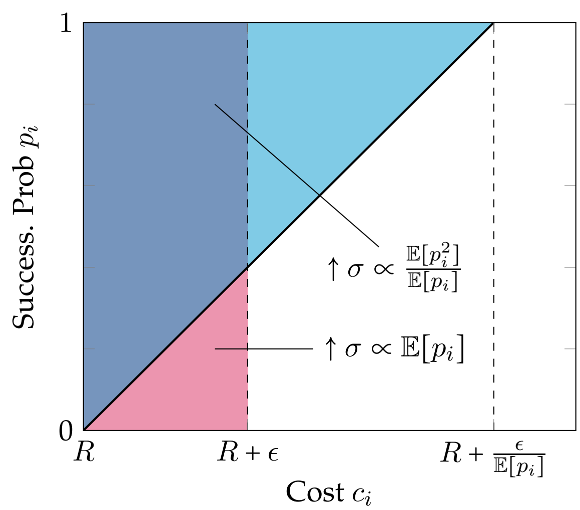

Figure 2: Improving Pull Contract When $x > \overline{c}$

The marginal increase from R to $R+\epsilon$ leads to a marginal increase in the area of the red rectangle by ϵ. The change in σ is reflected by taking the integral reflected by the integral of p_i over the red rectangle, which is $\mathbb{E}[p_i]\epsilon$. For every dollar d spends on push, σ increases by $\mathbb{E}[p_i]$.

However, for every dollar d spends on pull funding, she increases σ by $\frac{\mathbb{E}[p_i^2]}{\mathbb{E}[p_i]}$; moving one dollar to pull generates an increase of $\mathbb{E}[p_i^2]$, but the dollar is only paid out with probability $\mathbb{E}[p_i]$. Thus, by incurring the same expected cost with $x = \frac{\epsilon}{\mathbb{E}[p_i]}$, the integral of σ over the blue triangle is greater than over the red area.

Push Sometimes Improves Pure Pull Contracts When beginning with a pure pull contract, the logic for introducing a positive amount of push is more subtle, and depends on properties of the distribution of costs $H(c_i)$. Suppose d achieved θ with a pure pull contract, i.e. with $x > 0$ and $R = 0$. A marginal increase in x, in this case, leads to an *outwards rotation* in the curve $p_i = \frac{c_i}{x}$. This has the positive effect of including new firms, who tend to have higher probabilities of success than under push. However, this has the negative effect that existing firms are far costlier to include, as they reap more rents. In particular, firms with relatively high costs end up seeing the greatest increases for the designer in costs of inclusion.

This is shown in the left panel of Figure 3. An increase in pull from x to x' rotates the curve $p_i = \frac{c_i}{x}$ outwards to $c_i = \frac{p_i}{x'}$. The red shaded area indicates the new types of firms that are screened in. With this outward rotation, firms with higher costs c_i end up seeing the greatest *increase* in expected costs of inclusion for the designer. This is shown by the distance between the old and new curves, represented by the arrows between $p_i = \frac{c_i}{x}$ and $p_i = \frac{c_i}{x'}$.

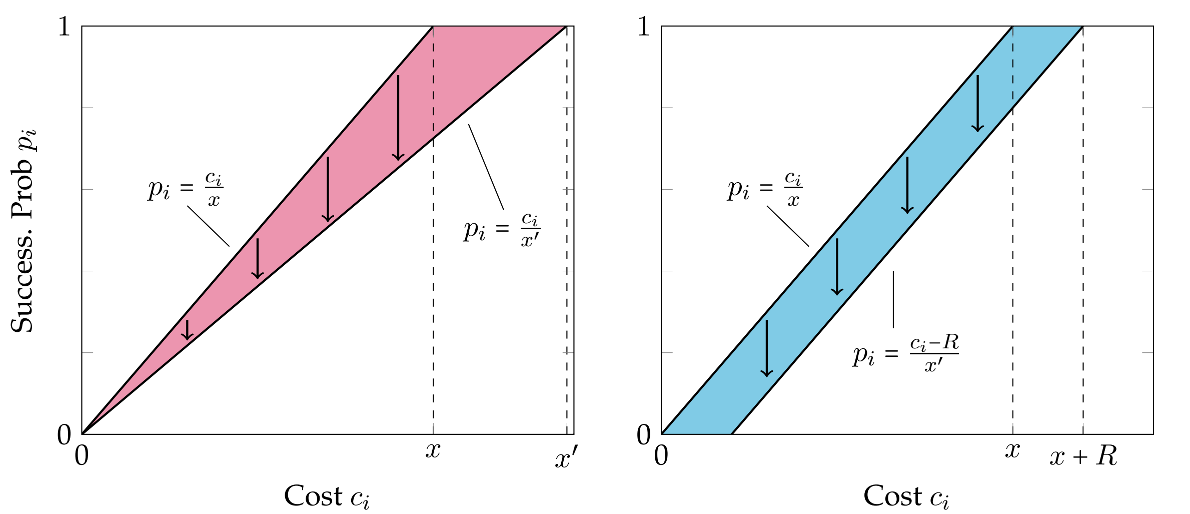

(a) Adding Pull $x' - x$ to Pure Pull Contract x

(b) Adding Push R to Pure Pull Contract x

Figure 3: Effects of Adding Pull and Push Funding to Pure Pull Contract

By contrast, the introduction of a push payment of R leads to an *outwards shift* in the curve $p_i = \frac{c_i}{x}$ to $p_i = \frac{c_i - R}{p_i}$. This is shown in the right panel of Figure 3, where the new firms being screened in are indicated by the blue shaded area. Unlike pull, push is less effective at selecting in firms with relatively high p_is. However, the increase in costs is relatively uniformly distributed across types c_i; the vertical distance between the new curve and old represented by the arrows is constant across firm types.

Thus, whether introducing a marginal amount of push is relatively cheaper than a marginal amount of pull depends on the curvature of costs $H(c_i)$ along the boundary line $c_i = p_i x$. This means that, unlike when beginning with pure push, the difference in the growth of the red and blue curves is not easy to compare. Theorem 1 provides a set of sufficient conditions for when adding in some push can/cannot be cheaper in expectation than a pure pull contract. If $H(c_i)$ is such that $\frac{H(c_i)}{h(c_i)c_i}$ is *strictly increasing* in c_i, then a pure pull contract can be improved upon by introducing a small amount of push funding. For example, the exponential distribution satisfies this property: $H(c_i) = 1 - e^{-\lambda c_i}$, with $\lambda > 0$.

If $H(c_i)$ is such that $\frac{H(c_i)}{h(c_i)c_i}$ is *decreasing* in c_i, then a pure pull contract cannot be improved by adding a small amount of push. For example, any distribution $H(c_i) = \frac{1}{\overline{c}^{\alpha}} c_i^{\alpha}$ with $\alpha > 0$ satisfies this property, including the continuous uniform distribution.

An important caveat to this latter condition is that pure pull can be improved on if $x > \overline{c}$, where $\overline{c}$ is the upper bound of the support of the distribution of costs $H(c_i)$. The intuition is that when $x > \overline{c}$, some pull money is being wasted on screening in a nonexistent mass of firms with $c_i \geq \overline{c}$ and $p_i \geq \frac{c_i}{x}$.

Reducing the amount of pull funding and replacing it with push funding saves on expected costs while leaving σ fixed. This is shown visually in Figure 4. Beginning with a pull amount $x > \overline{c}$ and no push funding, the designer can introduce an improvement by:

- adding in some push funding $R > 0$;
- and reducing the amount of pull funding from x to $\overline{c} - R$.

Intuitively, this change saves some costs by no longer attempting to screen in the nonexistent firms above $\overline{c}$ — shown in blue — thus avoiding overpayment of existing firms benefitting from pull. These cost savings are used to screen in a number of new types via push, indicated in the shaded red area. While the designer loses out on a sliver of firms in the green area immediately to the left of $\overline{c}$, they are made up for by the contribution of the new red area, and σ increases.

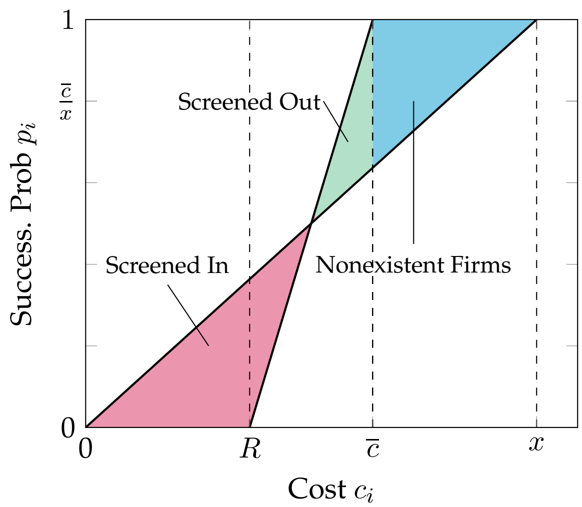

Figure 4: Increase σ From Marginal Increase in Pull Funding, Starting with Pure Push

3.3 Comparing Pure Push and Pure Pull Mechanisms

Next, we answer the following question. If d could only choose between pure push and pure pull mechanisms, which one would have a lower expected cost? This constraint may reflect a realistic set of policies available to a government or funder, who must often rely on simplicity when designing mechanisms. The following result characterizes when pure push is cheaper than pure pull, and vice-versa.

Theorem 2. *Consider a pure pull mechanism that achieves a target probability of succes θ with a transfer $\tilde{x}(\theta)$ and a pure push mechanism that achieves θ with a transfer $\tilde{R}(\theta)$. Pure pull is cheaper in expectation than pure push if and only if $\frac{\tilde{R}(\theta)}{\mathbb{E}[p_i]} \geq \tilde{x}(\theta)$. If $1-(1-\theta)^{\frac{1}{N}}$ is small, pure pull is cheaper in expectation. If $1-(1-\theta)^{\frac{1}{N}}$ is large, pure push is cheaper in expectation.*

The comparison between pure push and pure pull is given by a simple statistic. Let $\tilde{R}$ be the size of the push incentive needed to achieve a target probability θ and $\tilde{x}$ the size of the pull incentive needed to achieve θ. Pure pull is cheaper in expectation if and only if

$$\frac{\tilde{R}(\theta)}{\mathbb{E}[p_i]} \geq \tilde{x}(\theta).$$

This inequality has an intuitive interpretation. $\tilde{x}(\theta)$ is the ratio $\frac{c_i}{p_i}$ of a marginal firm under pure pull. $\tilde{R}(\theta)$ is the value of the cost c_i for a marginal firm under pure push, while $\mathbb{E}[p_i]$ is the *expected* value of p_i for the marginal firm under pure push. Whichever mechanism is cheaper depends on whichever ratio is lower.

Investigating this inequality visually provides a particularly useful characterization, shown in the two panels of Figure 5. Given the push payment $\tilde{R}$, there is a firm with type $(\tilde{p}, \tilde{R})$ such that $\tilde{x} = \frac{\tilde{R}}{\tilde{p}}$. That is, this firm is marginal under both pure *push* and *pull*. If $\tilde{p} \geq \mathbb{E}[p_i]$, pull is cheaper;

otherwise push is cheaper. The left panel of Figure 5 shows an example where pull is cheaper: $\tilde{p}$ is greater than $\mathbb{E}[p_i]$ at the intersection of the vertical red line $c_i = \tilde{R}$ (push) and the blue line $p_i = \frac{c_i}{\tilde{x}}$ (pull).

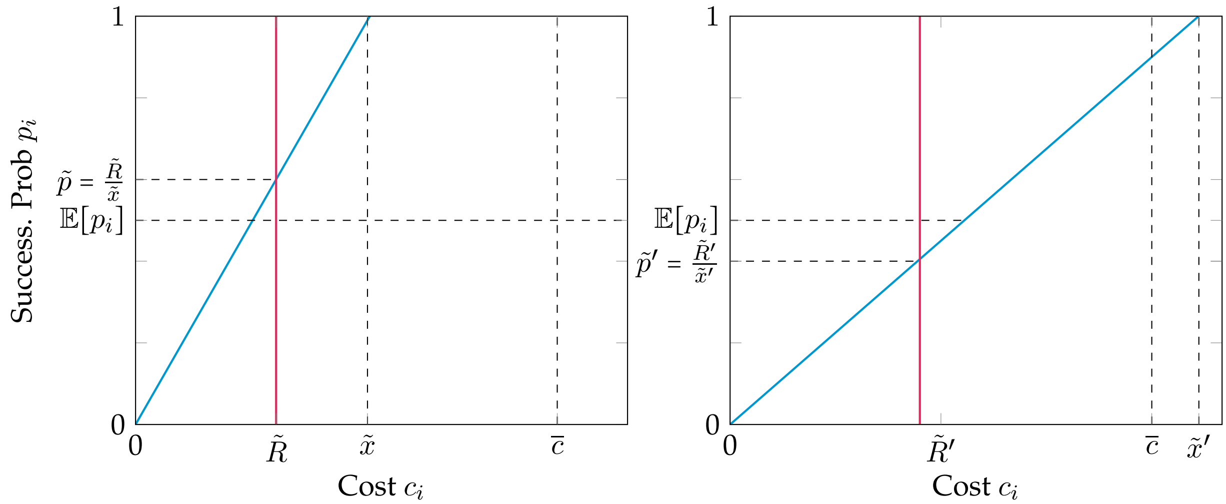

(a) Pure Pull is Cheaper

(b) Adding Push R to Pure Pull Contract x

Figure 5: Effects of Adding Pull and Push Funding to Pure Pull Contract

The right panel shows a different case when the designer must screen in more firms, and where push is cheaper. The intersection of $c_i = \tilde{R}'$ with the line $p_i = \frac{c_i}{\tilde{x}'}$ at $p_i = \tilde{p}'$ is less than $\mathbb{E}[p_i]$. Whichever mechanism is cheaper will depend closely on the distributions of probabilities and costs G and H. However, Theorem 2 shows that for low θ, pure pull will always be cheaper. Te designer can achieve small target probabilities of success by offering a small pull payment and select in firms with relatively high p_is, and the logic is similar to that in the proof of Theorem 1 showing that any pure push contract can be improved with some pull.

For high θ, pure push will always be cheaper. The intuition is as follows. Suppose the designer wanted to screen in all firm types. The cost of doing this with push would be $\bar{c}$ per firm. However, under pure pull, to screen in all firm types, even for very, very large values of x, there is still a small mass of firms (with p_i close to 0) who will have no incentive to participate. In fact, the only way to screen in *all* firms is to send x to $+\infty$. Thus, as θ approaches $1 - (1 - \mathbb{E}[p_i])^N$ — which is the largest value θ can take — the cost of pure pull approaches infinity, while the cost of pure push remains finite, meaning pure push is cheaper.

A corollary of this result is that as the market of firms N grows large, pure pull begins to dominate. The intuition is that for large N, the designer can "depend" on the existence of highly efficient firms, thus shrinking the set of firms she needs to screen in.

3.4 Fundraising Costs

A notable feature of pull funding is that it often requires the designer to *fundraise* or secure budget items well into the future. While the amount that needs to be paid in push funding can be known relatively quickly (firms in the present either participate or do not), fundraisers may have a greater difficulty committing larger funds in the future. How should pull funding be doled out to firms to minimize the amount that the designer has to fundraise?

In particular, recall that pull funding simply delivers an expected utility of x to successful firms. Let P_n be the expected probability that $n \leq N$ firms succeed; note that individual firms as well as the designer have symmetric beliefs about other firms' types, so that these are the same whether looking from firm i's perspective or the designer's. Let X_n be the amount of money in *pull funding* that a designer must pay out, conditional on n firms succeeding. The designer's fundraising cost is $\max_{i \in 1,\ldots,N} X_n$. Conditional on n firms succeeding, firms' expected rewards must be the same. Thus, if n firms succeed, their individual expected payout must be $\frac{X_n}{n}$. Firm i's expected reward *conditional on succeeding* is then given by:

$$\sum_{n=0}^{N-1} \frac{P_n}{n+1} X_{n+1}.$$

The designer's fundraising problem is thus

$$\min_{(X_1,\ldots,X_N)} \max_{i \in 1,\ldots,N} X_i \text{ s.t. } \sum_{n=0}^{N-1} \frac{P_n}{n+1} X_{n+1} = x \tag{7}$$

The solution to this fundraising problem is described in the following proposition

Proposition 1. *Consider a contract (R, x) that generates a probability θ of at least one success. The payout scheme with the lowest fundraising cost for the designer is given by a single pot of money whose size is $\frac{N \cdot \sigma(R,x) \cdot x}{\theta}$, which is paid out in full to all successful firms if at least one firms succeeds. I.e., the solution to (7) is given by $X_n = \frac{N \cdot \sigma(R,x) \cdot x}{\theta}$ for all n.*

The main consequence of the proposition is that, among all methods of doling out the (expected) payment of x to firms, the one that involves fundraising a large pot of money and paying out the entire pot to all successful firms is the cheapest.

Although the formal proof is more involved, the intuition can be gleaned as follows. Suppose that instead of a single shared pot, the designer tailored a prize of size x to all firms. The fundraising cost for this payment method would be $N \cdot x$. However, notice that the cost of a shared pot is strictly lower, i.e.

$$N \cdot x \geq \frac{N \cdot \sigma(R,x) \cdot x}{\theta}$$

since $\sigma \leq 1 - (1-\sigma)^N = \theta$ (strictly for $N > 1$). Intuitively, having firms share a pot of money allows firms to receive payouts far above x if few other firms succeed, which are counterbalanced by payouts below x if many other firms succeed.

3.5 Social Optimum

So far, we have assumed that the designer's target probability of success θ is exogenously given. This section shows that θ can be derived endogenously from a social planner's problem that the designer solves, and characterizes features of the solution at this optimum.

Suppose the development of an innovation has a value V to society. We make the following assumption.

Assumption 4. $V < N \cdot \overline{c}$.

This assumption says that the social value of an innovation will always be less than the cost of attempting to screen in all firms (via push).

The designer d can choose among contracts (R, x) that achieve the development of this innovation with a probability θ. Each of these contracts has its associated expected cost. Thus the designer's problem can be described as

$$\max_{\theta \in [0, \overline{\theta}]} V\theta - C(\theta) \text{ s.t. } C(\theta) = N \cdot \min_{R,x} R \cdot \int_0^1 g(p_i) H(R + p_i x) dp_i + x \cdot \sigma(R, x) \text{ and } \sigma(R, x) = 1 - (1-\theta)^{\frac{1}{N}} \tag{8}$$

Here, $C(\theta)$ represents the expected costs of achieving a target probability of θ. These costs are in turn derived from the minimization problem associated with (6). Finally, recall that $\overline{\theta}$ is the upper bound on the probability of at least one success. The following Theorem lays out the solution to this problem.

Proposition 2. *Suppose θ^* solves the designer's problem in (8). Then it follows that θ^* solves*

$$V = \lambda^* \frac{(1-\theta)^{\frac{1}{N}}}{1-\theta} \tag{9}$$

where λ^* is the Lagrange multiplier in the solution to (6) for $\theta = \theta^*$. Moreover, for every solution to (6) that achieves a target probability of success θ^*, there exists a V such that θ^* is a solution to (8)

The relatively simple form of (13) comes by utilizing an envelope argument. Since $C'(\theta)$ is the solution to the cost minimization problem (for a target probability θ), its variation with θ depends only on how the objective varies with θ, without worrying about variations in x or R due to θ. In this case, this only depends on variations in the constraint $\sigma(R, x) = 1 - (1-\theta)^{\frac{1}{N}}$. Thus, given a social value V, it is possible for the designer to determine what probability of success she should target. This also generates a mild connection to the literature on using market power to screen; settings with higher Vs denote markets that generate greater social surplus (even if that surplus cannot be reached by firms). If a designer must select between projects with different Vs and has information about firms' costs, she can thus compute the solution to (13) and choose the opportunity that delivers the highest expected social value (less costs).

Second, a nearly immediate consequence is that this characterization *microfounds* the problem described in (4): any target probability of success θ can be rationalized by some $V > 0$. Thus, the decisions policymakers who attempt to incentivize firm entry via target probabilities of success can be rationalized as a solution to this planning problem.

4 Discussion

We study the problem of a designer who wants to develop a socially beneficial innovation with a sufficiently high probability θ. To do this, the designer can offer transfers to private sector firms. These may consist of push funding — unconditionally paid ex-ante to cover firms' R&D costs; pull funding — conditionally paid ex-post *only if* a firm develops the innovation; or some combination of both. Firms vary in their costs of research and development and their probabilities of successfully developing the innovation, conditional on taking on those costs. These two variables characterize firms' types, which are private information known only to firms themselves, and not to the designer. We assume that the innovation would not occur in the absence of this funding, since firms would not find it profitable to take on the costs of research and development.

We show that, if the designer is concerned about minimizing expected costs, it is optimal for her to offer at least some pull funding. Under a simple assumption on distributions, we further show that in some cases it is optimal to additionally utilize no push funding or an interior amount of push funding. The intuition is that push and pull generate different trade offs vis a vis screening. Push funding allows firms to collect rents on only a single dimension; they enter if and only if their costs are below the offer of push funding from the designer. Pull funding allows firms to collect rents on two dimensions: their costs and their probabilities of success. In the case of pure pull, they enter if and only if the ratio of their cost to probability of success is below the amount of pull funding the designer offers. While this allows firms to collect rents on *two dimensions*, it also generates a benefit for the designer by selecting in only firms who have sufficiently high likelihoods of success, unlike push. The benefit from pull's selection on firms with high probabilities of success dominates the cost of rents that firms reap for small amounts of pull funding. For large amounts of pull funding, whether the cost of rents dominates depends on the ratio of the moments of the cost distribution.

We also compare two very simple mechanisms: pure push contracts (where firms are only paid with push) and pure pull contracts (where firms are only paid with pull). We show that if the target probability of success is large, pure push mechanisms become cheaper, because screening in larger shares of the market requires arbitrarily large pull payments to incentivize entry. When the target probability of success is small or, relatedly, when the number of firms in the market is large, the designer needs only select in a small set of firm types, meaning she can rely on pull.

We also show that, if the designer is concerned about minimizing the *maximal* payout of pull funding — e.g. because she wants to minimize the amount of money she must fundraise to finance pull transfers — the most efficient way to distribute pull payments is through a shared pot of money, as opposed to individually-tailored lump-sum prizes or other payment rules. Finally, we study a social planning problem for the designer, showing how the designer's pursuit of a target probability of success can be microfounded through the maximization of a social welfare function, where the designer trades off the achievement of some social surplus with the development of an innovation against the costs of achieving that innovation. The result also provides a succinct formula for what a socially-efficient target probability of success should be.

As suggested in our literature review, a key difference between our paper and the remaining literature on incentives for innovation is that the market as-is does not provide useful information to the designer. Put more explicitly, in the absence of transfers from the designer, there is little market for the innovation due to lack of intellectual property rights, hold-up problems, or other issues precluding innovation. Thus, the theoretical framework in this paper is also meant to be a practical guide for policymakers in governments and nonprofits who are looking to use combinations of push and pull funding to incentivize the development of socially beneficial innovations that the private sector has not pursued in the current status quo.

References

Acemoglu, D., & Linn, J. (2004). Market size in innovation: theory and evidence from the pharmaceutical industry. *The Quarterly journal of economics*, *119*(3), 1049–1090.

Ahlvik, L., & van den Bijgaart, I. (2024). Screening green innovation through carbon pricing. *Journal of Environmental Economics and Management*, *124*, 102932.

Blevins, E. G., Angert, J., Gallo, M. E., Harris, L., Kuiken, T., Lindbergh, R., & Sekar, K. (2025, December). *Federal research and development (r&d) funding: Fy2026* (CRS Report No. R48694). Washington, DC: Congressional Research Service, Library of Congress. Retrieved from `https://www.congress.gov/crs_external_products/R/PDF/R48694/R48694.5.pdf`

Bryan, K. A., & Lemus, J. (2017). The direction of innovation. *Journal of Economic Theory*, *172*, 247–272.

Budish, E., Durvasula, M. M., Roin, B. N., & Williams, H. L. (2025). *Missing markets for innovation: Evidence from new uses of existing drugs* (Tech. Rep.). National Bureau of Economic Research.

Carroll, G. (2017). Robustness and separation in multidimensional screening. *Econometrica*, *85*(2), 453–488.

Chari, V. V., Golosov, M., & 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., Iossa, E., & Rey, P. (2021). Prizes versus contracts as incentives for innovation. *The Review of Economic Studies*, *88*(5), 2149–2178.

Deb, R., & Roesler, A.-K. (2024). Multi-dimensional screening: Buyer-optimal learning and informational robustness. *The Review of Economic Studies*, *91*(5), 2744–2770.

Gates, B. (2022, July 13). *By 2026, the gates foundation aims to spend $9 billion a year.* Retrieved from `https://www.gatesnotes.com/work/work-section-topic/reader/commitment-to-the-gates-foundation`

Halac, M., Kartik, N., & Liu, Q. (2017). Contests for experimentation. *Journal of Political Economy*, *125*(5), 1523–1569.

Kremer, M. (2002). Pharmaceuticals and the developing world. *Journal of Economic perspectives*, *16*(4), 67–90.

Kremer, M., Levin, J., & Snyder, C. M. (2020). Advance market commitments: insights from theory and experience. In *Aea papers and proceedings* (Vol. 110, pp. 269–273).

Kremer, M., Levin, J., & Snyder, C. M. (2022). Designing advance market commitments for new vaccines. *Management Science*, *68*(7), 4786–4814.

Rochet, J.-C., & Stole, L. A. (2003). The economics of multidimensional screening. *Econometric Society Monographs*, *35*, 150–197.

U.S. Food and Drug Administration. (2020). *Frequently asked questions on patents and exclusivity.* Web Page. Retrieved from `https://www.fda.gov/drugs/development-approval-process-drugs/frequently-asked-questions-patents-and-exclusivity`

van der Pol, K. H., Aljofan, M., Blin, O., Cornel, J. H., Rongen, G. A., Woestelandt, A.-G., & Spedding, M. (2023). Drug repurposing of generic drugs: challenges and the potential role for government. *Applied health economics and health policy*, *21*(6), 831–840.

Veiga, A., & Weyl, E. G. (2016). Product design in selection markets. *The Quarterly Journal of Economics*, *131*(2), 1007–1056.

Weyl, E. G., & Tirole, J. (2012). Market power screens willingness-to-pay. *The Quarterly Journal of Economics*, *127*(4), 1971–2003.

XPRIZE Foundation. (2004). *Ansari xprize: Launching a new space industry.* `https://www.xprize.org/competitions/ansari`. Retrieved from `https://www.xprize.org/competitions/ansari`

Yang, F. (2025). Costly multidimensional screening. *Review of Economic Studies*, rdaf105.

A Proofs

Theorem 1. *Fixing the probability of success $\theta < \overline{\theta}$, the mechanism that minimizes the expected cost always involves a positive amount of pull funding $(x > 0)$. If $\frac{H(c_i)}{h(c_i)c_i}$ is strictly increasing in c_i, then the optimal mechanism involves a positive amount of push funding $(R > 0)$. If $\frac{H(c_i)}{h(c_i)c_i}$ is weakly decreasing in c_i and $x \geq \overline{c}$, then the optimal mechanism involves no push funding $(R = 0)$.*

Proof. The proof of this result follows from the following three lemmata. □

Lemma 1. *Any pure push contract can be improved upon with some pull. That is, any contract (R, x) with $R > 0$ and $x = 0$ can be made cheaper by a policy with $x > 0$.*

Proof. First, we write the derivative of $\sigma(R, x)$ with respect to R and x:

$$\sigma_R = \frac{d}{dR}\int_0^1 p_i g(p_i)H(R + p_i x)dp_i = \int_0^1 p_i g(p_i)h(R + p_i x)dp_i$$
$$\sigma_x = \frac{d}{dx}\int_0^1 p_i g(p_i)H(R + p_i x)dp_i = \int_0^1 p_i^2 g(p_i)h(R + p_i x)dp_i$$

Next, recall the Lagrangian describing the designer's problem:

$$\min_{R,x} R \cdot \int_0^1 g(p_i)H(R + p_i x)dp_i + x \cdot \sigma(R, x) + \lambda(1 - (1 - \theta)^{\frac{1}{N}} - \sigma(R, x)) \tag{6}$$

The first derivatives of equation (6) with respect to R and x are:

$$[R]: \quad \int_0^1 g(p_i)H(R + p_i x)dp_i + R \cdot \int_0^1 g(p_i)h(R + p_i x)dp_i + x \cdot \sigma_R - \lambda\sigma_R$$
$$[x]: \quad R \cdot \sigma_R + x \cdot \sigma_x + \sigma - \lambda\sigma_x$$

where the σ_R in the second line comes from the fact that

$$\frac{d}{dx}\int_0^1 g(p_i)H(R + p_i x)dp_i = \int_0^1 p_i g(p_i)h(R + p_i x)dp_i$$

Now, suppose by contradiction that a contract with $R > 0$ and $x = 0$ minimized expected costs. Then, it must be the case that

$$\int_0^1 g(p_i)H(R)dp_i + R \cdot \int_0^1 g(p_i)h(R)dp_i - \lambda\sigma_R = 0$$
$$R \cdot \sigma_R + \sigma - \lambda\sigma_x \geq 0$$

I.e., the first order condition with respect to R binds with equality and the first order condition with respect to x is weakly bigger than 0. Solving for the Lagrange multiplier λ in the first line gives

$$\lambda = \frac{\int_0^1 g(p_i)H(R)dp_i + R \cdot \int_0^1 g(p_i)h(R)dp_i}{\sigma_R} = \frac{H(R) + Rh(R)}{\sigma_R}.$$

The derivative with respect to x can then be rearranged as

$$R \cdot \frac{\sigma_R}{\sigma_x} + \frac{\sigma}{\sigma_x} \geq \lambda$$

This implies that

$$\begin{aligned} R \cdot \frac{\sigma_R}{\sigma_x} + \frac{\sigma}{\sigma_x} &\geq \frac{H(R) + Rh(R)}{\sigma_R} \\ R \cdot \frac{\mathbb{E}[p_i]h(R)}{\mathbb{E}[p_i^2]h(R)} + \frac{\mathbb{E}[p_i]H(R)}{\mathbb{E}[p_i^2]h(R)} &\geq \frac{H(R) + Rh(R)}{\mathbb{E}[p_i]h(R)} \\ \frac{\mathbb{E}[p_i]^2}{\mathbb{E}[p_i^2]}\left(R + \frac{H(R)}{h(R)}\right) &\geq R + \frac{H(R)}{h(R)} \\ \frac{\mathbb{E}[p_i]^2}{\mathbb{E}[p_i^2]} &\geq 1 \end{aligned}$$

which is a contradiction, since $\mathbb{E}[p_i^2] > \mathbb{E}[p_i]^2$. Hence, it is the case that the contract $(R, 0)$ that achieves θ can be improved upon by a contract with $x > 0$. □

Lemma 2. *Any pure pull contract can be improved upon with some push only if the expression*

$$\frac{\int_0^1 p_i g(p_i)H(p_i x)dp_i}{\int_0^1 p_i^2 g(p_i)h(p_i x)dp_i} - \frac{\int_0^1 g(p_i)H(p_i x)dp_i}{\int_0^1 p_i g(p_i)h(p_i x)dp_i} \tag{10}$$

is > 0. *That is, any contract* (R, x) *with* $R = 0$ *and* $x > 0$ *can be made cheaper by a policy with* $R > 0$ *only if the equation above holds.*

Proof. Suppose that a contract with $R = 0$ and $x > 0$ minimized expected costs. Utilizing a similar argument as in the previous lemma, we must have that the derivative of the Lagrangian with respect to R at $R = 0$ is ≥ 0, i.e.:

$$\int_0^1 g(p_i)H(p_i x)dp_i + x \cdot \sigma_R - \lambda\sigma_R \geq 0 \implies x + \frac{\int_0^1 g(p_i)H(p_i x)dp_i}{\sigma_R} \geq \lambda$$

Simultaneously, the derivative of the Lagrangian with respect to x must be $= 0$:

$$x \cdot \sigma_x + \sigma - \lambda\sigma_x = 0 \implies x + \frac{\sigma}{\sigma_x} = \lambda$$

Together, these imply that

$$\frac{\int_0^1 g(p_i)H(p_i x)dp_i}{\int_0^1 p_i g(p_i)h(p_i x)dp_i} \geq \frac{\int_0^1 p_i g(p_i)H(p_i x)dp_i}{\int_0^1 p_i^2 g(p_i)h(p_i x)dp_i}$$

in which case $R = 0$ remains an optimum. It fails only if the expression

$$\frac{\int_0^1 p_i g(p_i)H(p_i x)dp_i}{\int_0^1 p_i^2 g(p_i)h(p_i x)dp_i} - \frac{\int_0^1 g(p_i)H(p_i x)dp_i}{\int_0^1 p_i g(p_i)h(p_i x)dp_i} \tag{10}$$

is > 0. □

Lemma 3. *The inequality*

$$\frac{\int_0^1 p_i g(p_i)H(p_i x)dp_i}{\int_0^1 p_i^2 g(p_i)h(p_i x)dp_i} - \frac{\int_0^1 g(p_i)H(p_i x)dp_i}{\int_0^1 p_i g(p_i)h(p_i x)dp_i} \tag{10}$$

is > 0 if $\frac{H(c_i)}{h(c_i)c_i}$ is strictly increasing in c_i and is ≤ 0 if $\frac{H(c_i)}{h(c_i)c_i}$ is weakly decreasing in c_i and $x < \bar{c}$.

Proof. We rewrite the inequality in (10) as

$$\left(\int_0^1 p_i g(p_i)H(p_i x)dp_i\right)\left(\int_0^1 p_i g(p_i)h(p_i x)dp_i\right) > \left(\int_0^1 g(p_i)H(p_i x)dp_i\right)\left(\int_0^1 p_i^2 g(p_i)h(p_i x)dp_i\right)$$

By Fubini's Theorem, this can be rewritten as

$$\begin{aligned}\int_0^1\int_0^1 p_i g(p_i)H(p_i x)q_i g(q_i)h(q_i x)dq_i dp_i > \int_0^1\int_0^1 g(p_i)H(p_i x)q_i^2 g(q_i)h(q_i x)dq_i dp_i \\ \int_0^1\int_0^1 g(p_i)g(q_i)H(p_i x)h(q_i x)[p_i q_i - q_i^2]dq_i dp_i > 0 \\ \int_0^1\int_0^1 g(p_i)g(q_i)H(p_i x)h(q_i x)q_i(p_i - q_i)dq_i dp_i > 0\end{aligned}$$

Dividing by two and switching the order of integration allows us to rewrite the LHS as:

$$\begin{aligned}\frac{1}{2}\int_0^1\int_0^1 [g(p_i)g(q_i)H(p_i x)h(q_i x)q_i(p_i - q_i)] + [g(p_i)g(q_i)H(q_i x)h(p_i x)p_i(q_i - p_i)]dq_i dp_i \\ = \frac{1}{2}\int_0^1\int_0^1 g(p_i)g(q_i)\big[(p_i - q_i)(H(p_i x)h(q_i x)q_i - H(q_i x)h(p_i x)p_i)\big]dq_i dp_i\end{aligned}$$

Next, we analyze the sign of

$$(p_i - q_i)\big[H(p_i x)h(q_i x)q_i - H(q_i x)h(p_i x)p_i\big] \tag{11}$$

which in turns determines the sign of the expression above. First, assume that $R \leq \bar{c}$ Suppose

$p_i > q_i > 0$. (11) is ≥ 0 if and only if

$$\frac{H(p_i x)}{h(p_i x)p_i} \geq \frac{H(q_i x)}{h(q_i x)q_i}.$$

A sufficient condition for this is that $\frac{H(c_i)}{h(c_i)c_i}$ is increasing in c_i. The expression is ≤ 0 if $\frac{H(c_i)}{h(c_i)c_i}$ is *decreasing* in c_i. A similar argument shows that if $q_i > p_i > 0$, an identical logic holds.

Finally, suppose $x \geq \overline{c}$. This means that there exists $\overline{p} < 1$ such that for $p_i > \overline{p}$, we have $H(p_i x) = 1$ and $h(p_i x) = 0$, meaning (11) is always satisfied for $p_i \in [\overline{p}, 1]$. Hence, if $\frac{H(c_i)}{h(c_i)c_i}$ is decreasing in c_i and $x \leq \overline{c}$, (11) fails; but if $x \geq \overline{c}$, this may not necessarily be true.

□

Theorem 2. *Consider a pure pull mechanism that achieves a target probability of succes θ with a transfer $\tilde{x}(\theta)$ and a pure push mechanism that achieves θ with a transfer $\tilde{R}(\theta)$. Pure pull is cheaper in expectation than pure push if and only if $\frac{\tilde{R}(\theta)}{\mathbb{E}[p_i]} \geq \tilde{x}(\theta)$. If $1-(1-\theta)^{\frac{1}{N}}$ is small, pure pull is cheaper in expectation. If $1-(1-\theta)^{\frac{1}{N}}$ is large, pure push is cheaper in expectation.*

Proof. Fixing the target probability θ, consider a pure push mechanism $(\tilde{R}, 0)$ and pure pull mechanism $(0, \tilde{x})$ that achieve the same target probability of success. Pure pull is weakly cheaper than pure push if and only if

$$\tilde{R}\int_0^1 g(p_i)h(\tilde{R})dc_i dp_i = \tilde{R}H(\tilde{R}) \geq \tilde{x}\sigma(0, \tilde{x})$$

Next, define $\tilde{\theta}$ as $\tilde{\theta} = 1-(1-\theta)^{\frac{1}{N}}$, which is $= 0$ when $\theta = 0$ and $= 1$ when $\theta = 1$. Note that $\tilde{R}$ and $\tilde{x}$ are implicit functions of θ, defined as:

$$\sigma(\tilde{R}(\tilde{\theta}), 0) = \int_0^1 p_i g(p_i)H(\tilde{R}(\tilde{\theta}))dp_i = \mathbb{E}[p_i]H(\tilde{R}(\tilde{\theta})) = \tilde{\theta}$$
$$\sigma(0, \tilde{x}(\tilde{\theta})) = \int_0^1 p_i g(p_i)H(p_i\tilde{x}(\tilde{\theta}))dp_i = \tilde{\theta}$$

Putting the pieces together, we then have that pure pull is weakly cheaper if and only if

$$\tilde{R}(\tilde{\theta})\frac{\tilde{\theta}}{\mathbb{E}[p_i]} \geq \tilde{x}(\tilde{\theta})\tilde{\theta} \iff \tilde{R}(\tilde{\theta}) \geq \mathbb{E}[p_i]\tilde{x}(\tilde{\theta})$$

which provides the main condition for comparing the costs of pure push and pure pull. Next, note that

$$\mathbb{E}[p_i]H(\tilde{R}(\tilde{\theta})) = \tilde{\theta} \implies \mathbb{E}[p_i]h(\tilde{R}(\tilde{\theta}))\tilde{R}'(\tilde{\theta}) = 1 \implies \tilde{R}'(\tilde{\theta}) = \frac{1}{\mathbb{E}[p_i]h(\tilde{R}(\tilde{\theta}))}$$
$$\int_0^1 p_i g(p_i)H(p_i\tilde{x}(\tilde{\theta}))dp_i = \tilde{\theta} \implies \int_0^1 p_i^2 g(p_i)h(p_i\tilde{x}(\tilde{\theta}))\tilde{x}'(\tilde{\theta})dp_i = 1 \implies \tilde{x}'(\tilde{\theta}) = \frac{1}{\int_0^1 p_i^2 g(p_i)h(p_i\tilde{x}(\tilde{\theta}))dp_i}$$

Since, at $\tilde{\theta} = 0$, $\tilde{x} = \tilde{R} = 0$, we have that

$$\begin{aligned}\tilde{R}'(\tilde{\theta}) \geq \mathbb{E}[p_i]\tilde{x}'(\tilde{\theta}) \iff & \frac{1}{\mathbb{E}[p_i]h(\tilde{R}(\tilde{\theta}))} \geq \frac{\mathbb{E}[p_i]}{\int_0^1 p_i^2 g(p_i)h(p_i\tilde{x}(\tilde{\theta}))dp_i} \\ \iff & \frac{1}{\mathbb{E}[p_i]h(0)} \geq \frac{\mathbb{E}[p_i]}{\int_0^1 p_i^2 g(p_i)h(0)dp_i} \\ \iff & \mathbb{E}[p_i^2] \geq \mathbb{E}[p_i]^2,\end{aligned}$$

which always holds. Thus, for θ small, we have that pure pull is always weakly cheaper than pure push.

Next, consider $\tilde{\theta} \to \mathbb{E}[p_i]$, which characterizes the upper bound for $\tilde{\theta}$. The expected cost of push to achieve this target probability is $\tilde{C}$ — which screens in all N firms (i.e. $H(\tilde{C}) = 1$). The expected cost of pull must then solve

$$\int_0^1 p_i g(p_i)H(p_i x)dp_i = \mathbb{E}[p_i],$$

but note that for all $x > 0$, there exists $p_i \in (0,1)$ such that $H(p_i x) < 1$; i.e., we always have

$$\int_0^1 p_i g(p_i)H(p_i x)dp_i < \int_0^1 p_i g(p_i)H(1)dp_i = \mathbb{E}[p_i]$$

hence, as $\tilde{\theta} \to \mathbb{E}[p_i]$, the cost of pull approaches $+\infty$, while the cost of push approaches $\tilde{C}$. This means that, for $\tilde{\theta}$ large (i.e., for $\tilde{\theta}$ in a neighborhood of $\mathbb{E}[p_i]$, pure push is cheaper in expectation than pure pull. Inverting the normalization $\tilde{\theta} = 1-(1-\theta)^{\frac{1}{N}}$ provides the statement in the proof. □

Proposition 3. *Consider a contract (R,x) that generates a probability θ of at least one success. The payout scheme with the lowest fundraising cost for the designer is given by a single pot of money whose size is $\frac{N\cdot\sigma(R,x)\cdot x}{\theta}$, which is paid out in full to all successful firms if at least one firms succeeds. I.e., the solution to (7) is given by $X_n = \frac{N\cdot\sigma(R,x)\cdot x}{\theta}$ for all n.*

Proof. With slight abuse of notation, note that the probability that n other firms succeed given a contract (R,x) is:

$$\binom{N-1}{n}\sigma^n(1-\sigma)^{N-1-n}$$

where $\sigma = \sigma(R, x)$. This means that the left hand side of the constraint in (7) can be written as:

$$\sum_{n=0}^{N-1} \frac{P_n}{n+1} X_{n+1} = \sum_{n=0}^{N-1} \frac{1}{n+1} \binom{N-1}{n} \sigma^n (1-\sigma)^{N-1-n} X_{n+1}$$
$$= \sum_{n=0}^{N-1} \frac{1}{n+1} \frac{(N-1)!}{(N-1-n)!n!} \sigma^n (1-\sigma)^{N-1-n} X_{n+1}$$
$$= \sum_{n=0}^{N-1} \frac{1}{N \cdot \sigma} \frac{N(N-1)!}{(N-(n+1)!)!(n+1)!} \sigma^{n+1} (1-\sigma)^{N-(n+1)} X_{n+1}$$
$$= \frac{1}{N \cdot \sigma} \sum_{n=0}^{N-1} \frac{N!}{(N!)!(n+1)!} \sigma^{n+1} (1-\sigma)^{N-(n+1)} X_{n+1}$$
$$= \frac{1}{N \cdot \sigma} \sum_{n=1}^{N} \binom{N}{n} \sigma^n (1-\sigma)^{N-n} X_n,$$

Define $\beta_n = \binom{N}{n} \sigma^n (1-\sigma)^{N-n}$. Define $\alpha_i = \frac{\beta_i}{\sum_{i=1}^N \beta_i}$. Note that $\sum_{i=1}^N \alpha_i = 1$. Moreover, we have that $\sum_{i=1}^N \beta_i = 1 - (1-\sigma)^N = \theta$, i.e. the probability of at least one success. Using this, we rewrite the problem as:

$$\min_{X_1,\ldots,X_n} \max_{i \in 1,\ldots,N} X_i \text{ s.t. } \sum_{n=1}^{N} \alpha_n X_n = \frac{N \cdot \sigma \cdot x}{\theta} \tag{12}$$

i.e. it is minimizing the maximum value of X_i subject to the linear constraint $\sum_{n=1}^N \beta_n X_n = N \cdot \sigma \cdot x$.

Let X_m be the solution to (12). By construction, we must have

$$X_m = \sum_{i=1}^{N} \alpha_n X_m \geq \sum_{i=1}^{N} \alpha_n X_n = \frac{N \cdot \sigma \cdot x}{\theta}$$

i.e. $\frac{N \cdot \sigma \cdot x}{\theta}$ is a lower bound for X_m. But note that this lower bound can be achieved with equality by simply setting $X_m = X_i$ for all i. Thus, we have $X_n = \frac{N \cdot \sigma \cdot x}{\theta}$ for all n.

□

Proposition 4. *Suppose θ^* solves the designer's problem in (8). Then it follows that θ^* solves*

$$V = \lambda^* \frac{(1-\theta)^{\frac{1}{N}}}{1-\theta} \tag{13}$$

where λ^* is the Lagrange multiplier in the solution to (6) for $\theta = \theta^*$. Moreover, for every solution to (6) that achieves a target probability of success θ^*, there exists a V such that θ^* is a solution to (8)

Proof. Recall that our problem was given by:

$$\max_{\theta \in [0, \overline{\theta}]} V\theta - C(\theta) \text{ s.t. } C(\theta) = N \cdot \min_{R,x} R \cdot \int_0^1 g(p_i) H(R + p_i x) dp_i + x \cdot \sigma(R, x) \text{ and } \sigma(R, x) = 1 - (1-\theta)^{\frac{1}{N}} \tag{8}$$

The first order condition for this problem is given by $V = C'(\theta)$. Note that:

$$C(\theta) = N \cdot \min_{R,x} R \cdot \int_0^1 g(p_i) H(R + p_i x) dp_i + x \cdot \sigma(R, x) + \lambda(1 - (1-\theta)^{\frac{1}{N}} - \sigma(R, x))$$

By the envelope theorem we have that

$$C'(\theta) = \lambda^* \frac{(1-\theta)^{\frac{1}{N}}}{1-\theta}$$

where λ^* is the Lagrange Multiplier for the minimization problem. The value of λ^* falls into two cases. If

$$\frac{\int_0^1 p_i g(p_i) H(p_i x) dp_i}{\int_0^1 p_i^2 g(p_i) h(p_i x) dp_i} - \frac{\int_0^1 g(p_i) H(p_i x) dp_i}{\int_0^1 p_i g(p_i) h(p_i x) dp_i} \geq 0 \tag{10}$$

then $\lambda^* = \lambda$ is the solution to

$$\lambda = x + \frac{\int_0^1 g(p_i) H(p_i x) dp_i}{\int_0^1 p_i g(p_i) h(p_i x) dp_i} \text{ and } \int_0^1 g(p_i) H(p_i x) dp_i = 1 - (1-\theta)^{\frac{1}{N}}.$$

x is given implictly as a function of θ by the latter equation, which is in turn plugged into the former equation. If (10) is < 0, then $\lambda^* = \lambda$ is the solution to

$$\lambda = x + \frac{\int_0^1 g(p_i) H(R + p_i x) dp_i}{\int_0^1 p_i g(p_i) h(R + p_i x) dp_i} + R \cdot \frac{\int_0^1 g(p_i) h(R + p_i x) dp_i}{\int_0^1 p_i g(p_i) h(R + p_i x) dp_i}$$

$$\lambda = x + \frac{\int_0^1 p_i g(p_i) H(R + p_i x) dp_i}{\int_0^1 p_i^2 g(p_i) h(R + p_i x) dp_i} + R \cdot \frac{\int_0^1 p_i g(p_i) h(R + p_i x) dp_i}{\int_0^1 p_i^2 g(p_i) h(R + p_i x) dp_i}$$

$$\text{and } \int_0^1 p_i g(p_i) H(R + p_i x) dp_i = 1 - (1-\theta)^{\frac{1}{N}}.$$

where R and x are similarly determined endogenously. In both cases, λ is bounded; as $\theta \to 1$, $C'(\theta) \to \infty$; as $\theta \to 0$, $C'(\theta) \to \lambda^* \cdot 1 \to 0$.

Finally, note that as V goes to 0, θ goes to 0. As V goes to $+\infty$, since λ is bounded, there exists a value of V such that $\theta = \overline{\theta}$. This gives the latter part of the result. □