Innovation Contests with Distinct Approaches*

Simon Block

University of Bonn

simon.block@uni-bonn.de

February 28, 2023

Abstract

This paper introduces a model of innovation contests where agents can produce the same innovation via distinct research approaches. Approaches differ in their overall viability based on their costs and success probabilities, and in the timing of costs and possible successes. There are two reasons why inefficient equilibrium behavior may arise when rewards are paid independent of the used approach: over-investment in approaches where successes are correlated; and research efforts on fast approaches crowding out efforts on slower but more effective approaches. Further, it is shown that a greedy algorithm determines the first best if costs are equal between all approaches. In a static setting, approach-specific rewards can uniquely implement the social optimum and extract the entire surplus. In a dynamic setting, approach-specific rewards can be augmented with an "efficient-stopping condition" to achieve the same result, provided the benefit of a successful innovation and the number of agents are both large enough.

JEL codes: D81, O31, O38

*I thank Kailin Chen, Francesc Dilmé, Marina Halac, Paul Heidhues, Thomas Kohler, Stephan Lauermann, Benny Moldovanu, Axel Niemeyer, Georg Nöldeke, Tobias Rachidi, Sven Rady, Anna Schulze Tilling, Johanna Tassot, Alexander Winter, and workshop participants in Bonn and Cologne for their helpful comments and suggestions. I am grateful to my advisors Stephan Lauermann and Sven Rady for their continuous support and guidance. I am thankful for the hospitality of the Yale Economics department, where parts of this paper were written. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany´s Excellence Strategy – EXC 2126/1– 390838866.

1 Introduction

In March 2020, when the World Health Organization declared COVID-19 to be a global pandemic, the race for a safe and effective vaccine against the novel Corona-virus was already well underway. After the genetic sequence of SARS-CoV-2 was published on January 11$^{\text{th}}$ 2020, multiple vaccine prototypes were developed within a short period of time. In addition to established approaches to vaccine development, such as using an inactivated or live-attenuated version of the virus itself, there were novel approaches, such as the mRNA technology used by BioNTec/Pfizer and Moderna, and the viral-vector technology used by AstraZeneca, that were considered very promising.[1] Perhaps most importantly, these novel approaches promised *fast* development: the genetic code of Moderna's mRNA vaccine was designed within a single day, and Phase 3 clinical trials were successfully completed by both BioNTech/Pfizer and Moderna in November 2020, before the fastest live attenuated COVID vaccine could even start Phase 1 clinical trials in December 2020.[2]

Vaccine development is not the only environment where there are distinct approaches to creating a specific innovation. Another important setting with clearly distinguishable approaches is the development of carbon capture technologies: Pires et al. (2011) assert that "CO_2 capture can be performed following three different technological concepts: post-combustion capture systems, pre-combustion capture systems, and oxy-fuel capture systems."

Even if the potential end product of all approaches has (approximately) the same functionality, there are (at least) three other dimensions in which approaches may be different from one another: *viability*, *correlation*, and *timing*. First, some approaches may be more viable than others overall, because they have a higher probability of success, or because they have lower costs. Second, the correlation of successes within an approach can be different between approaches. For example, approaches that depend on applying novel technologies may have a higher idiosyncratic risk of failure than those that make use of long-established methods. Third, there can be differences in the distribution of possible successes and costs across time.

In the wake of the COVID-19 pandemic, there were multiple calls for a large cash prize for the company that is the first to develop a safe and effective vaccine.[3] Such an approach-independent contest would have stood in the tradition of famous innovation contests like the British Longitude Act of 1714, the Netflix Prize, or the Google Lunar X-Prize. However, in

[1] Le et al. (2020) categorize the COVID-19 vaccines that were in development in April 2020 into a total of nine different approaches.

[2] Polack et al. 2020; Baden et al. 2021; Wang et al. 2021.

[3] For example by Hemel and Ouellette (2020), and Callaghan (2020).

a setting where agents must choose between distinct research approaches, it is questionable whether such an approach-independent contest is the ideal method. This paper identifies potential inefficiencies induced by such approach-independent contests.

Moreover, this paper addresses the portfolio choice problem of efficiently allocating research efforts to different approaches depending on the aforementioned three aspects of viability, correlation, and timing. Finally, I study how a principal can utilize innovation contests to steer the behavior of a set of agents who choose between different research approaches.

In the model, following Halac, Kartik, and Liu (2017), a principal and a set of identical agents are uncertain about the true state of the world, and the principal has a fixed valuation for obtaining the innovation, but there is no additional benefit if the innovation is produced more than once. There are two periods[4] and in each period the agents can choose whether to exert costly effort to produce the innovation with a probability that depends on the state of the world. The novel feature of my model is that agents do not just choose whether to exert effort, they also chose from a set of approaches where to spend their effort. The costs of effort and the probability of success can vary both between approaches and between periods. In particular, some research approaches may be faster than others, in the sense that both costs and instantaneous probability of success are higher in the first period. It is assumed that all actions and successes are publicly observed.

In this setting, a contest is defined to be a reward rule that only rewards agents if they succeed in producing the innovation, and that does not discriminate against agents based on their identity. Nevertheless, a contest may discriminate between agents based on the approach that leads them to success.

This paper identifies two separate reasons why approach-independent contests may induce inefficient equilibrium behavior. The first reason is that too many agents may choose to follow approaches with high overall viability. If successes within these approaches are sufficiently correlated, efficiency would require agents to spread out across more approaches instead, even if some of them are less attractive. The second reason is due to differences in the timing of successes. If some approaches are inherently faster than other approaches, then agents following those fast approaches may crowd out agents on slower approaches. This can happen even if the slower approaches are much more efficient. In some instances, due to this "crowding-out effect", increasing the total prize of an approach-independent contest even induces a *decrease* in the total probability of obtaining a success.

To find the socially optimal allocation, a simple greedy algorithm—resembling the

[4] In Halac, Kartik, and Liu (2017) The motivating example works with two periods, whereas the main model is formulated in continuous time.

marginal improvement algorithm as introduced by Chade and Smith (2006)—can be used to identify the socially optimal allocation of agents to approaches under some additional assumptions on the nature of approaches. The greedy algorithm sequentially assigns agents to the approach where their marginal contribution to the expected social surplus is the largest, given the behavior of the already assigned agents. It stops when it is no longer possible to assign an agent with a positive marginal contribution, or all agents are assigned. In the static case where there is only a single period, a sufficient condition for the greedy algorithm to work is that all approaches have the same costs. This condition would, for example, be fulfilled if costs are caused by a standardized testing procedure, like clinical trials for a vaccine. In the case with two periods, in addition to an equal costs assumption, an additional assumption that allows splitting the set of approaches into "fast" and "slow" approaches is sufficient to show that an iterated version of the greedy algorithm can identify the social optimum.

I show that in a static setting, the principal can implement any desired equilibrium behavior as the unique Nash equilibrium and extract almost the entire social surplus at the same time. To achieve this, the principal can use an approach-specific contest: Agents that follow the same approach compete for a fixed reward, but no agent is affected by the actions and successes of other agents that follow different approaches. Because the marginal contribution of every additional agent following the same approach is diminishing in the number of agents, any fixed prize always corresponds to a certain number of agents following an approach. If less than this number of agents follow the approach, it is profitable for an additional agent to join, and if more agents were to follow the approach, their expected payoff would be negative.

With multiple periods, early successes may lead to efficiency gains because they allow all agents to stop incurring costs. But simply augmenting an approach-specific contest with an "efficient-stopping condition", which eliminates rewards for successes that occur after the first success, does in general not allow the principal to harvest these gains in efficiency. The reason is that such a condition makes the expected payoffs of agents following different approaches interdependent. However, if the principal's valuation and the number of agents are both large, an approach-specific contest with efficient-stopping condition can implement the social optimum and extract the full surplus.

The rest of this paper is structured as follows. In Section 2, I introduce the baseline model and briefly discuss the most important assumptions. Section 3 analyzes the static case where there is only a single relevant period, whereas Section 4 analyzes the dynamic aspects of the model. Both sections follow approximately the same structure: First, I illustrate inefficiencies caused by approach-independent contests, then I show how a greedy algorithm can be used to identify the social optimum under certain conditions, and then I address the question of implementation. Finally, Section 5 concludes.

Related Literature

There is a vast theoretical literature on contest design; Fu and Wu (2019) provide a good overview. Part of this literature analyzes research contests. In a research contest, agents compete in terms of the quality of the innovation they produce, and the contest designer aims to obtain a high quality. Taylor (1995), and Che and Gale (2003) are exemplary for this branch of literature.

This paper falls into the closely related branch of literature that studies innovation contests. These are contests where agents compete to achieve a pre-specified quality before their competitors do. Often, innovation contests are studied in a dynamic setting where agents learn about an underlying state of the world. For example, Choi (1991), and Malueg and Tsutsui (1997) study R&D races where agents learn about the "hazard rate" that governs the arrival of innovations. Halac, Kartik, and Liu (2017) expand on the literature on innovation contests with learning, by studying contest design where the principal can jointly vary the prize-sharing scheme and information-disclosure policy.

The environment of this paper has been inspired by Halac, Kartik, and Liu (2017). Therefore, there are many similarities. Most notably, there is an underlying state of the world, unknown to the principal and the agents, that determines whether innovation is feasible. In contrast to Halac, Kartik, and Liu (2017), I assume that all actions and successes are publicly observable. Hence the principal cannot make use of different information-disclosure policies.

Instead, I introduce distinct approaches and endow the principal with the ability to discriminate between agents based on the approach they use. There are several other papers where agents may choose between different research projects or approaches. Acemoglu, Bimpikis, and Ozdaglar (2011) study a setting where firms choose between implementing one of multiple uncertain research projects. Alternatively, firms can wait and copy the innovation of other firms that have successfully completed a research project. Even though the environment is quite similar to the one I study, their findings are mainly about (preventing) free-riding due to firms being able to copy the innovation produced by others.

Akcigit and Liu (2016) study the case where firms may switch between safe and risky research projects. In their model actions and outcomes are private information, and firms may gain an advantage by hiding information about dead-ends they discover. Letina and Schmutzler (2019) study a setting with a continuum of research approaches, where the optimal approach is ex-ante unknown. After uncertainty dissolves, the quality of an approach depends on the distance to the optimal approach. In their setup, the principal tries to induce variety in the approaches that agents choose because it generates an option value. Also, Letina (2016) studies a related setting, where firms simultaneously choose a set of research projects from a continuum and there is only a single approach that leads to innovation.

2 Model

2.1 Model

There is a principal who wants to obtain a specific innovation. In addition, there are $n \in \mathbb{N}$ identical, risk-neutral agents, with $\mathcal{N} := \{1, ..., n\}$ denoting the set of all agents. Only the agents can produce the innovation by following one of $K \in \mathbb{N}$ distinct *approaches* $a_1, ..., a_K$, with $\mathcal{A} \coloneqq \{a_1, ..., a_K\}$ denoting the set of all approaches.

A persistent, unobservable state of the world $(\theta_{a_1}, ..., \theta_{a_k}) \in \{G, B\}^K$ determines the *feasibility* of each approach.[5] If $\theta_a = G$, approach $a \in \mathcal{A}$ is *feasible*, else approach a is *not feasible*. The variables $\theta_{a_1}, ..., \theta_{a_k}$ are independently distributed with

$$P_a \coloneqq \mathbb{P}(\theta_a = G) \in (0, 1] \qquad \forall a \in \mathcal{A}.$$

There are two time periods, $t = 1$ and $t = 2$. An approach $a \in \mathcal{A}$ is characterized by its prior probability of being feasible, P_a, together with two ordered pairs

$$[(\lambda_{a,1}, c_{a,1}), (\lambda_{a,2}, c_{a,2})],$$

which specify how the costs of following approach a, and the probability of successfully producing the innovation when following approach a are distributed across the two periods. More specifically, in each period $t = 1, 2$ in which an agent follows a, this agent incurs cost $c_{a,t} \in \mathbb{R}_+$, and conditional on $\theta_a = G$ succeeds with probability $\lambda_{a,t} \in [0, 1]$. If $\theta_a = B$, the conditional probability of success of all agents following a is 0 in both periods. Moreover, conditional on the state of the world, successes are independent across agents.

The action of an agent i in period $t = 1, 2$ is denoted by $\alpha_{i,t}$. In period 1, each agent i may choose to either *follow one approach* $a \in \mathcal{A}$, denoted by $\alpha_{i,1} = a$, or to *abstain*, in which case $\alpha_{i,1} = \alpha_{i,2} \coloneqq \emptyset$. In period 2, each agent i that followed an approach in the first period may choose to *continue following* $\alpha_{i,1}$, in which case $\alpha_{i,2} = \alpha_{i,1}$, or to *quit* in which

[5] Here G stands for *good* and B stands for *bad*.

case $\alpha_{i,2} = \emptyset$. Abstaining and quitting are costless actions. Agents cannot switch between approaches or re-enter after abstaining in the first period.

An *action profile at time* $t = 1, 2$, denoted by α_t, is a list of all actions at t, i.e.,

$$\alpha_t = (\alpha_{1,t}, ..., \alpha_{n,t}).$$

Whether an agent successfully produced an innovation in period $t = 1, 2$ is encoded by $s_t = (s_{1,t}, ..., s_{n,t})$, where $s_{i,t} = 1$ means that agent i succeeded and $s_{i,t} = 0$ means that she did not. All actions and all successes are publicly observable. In the first period, the principal and the agents share a common prior belief $P = (P_{a_1}, ..., P_{a_K})$ about the state of the world. In the second period, they share a public belief $\mu = (\mu_{a_1}, ..., \mu_{a_K})$.[6]

Success is only possible when an approach is feasible. Thus, whenever a success on approach a is observed, it holds that $\mu_a = 1$. In the absence of any successes on approach a in the first period, Bayesian updating demands that

$$\mu_a = \frac{P_a(1-\lambda_{a,1})^{n_a}}{P_a(1-\lambda_{a,1})^{n_a(1)} + 1 - P_a},$$

where n_a is the number of agents that followed approach a in period 1.

Agents do not directly benefit from successes. For a given terminal history $h^T = (\alpha_1, s_1, \alpha_1, s_2)$, the payoff of an agent i is equal to the *reward* $w_i(h^T)$ that she receives from the principal, minus the costs she incurred. The principal receives a payoff of $v \in \mathbb{R}_+$ if at least one agent succeeds, minus the sum of the rewards paid to the agents. The principal is not budget constrained.

Denote the set of terminal histories by H^T. Attention is restricted to reward functions $(w_i : H^T \to \mathbb{R}_+)_{i \in \mathcal{N}}$ called *contests*. A contest has two properties: it is anonymous, and it only rewards agents that actually succeed. Formally, a contest must fulfill: (i) $w_i(h^T) = w_j(\tau_{ij}(h^T))\ \forall i, j \in \mathcal{N}$, where $\tau_{ij} : \mathcal{H}^T \to \mathcal{H}^T$ is the permutation that switches $\alpha_{i,t}$ with $\alpha_{j,t}$ and $s_{i,t}$ with $s_{j,t}$ for $t = 1, 2$; and (ii) $s_{i,t} = 0\ \forall t \implies w_i(h^T) = 0$.

The solution concept is Nash equilibrium and the analysis focuses on equilibria in pure strategies.

2.2 Discussion of Assumptions

The presence of multiple approaches is the defining feature of this model. By allowing the costs and conditional probabilities of success to vary over time, it becomes possible to

[6] Since successes are independent conditional on the state of the world, the public belief can be written as a list of beliefs about the state of the world for every approach.

study the effects of their timing being different for different approaches. For example, it is interesting to compare "fast" and "slow" approaches.

The assumption that agents cannot switch between approaches and that abstaining and quitting are irreversible helps limit the action space in the second period. This allows focusing the attention on the coordination in the first period. Moreover, since agents cannot switch between approaches, the success probabilities and costs for period 2 can be interpreted as depending on the cumulative effort that agents exert. Similarly, the assumption that every action and success is observable by all agents and the principal ensures that there is a public belief that all share. Taken together, these assumptions give the principal a lot of power.

The limitations to the principal's power are built into the definition of a contest. The principal is unable to discriminate against agents based on their identity. Additionally, the principal can only reward agents that are successful, which means that the principal cannot directly compensate agents for their efforts. This indirectly captures the moral-hazard problem the principal would face if effort were not observable.

3 Static Case: One Period

In this section, the model from Section 2 is restricted to a static setting, by considering the case where the idiosyncratic chance of success $\lambda_{a,2}$ equals 0 for all approaches a. Since there is only a single relevant period, the timing of costs and possible successes does not come into play. The central aspect is the *correlation* between successes of agents following the same approach.

Recall that an approach a is characterized by its probability of being feasible P_a, its cost $c_{a,1}$, and the probability of success conditional on being feasible $\lambda_{a,1}$. Since there is only one relevant period, in this section the index referring to the period is omitted. That means c_a is used instead of $c_{a,1}$, and λ_a instead of $\lambda_{a,1}$. Also the action of an agent i is simply denoted by α_i, an action profile by $\alpha = (\alpha_1, ..., \alpha_n)$, and the successes and failures of the agents are described by $s = (s_1, ..., s_n)$.

Since every agent acts only once, simultaneously with all other agents, a pure strategy of an agent i is simply the action α_i which i chooses.

3.1 Correlation Representation

Instead of describing the distribution of successes by P_a and λ_a, it can be equally well described by a pair (ϕ_a, ρ_a), where $\phi_a \coloneqq P_a\lambda_a$ is the *unconditional probability of success*, and ρ_a is the *correlation coefficient between the successes of any two agents following* a. Suppose two agents follow approach a, with $0 < P_a\lambda_a < 1$, then the correlation coefficient of their successes is

$$\rho_a = \frac{P_a\lambda_a^2 - (P_a\lambda_a)^2}{P_a\lambda_a(1 - P_a\lambda_a)} = \frac{\lambda_a - P_a\lambda_a}{1 - P_a\lambda_a}.$$

Hence, except for the cases of certain success ($P_a\lambda_a = 1$) or certain failure ($P_a\lambda_a = 0$), where the correlation coefficient is not defined because the variance is 0, the parameters ϕ_a and ρ_a are uniquely defined for any pair P_a and λ_a.[7] It follows that the model allows for all levels of positive correlation between successes on a given approach, no matter the unconditional probability of success. The extreme case of perfect correlation occurs if $\lambda_a = 1$, and successes are uncorrelated if $P_a = 1$.

3.2 Approach-Independent Contests: Inefficiencies

There are two reasons why it is interesting to study contests with rewards that are independent of the approach used by a successful agent. First, they are the simplest and most common way of designing an actual innovation contest. Second, they can be thought of as representing market incentives. Since all successful innovations have the same functionality, it seems plausible that each successful agent can capture an equal share of the market, independently of the approach that they used. The following example illustrates how inefficiencies may arise when there are multiple approaches. Such inefficiencies only arise when the correlation of successes of agents following the same approach is relatively high.

Example 1. *There are two agents,* 1 *and* 2*, and two approaches,* A *and* B*. Both approaches have the same costs* $c = c_A = c_B$*, but* A *has a higher unconditional probability of success, that is,* $P_A\lambda_A > P_B\lambda_B$*. There is a fixed reward* $\bar{w}$*, with* $\frac{2c}{P_B\lambda_B} < \bar{w} < v$[8]*, that is shared equally between all successful agents, irrespective of the approach that leads them to success.*

In this setting, a single agent would always prefer to follow A. Consider the case where agent 1 follows A, what is agent 2's best reply? Intuitively that depends on two things: first, on how much larger $P_A\lambda_A$ is compared to $P_B\lambda_B$; second, if the difference between $P_A\lambda_A$ and $P_B\lambda_B$ is not too large, on the level of correlation between successes on A.

[7] Conversely, given any $\phi_a > 0$ and ρ_a, it holds that $\lambda_a = 1 - (1 - \rho_a)(1 - \phi_a) = \phi_a + \rho_a(1 - \phi_a)$, and $P_a = \frac{\phi_a}{\lambda_a} = \frac{\phi_a}{\phi_a + \rho_a(1-\phi_a)}$ are uniquely determined.

[8] This condition simply ensures that agents always prefer following an approach to abstaining from the contest and that this is also in the interest of the principal.

If agent 2 succeeds with either approach, she is certain to receive at least half of the reward $\bar{w}$. In addition, if agent 1 fails, she receives the entire reward. If agent 2 follows A and succeeds, the probability that she is the only one to succeed is

$$1 - \lambda_A = (1 - P_A\lambda_A)(1 - \rho_A).$$

In contrast, conditional on succeeding on approach B, the probability that agent 2 is the only one to succeed is $1 - P_A\lambda_A$. For $P_A < 1$, the probability of receiving the entire reward is larger when succeeding with approach B because successes on A are correlated, while between the two different approaches they are independent.

If there is no correlation on A, then clearly agent 2's best reply is to also follow A. This is not an inefficiency. If successes on the most promising approach are uncorrelated, then it is in the interest of efficiency that all agents that follow an approach take the most promising one. In contrast even if successes on A are perfectly correlated, agent 2 will still strictly prefer $\alpha_2 = A$ as long as

$$\frac{P_A\lambda_A}{2} > \frac{P_B\lambda_B}{1 + P_B\lambda_B}.$$

This is a clear inefficiency. If successes on an approach are perfectly correlated, at most one agent should follow it because the marginal social benefit (MSB) of an additional agent is 0. Solving for ρ_A yields that agent 2 strictly prefers A to B if and only if

$$\rho_A < \frac{P_A\lambda_A - P_B\lambda_B}{P_A\lambda_A} \cdot \frac{2 - P_A\lambda_A}{1 - P_A\lambda_A}.$$

From the perspective of a social planner, given that $\alpha_1 = A$, the aim is to assign that approach to agent 2 which provides the highest MSB. Since there is no benefit from having more than one success, agent 2 only contributes to the social surplus in the event that agent 1 does not succeed on approach A. Conditional on this event, it is less likely that approach A is feasible. Let

$$\mu_{A,1} = \mathbb{P}(\text{A is feasible}|\text{one failed attempt on A}) = \frac{P_A(1 - \lambda_A)}{P_A(1 - \lambda_A) + 1 - P_A}.$$

Then the MSB of agent 2 also following A can be written as

$$\mu_{A,1}\lambda_A(1 - P_A\lambda_A)v - c,$$

whereas the MSB of agent 2 following B is

$$P_B\lambda_B(1-P_A\lambda_A)v-c.$$

Hence $\alpha_2 = A$ is efficient if and only if $\mu_{A,1}\lambda_A > P_B\lambda_B$. This is equivalent to

$$\rho_A < \frac{P_A\lambda_A - P_B\lambda_B}{P_A\lambda_A}.$$

Consequently, whenever $\frac{P_A\lambda_A - P_B\lambda_B}{P_A\lambda_A} < \rho_A < \frac{P_A\lambda_A - P_B\lambda_B}{P_A\lambda_A} \cdot \frac{2-P_A\lambda_A}{1-P_A\lambda_A}$, both agents will follow approach A in equilibrium, even though it is socially optimal that they pursue different approaches. Put differently, there is over-investment in the most viable approach A because equilibrium behavior is determined by the average profit of all agents following A, instead of the MSB.

3.3 First Best: Greedy Algorithm

For the analysis of the social planner's problem, some additional notation is useful. For a given action profile $\alpha = (\alpha_1, ..., \alpha_n)$, denote the number of agents following approach a by $n_a(\alpha)$. Moreover, denote by $\pi(\alpha)$ the probability that there is at least one success. Finally, let the belief that approach a is feasible conditional on i failures on a be denoted by $\mu_{a,i}$, that is,

$$\mu_{a,i} = \frac{P_a(1-\lambda_a)^i}{P_a(1-\lambda_a)^i + 1 - P_a}.$$

Notice that the probability that α does not lead to a success can be written as

$$1 - \pi(\alpha) = \prod_{a\in\mathcal{A}} \prod_{i=0}^{n_a(\alpha)-1} (1-\mu_{a,i}\lambda_a).$$

The social planner's problem is simply

$$\max_{\alpha}\ \pi(\alpha)v - \sum_{i\in\mathcal{N}} c_{\alpha_i}.$$

For a given action profile α, the *marginal social benefit of an additional agent following approach* a, denoted $MSB_a(\alpha)$, is the probability that only the additional agent succeeds multiplied with the principal's valuation v, minus the cost of following a, that is,

$$MSB_a(\alpha) = \mu_{a,n_a(\alpha)}\lambda_a(1-\pi(\alpha))v - c_a.$$

The composition of α affects $MSB_a(\alpha)$ through two separate channels. One channel is that the more attempts on the same approach $n_a(\alpha)$ there are, the lower the belief $\mu_{a,n_a(\alpha)}$ will be. This means that conditional on all other attempts failing, the probability that an additional attempt on a succeeds decreases in $n_a(\alpha)$. The other channel is that the larger the probability that α already leads to a success, the smaller will be the "residual value" of a success $(1 - \pi(\alpha))v$.[9]

To separate these two channels, it is helpful to consider the *hypothetical MSB of an additional attempt on a, given i failures*, which is the function $MSB_{a,i} : [0, 1] \to \mathbb{R}$ given by

$$MSB_{a,i}(\pi) = \mu_{a,i}\lambda_a\pi - c_a.$$

Figure 1 plots the hypothetical MSB for the first, second and third attempt on some approach a.

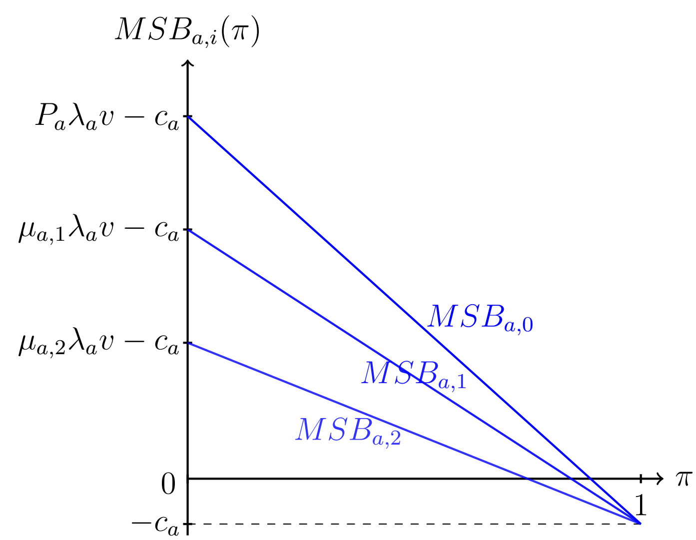

Figure 1: Hypothetical marginal social benefit (MSB) of the first, second, and third attempt on approach a for varying success probability π.

At first sight, a greedy algorithm that sequentially assigns agents to the approach where their MSB is the greatest seems appealing. Such an algorithm would run until there is no approach left where an agent would have a positive MSB, or there is no agent left to assign.[10]

[9] This second channel would disappear if attempts were made sequentially.

[10] Chade and Smith (2006) show that in a similar setting a greedy algorithm, which they call marginal improvement algorithm, identifies the social optimum. In their setting, each option can be used at most once. Here, in contrast, multiple agents can follow the same approach and their successes are correlated. Therefore, an approach cannot just be replaced by n identical copies.

Definition 1 (Greedy algorithm). *Start with an action profile* $\alpha^0 = (\emptyset, ..., \emptyset)$ *where no agent follows an approach, and* $i = 1$.

Step 1 Choose any $\alpha_i^* \in \arg\max_a MSB_a(\alpha^{i-1})$.

Step 2 If $MSB_{\alpha_i^*}(\alpha^{i-1}) \leq 0$, *then set* $\alpha^* = \alpha^{i-1}$ *and stop.*

Step 3 Set $\alpha^i = (\alpha_1^*, ..., \alpha_i^*, \emptyset, ..., \emptyset)$.

Step 4 If $i < n$, *increase* i *by* 1 *and go to Step 1, else set* $\alpha^* = \alpha^n$ *and stop.*

However, a simple example shows that it is not generally possible to attain the social optimum by applying the greedy algorithm.

Example 2. *There are two agents,* 1 *and* 2*, and two approaches,* A *and* B. *Both approaches are feasible, that is,* $P_A = P_B = 1$. *The approaches differ regarding the (conditional) probabilities of successes, with* $\lambda_A = 1$ *and* $\lambda_B = 0.8$. *Also, the approaches differ regarding their costs, with* $c_A = 1$ *and* $c_B = \varepsilon > 0$. *The principal has a valuation of* $v = 10$.

In this example, the marginal social benefit of the first agent that is assigned to an approach will be $v - c_A = 9$ if the agent is assigned to A, and only $0.8v - c_B = 8 - \varepsilon$ if the agent is assigned to B. Thus the greedy algorithm would assign the first agent to A. Moreover once an agent is assigned to A, success is already certain and the algorithm would stop.

However, if both agents follow approach B instead, the social surplus will be

$$(1 - (1 - \lambda_B)^2)v - 2c_B = 9.6 - 2\varepsilon.$$

Hence for $\varepsilon < 0.3$, the greedy algorithm fails to identify the first best.

Figure 2 plots the hypothetical MSB of both approaches in Example 2 for varying levels of $\pi(\alpha)$. Since $P_A = P_B = 1$, the MSB of an additional attempt does not depend on the number of attempts on the same approach in α, but only on $\pi(\alpha)$. The important thing to notice is the order reversal of the MSB of both approaches. For low values of $\pi(\alpha)$, approach A provides a higher MSB, but for $\pi(\alpha) > \frac{1+\varepsilon}{2}$ the MSB of B is larger.

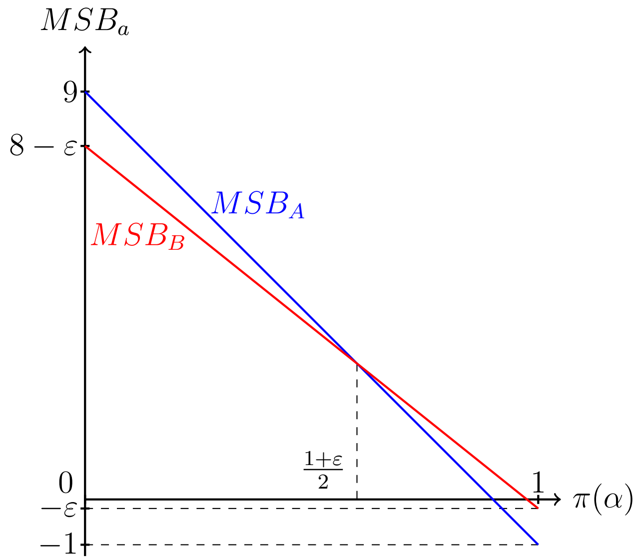

Figure 2: Hypothetical marginal social benefit (MSB) of additional attempts on approaches A and B from Example 2 for varying success probability $\pi(\alpha)$.

Example 2 already suggests that the greedy algorithm fails because the two approaches have significantly different costs. Indeed, the following result shows that when the costs of all approaches are identical, the greedy algorithm always identifies the first best.

Proposition 1. *If all approaches have equal costs $c > 0$, then the greedy algorithm always attains a socially optimal action profile α^*.*

Proof. See Appendix.

An obvious shortcoming of Proposition 1 is that it only applies when costs are identical for all approaches. However in applications where the greatest part of the research costs actually stems from standardized testing procedures, like clinical trials for a vaccine, assuming equal costs is not too far-fetched.

The logic behind Proposition 1 is that when costs are identical for all approaches, the ordering of the MSB of additional attempts is unaffected by $\pi(\alpha)$. Hence a socially optimal action profile can never contain an attempt that has a lower MSB than another that is not contained. Such an action profile could always be improved by removing the former and including the latter attempt since the ordering of the MSB is the same everywhere. Figure 3 illustrates this graphically.

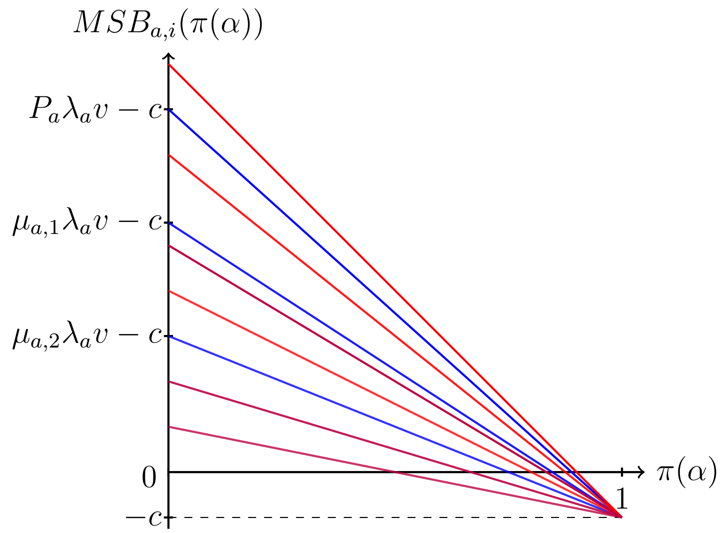

Figure 3: Illustration of Proposition 1 in a scenario with three different approaches (color coded as blue, red, and purple). The ordering of the hypothetical marginal social benefit (MSB) is unaffected by $\pi(\alpha)$. Hence the greedy algorithm attains the first best.

3.4 Implementation: Approach-Specific Rewards

The next result shows that any desired action profile can be implemented as a unique[11] Nash equilibrium, through a suitable *approach-specific contest*. A contest is approach-specific if for all approaches a there is a reward $\bar{w}_a \in \mathbb{R}_+$ which is shared equally among all agents that succeed by following approach a.

Proposition 2. *Consider an arbitrary action profile α. If $c_a > 0$ and $P_a\lambda_a > 0$ for all approaches a*[12]*, then there exists $\gamma > 0$ such that for all $\varepsilon \in (0, \gamma)$, α and every permutation of α constitute Nash equilibria of an approach-specific contest with rewards*

$$\bar{w}_a := \begin{cases} \frac{n_a(\alpha)(c_a+\varepsilon)}{1-\prod_{i=0}^{n_a(\alpha)-1}(1-\mu_{a,i}\lambda_a)} & \text{if } n_a(\alpha) > 0, \\ 0 & \text{if } n_a(\alpha) = 0, \end{cases}$$

Furthermore, there are no Nash equilibria that are not permutations of α.

Proof. See Appendix.

[11] Excluding permutations of the action profile.

[12] These mild assumptions are necessary since agents cannot be induced to follow an approach that is guaranteed to fail, or dissuaded from following a costless approach.

Proposition 2 does not only show that any action profile is implementable as a unique equilibrium, but it also shows that the principal can extract (almost) the entire social surplus at the same time. This is the case since the expected payoff of every agent following an approach a is

$$\frac{1}{n_a(\alpha)} \underbrace{\left(1 - \prod_{i=0}^{n_a(\alpha)-1} (1 - \mu_{a,i}\lambda_a)\right)}_{=\mathbb{P}(\text{One of the } n_a(\alpha) \text{ attempts on } a \text{ succeeds})} \bar{w}_a = \varepsilon$$

The symmetry of all agents following the same approach implies that they all receive an equal share of the expected reward that is paid out for successes on their approach. However, research efforts on the same approach are strategic substitutes, that is, the expected reward of an agent following some approach a is decreasing in $n_a(\alpha)$. By selecting a reward that is slightly above the break-even point when $n_a(\alpha)$ agents follow a, the principal uses the forces of competition in her favor. If fewer than $n_a(\alpha)$ agents were to follow a, it would be profitable for another agent to join. Conversely, in equilibrium no more than $n_a(\alpha)$ agents can follow a because every agent's expected profit would be negative.

What is special about approach-specific contests is neither that they can implement any desired behavior as a Nash equilibrium, nor that they allow extracting the full social surplus. Both could be done more simply by setting

$$w_i(\alpha, s) = \begin{cases} \frac{c_a}{P_a \lambda_a} & \text{if } \alpha_i = a \text{ and } s_i = 1, \\ 0 & \text{else.} \end{cases}$$

But with such a reward scheme, *any* action profile would be a Nash equilibrium. It is special that approach-specific contests can use the competition within approaches to *uniquely* implement the desired equilibrium behavior.

Since Proposition 2 shows that the principal can uniquely implement any desired behavior and extract almost the full social surplus at the same time, the following corollary follows almost immediately.

Corollary 1. *It is optimal for the principal to implement the first best.*

Proof. See Appendix.

4 Dynamic Aspects: Two Periods

When there is more than one period, it matters how the costs and possible successes of approaches are distributed across time. If one approach is inherently "faster" than other approaches, this has two dissimilar effects. On the one hand, an early success can decrease the total costs incurred by all agents. These savings can be enforced by never rewarding agents for a success that occurs after the period of the first success.

On the other hand, when there are agents following fast approaches, this competition may discourage others from following slower approaches. In the case where a slow approach has a comparatively higher total probability of success, this crowding-out effect can lead to inefficient equilibrium behavior.

4.1 Approach-Independent Contests: Crowding-Out Effect

Consider a fixed reward $\bar{w}$ that is awarded to the first successful agent, irrespective of the approach, and shared equally if multiple agents succeed first in the same period. This can be thought of as representing a market structure where the first agents that succeed can capture the entire market. Economic intuition suggests that increasing the reward $\bar{w}$ increases agents' incentives to participate in the contest and should therefore always at least weakly increase the probability of obtaining a success. The following example illustrates that, in fact, increasing the reward can sometimes even decrease the total probability of success.

Example 3. *There are two agents, 1 and 2, and two approaches, A and B. Both approaches are feasible, that is,* $P_A = P_B = 1$*. Both approaches also have the same costs, and all costs are incurred in the first period. approach B's probability of success is* $\frac{2}{3}$*. This is higher than the probability of success on approach A, which is* $\frac{1}{2}$*. However, approach A is "faster": successes on approach A always occur in period* 1*, while successes on approach B always occur in period* 2*. In summary,*

$$[(\lambda_{A,1}, c_{A,1}), (\lambda_{A,2}, c_{A,2})] = [(\frac{1}{2}, c), (0, 0)], \text{ and}$$
$$[(\lambda_{B,1}, c_{B,1}), (\lambda_{B,2}, c_{B,2})] = [(0, c), (\frac{2}{3}, 0)].$$

For $\bar{w} < \frac{c}{\lambda_{B,2}} = \frac{3}{2}c$, it is a dominant strategy for both agents to abstain.

For $\frac{3}{2}c < \bar{w} < 2c$, approach B is profitable, while A is not. Thus in any equilibrium, one of the agents follows B while the other abstains. Both agents following B cannot be an equilibrium because $\frac{1-(1-\lambda_{B,2})^2}{2}\bar{w} = \frac{4}{9}\bar{w} < c$. Hence the probability of a success is $\lambda_{B,2} = \frac{2}{3}$.

For $\bar{w} > 2c$, approach A becomes profitable, at least for a single agent. However for $\bar{w} < 3c$, if one agent follows approach A, it is no longer profitable for the other to follow B because $\bar{w}\lambda_{B,2}(1-\lambda_{A,1}) = \frac{\bar{w}}{3} < c$. Additionally for $\bar{w} < \frac{8}{3}c$, both agents following A cannot be an equilibrium because $\frac{1-(1-\lambda_{A,1})^2}{2}\bar{w} = \frac{3}{8}\bar{w} < c$. Thus for $\bar{w} \in (2c, \frac{8}{3})$, in any equilibrium one agent follows A while the other abstains. Hence the probability of success is only $\lambda_{A,1} = \frac{1}{2}$.

Therefore, by increasing the reward $\bar{w}$ from somewhere in the interval $\left(\frac{3}{2}c, 2c\right)$ to a reward in $\left(2c, \frac{8}{3}c\right)$, the principal actually decreases the probability of a success. For a valuation $v > \frac{8}{3}c$, this implies that both the principal's expected payoff and total welfare decrease as well.

Recall that in the static case an approach-independent contest only induced inefficient equilibrium behavior when successes within approaches were sufficiently correlated. In Example 3, for $\bar{w} \in (2c, \frac{8}{3})$, equilibrium behavior is clearly inefficient, even though the successes within approaches are uncorrelated. This is due to the different timing of successes. Agents following the faster but less promising approach A crowd out their competition on the more effective approach B.

To analyze this crowding-out effect more generally, the next example considers a similar, but more general setup than Example 3.

Example 4. *Modifying the model slightly, suppose there is an infinite number of agents. There are two approaches A and B, that are both feasible ($P_A = P_B = 1$). Costs are identical for both approaches, and A is faster than B, while B is more effective:*

$$[(\lambda_{A,1}, c_{A,1}), (\lambda_{A,2}, c_{A,2})] = [(\phi_A, c), (0, 0)],$$
$$[(\lambda_{B,1}, c_{B,1}), (\lambda_{B,2}, c_{B,2})] = [(0, c), (\phi_B, 0)],$$

with $\phi_B > \phi_A > 0$, and $c > 0$.

For $\bar{w} > \frac{c}{\phi_A}$, approach A is profitable. Agents following A are not affected by those following B because they always succeed first. Hence in any equilibrium, the number of agents following A, denoted n_A, must fulfill

$$\frac{\bar{w}}{n_A}\left(1-(1-\phi_A)^{n_A}\right) \geq c \geq \frac{\bar{w}}{n_A+1}\left[1-(1-\phi_A)^{n_A+1}\right]. \tag{1}$$

Since $\frac{1-(1-\phi_A)^{n_A}}{n_A}$ is decreasing in n_A, the number of agents following A is weakly increasing in $\bar{w}$. This reflects the intuition that a larger reward will make it profitable for more agents to participate in the contest.

Given n_A, the number of agents following B must then satisfy

$$\frac{(1-\phi_A)^{n_A}\bar{w}}{n_B}\left(1-(1-\phi_B)^{n_B}\right) \geq c \geq \frac{(1-\phi_A)^{n_A}\bar{w}}{n_B+1}\left(1-(1-\phi_B)^{n_B+1}\right).$$

Here, $(1-\phi_A)^{n_A}\bar{w}$ can be viewed as the "effective reward" for a success on approach B, because agents following B are only rewarded if all agents on A fail.

How does the effective reward $(1-\phi_A)^{n_A}\bar{w}$ behave as $\bar{w}$ increases? Since n_A depends on $\bar{w}$, this is non-trivial. Clearly, it is increasing on every interval where increasing $\bar{w}$ does not increase n_A. Also, there is a discontinuous drop whenever n_A increases by one. Ultimately, though, the increase in n_A dominates, and the effective reward tends to zero, which implies the following remark.

Remark 1. *In the setting of Example 4, for $\bar{w} \to \infty$, the number of agents following approach B goes to zero.*

Proof. See Appendix.

The intuition is that the crowding-out effect of agents following approach B builds up exponentially, while the reward $\bar{w}$ grows linearly. It may be that at first the increase in $\bar{w}$ outweighs the effects of increased competition from agents on the faster approach A. Consider a prize $\bar{w}$ and the corresponding number n_A of agents following A that fulfills (1). Increasing $\bar{w}$ with the factor $\frac{n_A+1}{n_A}$ must always increase the number of agents following A by at least one. Thus, this increase in $\bar{w}$ changes the effective reward for approach B by the factor $\frac{n_A+1}{n_A}(1-\phi_A)$. For n_A large enough, this factor is smaller than 1. Hence after a certain point, the effective reward for approach B grows smaller with every increase in n_A.

4.2 First Best: Iterative Greedy Algorithm

In a dynamic setting, early successes can reduce the total costs incurred by all agents. Since only one success is socially beneficial, efficiency demands that all agents stop following their approach in the period after the first success occurs.

To be able to provide a (polynomial time) algorithm that determines the first best in the two-period model, I make the following two assumptions.

Assumption 1 (Fast and slow approaches)**.** *All approaches $a \in \mathcal{A}$, are either fast, which means $\lambda_{a,2} = 0$, or slow, which means $\lambda_{a,1} = 0$.*

This means that fast approaches can only lead to success in period 1, and slow approaches can only lead to success in period 2.[13] Denote the set of all fast approaches by

[13] If an approach never leads to success, it can be both fast and slow, but it is also useless.

$A_F \subset \mathcal{A}$, and the set of all slow approaches by $A_S \subset \mathcal{A}$.

Assumption 2 (Equal costs)**.** *There exist $c_F, c_{S,1}, c_{S,2} > 0$ such that $c_{a_F,1} = c_F$[14] for all $a_F \in A_F$, and $c_{a_S,1} = c_{S,1}$ and $c_{a_S,2} = c_{S,2}$ for all $a_S \in A_S$.*

Assumption 1 makes it possible to split the set of agents into two groups, those that follow slow approaches and those that follow fast approaches. Furthermore, keeping the behavior of one of the groups fixed, it is possible to analyze the behavior of the other group in isolation.

Consider a fixed set of agents $N_S \subset \mathcal{N}$ that follow slow approaches. Conditional on no other agent succeeding in period 1, efficiency requires that all agents in N_S continue on their approach in period 2. Otherwise, the costs $c_{S,1}$ they incur in period 1 would be wasted. Denote by π_S the probability that one of the agents in N_S succeeds if all of them follow their approach for both periods. Also, denote by n_S the number of agents in N_S. Taking P_S and n_S as given, the expected benefit of a success in period 1 is

$$v_F = (1 - \pi_S)v + n_S c_{S,2}.$$

The early success allows all agents on a slow approach to quit and save the costs $c_{S,2}$. However, the early success does not generate the full value v, because with probability π_S an agent on a slow approach would still succeed.

Thus for a given behavior of agents following slow approaches, the greedy algorithm from the static case, with v_F instead of v and c_F instead of c, can be used to determine the corresponding optimal behavior of the remaining agents when they may only use fast approaches.

Conversely, consider a fixed set of agents $N_F \subset \mathcal{N}$ that follow fast approaches. Denote by π_F the probability that one of the agents in N_F succeeds and by n_F the number of agents in N_F. Taking π_F and n_F as given, the (unconditional) expected benefit of a success in period 2 is

$$v_S = (1 - \pi_F)v.$$

However, the expected total cost of following a slow approach is only $c_{S,1} + (1 - \pi_F)c_{S,2}$ because with probability π_F an agent following a fast approach will succeed in period 1, and then agents on slow approaches can quit in period 2 and save $c_{S,2}$.

Thus for a given behavior of agents following fast approaches, the greedy algorithm can be used to determine the corresponding optimal behavior of the remaining agents if they

[14] The costs in period 2, $c_{a_F,2}$, are irrelevant because $\lambda_{a_F,2} = 0$ for fast approaches.

are restricted to slow approaches. In this case, v_S takes the role of v and $c_{S,1} + (1 - \pi_F)c_{S,2}$ the role of c.

It remains to find the optimal combination of the number of agents following slow and fast approaches. This can be addressed by iteratively considering all candidate solutions. A candidate solution is one where $n_F \in \{0, ..., n\}$ agents are assigned optimally to fast approaches only, and the remaining $n_S := n - n_F$ are then assigned to slow approaches only, in optimal fashion given the probability P_F that one of the n_F agents assigned to fast approaches succeeds. The following paragraph introduces notation to formally describe this iterative procedure.

For all $a_F \in A_F$, let $\lambda_{a_F} = \lambda_{a_F,1}$, and for all $a_S \in A_S$, let $\lambda_{a_S} = \lambda_{a_S,2}$. Moreover, for all $a \in \mathcal{A}$, and all $i \in \{1, ..., n\}$, let

$$\mu_{a,i} = \frac{P_a(1-\lambda_a)^i}{P_a(1-\lambda_a)^i + 1 - P_a}.$$

Denote by $\sigma_i = \emptyset$ the strategy where agent i abstains in period 1. Denote by $\sigma_i = a_F \in \mathcal{A}_F$ the strategy where i follows a_F in period 1 and then quits in period 2. Denote by $\sigma_i = a_S \in \mathcal{A}_S$ the strategy where i follows a_S in period 1 and then continues in period 2 if and only if no agent succeeds in period 1. To determine the social optimum, it is without loss to focus only on these three types of strategy.

The *expected marginal social benefit of an additional agent using strategy* $\sigma_i = a \in \mathcal{A}$, denoted by $MSB_a(\sigma)$, is

$$MSB_a(\sigma) = \begin{cases} \mu_{a,n_a}\lambda_a(1-\pi_F)\left[(1-\pi_S)v + n_S c_{S,2}\right] - c_F & \text{if } a \in \mathcal{A}_F, \\ \mu_{a,n_a}(1-\pi_F)(1-\pi_S)v - c_{S,1} - (1-\pi_F)c_{S,2} & \text{if } a \in \mathcal{A}_S. \end{cases}$$

Finally, the *expected social surplus* is

$$W(\sigma) = \left[1 - (1-\pi_F)(1-\pi_S)\right]v - n_F c_F - n_S(c_{S,1} + (1-P_F)c_{S,2}).$$

Definition 2 (Iterative greedy algorithm)**.** *Start with strategy profile* $\sigma^0 = (\emptyset, ..., \emptyset)$*,* $i = 0$*, and* $W^* = 0$*.*

Step 1: Start with $j = i + 1$ *and* $\sigma^{i,j} = \sigma^i = (\sigma_1, ..., \sigma_i, \emptyset, ..., \emptyset)$*.*

(a) Choose any $\sigma_j \in \arg\max_{a \in A_S} MSB_a(\sigma^{i,j})$*.*

(b) If $MB_{\sigma_j}(\sigma^{i,j}) \leq 0$*, then set* $\sigma^{i,*} = \sigma^{i,j}$ *and go to Step 2.*

(c) Set $\sigma^{i,j+1} = (\sigma_1, ..., \sigma_j, \emptyset, ..., \emptyset)$*.*

(d) If $j < n$*, increase* j *by* 1 *and go to (a), else set* $\sigma^{i,*} = \sigma^{n,*}$ *and go to Step 2.*

Step 2: If $W(\sigma^{i,*}) > W^*$*, set* $W^* = W(\sigma^{i,*})$ *and* $\sigma^* = \sigma^{i,*}$*.*

Step 3: Choose any $\sigma_{i+1} \in \arg\max_{a \in A_F} MSB_a(\sigma^i)$.
Step 4: Set $\sigma^{i+1} = (\sigma_1, ..., \sigma_{i+1}, \emptyset, ..., \emptyset)$.
Step 5: If $i + 1 < n$*, increase* i *by* 1 *and go to Step 1.*
Step 6: If $W(\sigma^n) > W^*$*, set* $W^* = W(\sigma^n)$*, and* $\sigma^* = \sigma^n$.

Proposition 3. *If Assumptions 1 and 2 are fulfilled, then the iterative greedy algorithm identifies a socially optimal strategy profile* σ^*.

Proof. See Appendix.

4.3 Implementation: Efficient-Stopping Condition

Efficiency demands that agents only continue following their approach in period 2 if no one succeeded in period 1. It seems reasonable that the principal should attempt to harvest these efficiency gains, by selecting a contest that never rewards agents for a success in period 2 if there has already been a success in period 1. In contrast to the approach-specific contests from the static case, such reward schemes create inter-dependencies between the actions and expected rewards of agents following different approaches.

A natural extension of approach-specific contests from the static case to two periods is a contest that consists of *within-period approach-specific contests with efficient-stopping condition*, that is, a reward scheme with prizes $\bar{w}_{a,1}$ and $\bar{w}_{a,2}$ for each approach a such that for a terminal history $h^T = (\alpha_1, s_1, \alpha_2, s_2)$,

$$w_i(h^T) = \begin{cases} \frac{1}{\sum_{j \in \mathcal{N}} \mathbb{1}_{\{\alpha_{j,1} = \alpha_{i,1} \text{ and } s_{j,1} = 1\}}} \bar{w}_{\alpha_{i,1},1} & \text{if } s_{i,1} = 1, \\ \frac{1}{\sum_{j \in \mathcal{N}} \mathbb{1}_{\{\alpha_{j,2} = \alpha_{i,2} \text{ and } s_{j,2} = 1\}}} \bar{w}_{\alpha_{i,2},2} & \text{if } s_{1,1} = ... = s_{n,1} = 0 \text{ and } s_{i,2} = 1, \\ 0 & \text{if } s_{i,1} = s_{i,2} = 0, \\ 0 & \text{if } s_{i,1} = 0, \text{ and there exists } j \neq i \text{ s.t. } s_{j,1} = 1. \end{cases}$$

Such a reward scheme gives the principal a lot of flexibility to tailor agents' incentives. The principal can set different rewards for different approaches, and also set different rewards for the same approach in different periods.

To fix ideas, consider the minimal case with one agent and a single approach a with $P_a \in (0, 1)$, and $\lambda_{a,1}, \lambda_{a,2} > 0$, and v sufficiently large. If the principal sets prizes

$$w_{a,1} = \frac{c_{a,1}}{P_a \lambda_{a,1}}$$

and

$$w_{a,2} = \frac{c_{a,2} + \varepsilon}{\mu_{a,1}\lambda_{a,2}}$$

with $\mu_{a,1} = \frac{P_a(1-\lambda_{a,1})}{P_a(1-\lambda_{a,1})+1-P_a}$, and $\varepsilon > 0$, it is optimal for the agent to follow a in period 1, and continue in period 2 if and only if she did not succeed in period 1. The expected payoff from following this strategy is $(1 - P_a\lambda_a)\varepsilon > 0$.

For ε sufficiently small, $c_{a,1} = c_{a,2}$ and $\lambda_{a,1} = \lambda_{a,2}$, the above prizes satisfy $w_{a,1} < w_{a,2}$, which means that the principal increases the reward over time to compensate the agent for a lower belief. This allows the principal to extract (almost) the entire social surplus. In effect, the principal exploits the observability of successes and the agent's inability of waiting in period 1 and postponing effort to period 2.

Interestingly, with multiple agents and approaches, it is in general not possible to implement the first best with within-period approach-specific contests with efficient-stopping condition, as the following example shows.

Example 5. *There are 3 agents,* 1, 2 *and* 3, *and two approaches, A and B. approach A is fast and approach B is slow, and successes on both approaches are fully correlated. More specifically, $P_A = \frac{3}{4}$ and $P_B = \frac{1}{2}$, and*

$$[(\lambda_{A,1}, c_{A,1}), (\lambda_{A,2}, c_{A,2})] = [(1, c), (0, c)],$$
$$[(\lambda_{A,1}, c_{A,1}), (\lambda_{A,2}, c_{A,2})] = [(0, \frac{3}{2}c), (1, \frac{1}{2}c)],$$

with $c > 0$ and $v = 100c$.

In this setting, efficiency requires that one agent follows A in period 1 and then quits in period 2, and a second agent follows B in period 1 and continues in period 2 if and only if the agent following A did not succeed in period 1. Because successes are fully correlated within approaches, the third agent has to abstain in period 1.

If the principal wants to implement the first best with within-period approach-specific contests with efficient-stopping condition, she has to set $w_{A,1} \geq \frac{4}{3}c$ so that it is profitable for one agent to follow A.

Given that one agent follows A, the principal has to select

$$w_{B,2} \geq \frac{c_{B,1} + (1 - P_A\lambda_{A,1})c_{B,2}}{(1 - P_A\lambda_{A,1})P_B\lambda_{B,2}} = 13c$$

so as to induce at least one agent to follow B in period 1.

Suppose the principal exactly selects $w_{A,1} = \frac{4}{3}c$, and $w_{B,2} = 13c$. In this case, the efficient behavior is not a Nash equilibrium. There is no profitable deviation for the agent that abstains or for the agent on the slow approach B. However, there is a profitable deviation for the agent that follows the fast approach A. By not following A, this agent increases the probability that $w_{B,2}$ is paid out. If the agent deviates to instead following approach B for both periods, her expected payoff is

$$\frac{w_{B,2}}{2} P_B \lambda_{B,2} - c_{B,1} - c_{B,2} = \frac{13c}{2} \cdot \frac{1}{2} - 2c = \frac{5}{4}c > 0.$$

It stands to reason that by increasing $w_{A,1}$ the principal could dissuade the agent from switching from A to B. But to achieve this, it is necessary that $\frac{3}{4}w_{A,1} - c \geq \frac{5}{4}c$, which is equivalent to $w_{A,1} \geq 3c$. When $\bar{w}_{A,1}$ is this high, it is no longer optimal for the third agent to abstain. If the third agent deviates to following A in period 1, she has an expected payoff of $\frac{3}{4} \cdot \frac{\bar{w}_{A,1}}{2} - c \geq \frac{1}{8}c > 0$.

Example 5 shows that it is in general not possible to implement the first best by using within-period approach-specific contests with efficient-stopping condition. This failure is caused by the fact that agents that follow a fast approach can increase the effective reward of slower approaches by deviating to those slower approaches.

Such deviations are relatively more profitable if there are only a few agents following the slow approach because the deviating agent receives a larger share of the increase in the effective reward. The following result shows that for this reason, if the principal's valuation v and the number of agents are both large enough, the social optimum can always be implemented through within-period approach specific contests with an efficient-stopping condition.

Proposition 4. *Suppose Assumption 1 is fulfilled and let $c_{a,1}, c_{a,2} > 0$, as well as $P_a, \lambda_{a,1}, \lambda_{a,2} < 1$ for all $a \in \mathcal{A}$. Then there exist $\hat{v}$ and $\hat{n}$ such that for all valuations $v > \hat{v}$, and all $n > \hat{n}$, there exists $\gamma > 0$ such that for all $\varepsilon \in (0, \gamma)$, the principal can implement the socially optimal strategy profile σ^*, via an approach-specific contest with efficient-stopping condition, featuring*

$$w_a = \frac{n_a^*(c_{a,1} + \varepsilon)}{P_a \left(1 - (1 - \lambda_{a,1})^{n_a^*}\right)} \qquad \text{for all } a \in A_F \text{ and}$$

$$w_a = \frac{n_a^*(c_{a,1} + (1 - \pi_F^*)c_{a,2} + \varepsilon)}{(1 - \pi_F^*) P_a \left(1 - (1 - \lambda_{a,2})^{n_a^*}\right)} \qquad \text{for all } a \in A_S.$$

Here, n_a^* denotes the socially optimal number of agents following approach a in period 1 for $a \in A_F$, while for $a \in A_S$ it denotes the socially optimal number of agents that follow a in period 1, and continue in period 2 if and only if no success occurred in period 1. π_F^* denotes the probability that one of the agents following fast approaches succeeds.

Proof. See Appendix.

5 Conclusion

I have studied the design of innovation contests in the presence of distinct approaches. It turned out that approach-independent contests may induce inefficient equilibrium behavior for two different reasons. When some approaches are more promising than others, but successes on these approaches are correlated, there may be over-investment in these approaches. In addition, when some approaches are faster than others, agents following the faster approaches may crowd out agents on slower, more effective approaches.

Furthermore, I found that in a static setting a greedy algorithm can be used to determine the socially optimal behavior if all approaches have the same costs. In a two-period setting, in addition to the equal-cost assumption, the assumption that all approaches are either fast or slow was sufficient to show that an iterated version of the greedy algorithm can be used to determine the social optimum.

In the static case, approach-specific contests turned out to be a powerful tool for the principal to uniquely implement any desired behavior and extract the social surplus at the same time. In the dynamic model, augmenting approach-specific contests with an efficient-stopping condition may fail to implement the desired behavior because it creates inter-dependencies between the payoffs of agents following different approaches. However, when the principal's valuation and the number of agents are both large enough, the socially optimal behavior *can* be implemented in this way, and hence the principal can extract the full social surplus.

References

Acemoglu, Daron, Kostas Bimpikis, and Asuman Ozdaglar (2011). "Experimentation, Patents, and Innovation". In: *American Economic Journal: Microeconomics* 3.1, pp. 37–77.

Akcigit, Ufuk and Qingmin Liu (2016). "The Role of Information in Innovation and Competition". In: *Journal of the European Economic Association* 14.4, pp. 828–870.

Baden, Lindsey R. et al. (2021). "Efficacy and Safety of the mRNA-1273 SARS-CoV-2 Vaccine". In: *New England Journal of Medicine* 384.5, pp. 403–416.

Callaghan, Chris (2020). *Would a Longitude Prize Speed Production of a Covid-19 Vaccine?* Times Higher Education (THE). URL: https://www.timeshighereducation.com/blog/would-longitude-prize-speed-production-covid-19-vaccine (visited on 11/02/2021).

Chade, Hector and Lones Smith (2006). "Simultaneous Search". In: *Econometrica* 74.5, pp. 1293–1307.

Che, Yeon-Koo and Ian Gale (2003). "Optimal Design of Research Contests". In: *American Economic Review* 93.3, pp. 646–671.

Choi, Jay P. (1991). "Dynamic R&D Competition under "Hazard Rate" Uncertainty". In: *The RAND Journal of Economics* 22.4, p. 596.

Fu, Qiang and Zenan Wu (2019). "Contests: Theory and Topics". In: *Oxford Research Encyclopedia of Economics and Finance*. Oxford University Press.

Halac, Marina, Navin Kartik, and Qingmin Liu (2017). "Contests for Experimentation". In: *Journal of Political Economy* 125.5, pp. 1523–1569.

Hemel, Daniel and Lisa Larrimore Ouellette (2020). *Want a Coronavirus Vaccine, Fast? Here's a Solution*. Time. URL: https://time.com/5795013/coronavirus-vaccine-prize-challenge/ (visited on 11/02/2021).

Le, Tung Thanh et al. (2020). "The COVID-19 Vaccine Development Landscape". In: *Nature Reviews Drug Discovery* 19.5, pp. 305–306.

Letina, Igor (2016). "The Road Not Taken: Competition and the R&D Portfolio". In: *The RAND Journal of Economics* 47.2, pp. 433–460.

Letina, Igor and Armin Schmutzler (2019). "Inducing Variety: A Theory of Innovation Contests". In: *International Economic Review* 60.4, pp. 1757–1780.

Malueg, David A. and Shunichi O. Tsutsui (1997). "Dynamic R&D Competition with Learning". In: *The RAND Journal of Economics* 28.4, pp. 751–772.

Pires, J.C.M. et al. (2011). "Recent Developments on Carbon Capture and Storage: An Overview". In: *Chemical Engineering Research and Design* 89.9, pp. 1446–1460.

Polack, Fernando P. et al. (2020). "Safety and Efficacy of the BNT162b2 mRNA Covid-19 Vaccine". In: *New England Journal of Medicine* 383.27, pp. 2603–2615.

Taylor, Curtis R (1995). "Digging for Golden Carrots: An Analysis of Research Tournaments". In: *American Economic Review* 85.4, pp. 872–890.

Wang, Ying et al. (2021). "Scalable Live-Attenuated SARS-CoV-2 Vaccine Candidate Demonstrates Preclinical Safety and Efficacy". In: *Proceedings of the National Academy of Sciences* 118.29, e2102775118.

Appendix

Proof of Proposition 1

Step 1: (Defining a relaxed optimization problem.)

Recall that

$$\mu_{a,i} := \frac{P_a(1-\lambda_a)^i}{P_a(1-\lambda_a)^i + 1 - P_a},$$

and that

$$\pi(\alpha) = 1 - \prod_{a\in\mathcal{A}} \prod_{i=0}^{n_a(\alpha)-1} (1-\mu_{a,i}\lambda_a).$$

An optimal action profile must solve

$$\max_{\alpha} \pi(\alpha)v - c\sum_{a\in\mathcal{A}} n_a(\alpha).$$

Define $\psi_{a,i} := \lambda_a\mu_{a,i-1}$, and let $\Psi := \{\psi_{a,i} | a \in \mathcal{A}, i \in \mathcal{N}\}$, and consider a solution Σ^* to a relaxed optimization problem:

$$\Sigma^* \in \arg\max_{\Sigma\subset\Psi} \left(1 - \prod_{\psi\in\Sigma}(1-\psi)\right) v - |\Sigma|c \qquad \text{s.t. } |\Sigma| \leq n.$$

It is a relaxed problem because for every action profile α, there is a *corresponding set* $\Sigma^\alpha = \{\psi_{a,i} : a \in \mathcal{A}, i \leq n_a(\alpha)\} \subset \Psi$ that yields the same value in the relaxed problem, as α in the original problem.

Step 2: (Solution to the relaxed problem must be an interval.)

However it must be that for all $\psi, \psi' \in \Psi$, if $\psi > \psi'$ and $\psi' \in \Sigma^*$ hold, then also $\psi \in \Sigma^*$. Otherwise one could improve on Σ^* by removing ψ' and adding ψ. This means that every solution Σ^* is an *interval*, that is,

$$\text{for all } \psi^* \in \Sigma^* \text{ and all } \psi \in \Psi\backslash\Sigma^*, \text{ it holds that } \psi^* \geq \psi.$$

This ensures that there always exists a solution to the relaxed problem that corresponds to an action profile because for all $a \in \mathcal{A}$ and $i \in \mathcal{N}$, it is true that $\psi_{a,i} \geq \psi_{a,i+1}$.[15]

[15] At this point one could theoretically stop, observe that the payoff function $\left(1 - \prod_{\psi\in\Sigma}(1-\psi)\right) v$ is (downward) recursive, and invoke the result about the marginal improvement algorithm of Chade and Smith (2006) except for the possibility of multiple elements in Ψ taking the same value which they exclude.

Step 3: (The greedy algorithm produces an interval in the relaxed problem.)

The marginal social benefit of an additional agent following approach a given some action profile α is

$$MB_a(\alpha) = \left(\psi_{a,n_a(\alpha)+1} \prod_{a \in A} \prod_{i=1}^{n_a(\alpha)} (1 - \psi_{a,i}) \right) v - c,$$

which is maximized when selecting a such that $\psi_{a,n_a(\alpha)+1}$ is maximal. Hence the greedy algorithm will always produce an action profile α^* such that, for some $\eta \leq n$ the corresponding set Σ^{α^*} contains the largest η elements of Ψ.

Step 4: (No interval improves on the one the greedy algorithm produces.)

When the algorithm stops at α^*, it holds that for all $a \in \mathcal{A}$ with $n_a(\alpha^*) > 0$

$$c < \left(\frac{\psi_{a,n_a(\alpha^*)}}{1 - \psi_{a,n_a(\alpha^*)}} \prod_{\tilde{a} \in A} \prod_{i=1}^{n_{\tilde{a}}(\alpha^*)} (1 - \psi_{\tilde{a},i}) \right) v,$$

and additionally either $\eta = n$, or for all $a \in \mathcal{A}$,

$$\left(\psi_{a,n_a(\alpha^*)+1} \prod_{\tilde{a} \in A} \prod_{i=1}^{n_{\tilde{a}}(\alpha^*)} (1 - \psi_{\tilde{a},i}) \right) v \leq c.$$

This means that all assigned agents have a positive marginal social benefit, while it is not possible to assign another agent with a positive marginal social benefit. Since for all $a \in \mathcal{A}$ and $i \in \mathcal{N}$, it holds that $\psi_{a,i} \geq \psi_{a,i+1}$, this implies that for all $\psi' \in \Psi \backslash \Sigma^{\alpha^*}$ and all $\psi^* \in \Sigma^{\alpha^*}$,

$$\left(\psi' \prod_{\psi \in \Sigma^{\alpha^*}} (1 - \psi) \right) v \leq c < \left(\psi^* \prod_{\psi \in \Sigma^{\alpha^*} \backslash \{\psi^*\}} (1 - \psi) \right) v.$$

Combining this with the observation that adding additional elements to Σ^{α^*} will only decrease $\left(\psi' \prod_{\psi \in \Sigma^{\alpha^*}} (1 - \psi) \right) v$, while removing elements from Σ^{α^*} will only increase $\left(\psi^* \prod_{\psi \in \Sigma^* \backslash \{\psi^*\}} (1 - \psi) \right) v$, there can be no interval Σ with $|\Sigma| \leq n$ that improves on Σ^{α^*}.

Thus the greedy algorithm produces an action profile that corresponds to a solution to the relaxed problem. Therefore, the action profile itself must solve the original optimization problem of the social planner. $\square$

Proof of Proposition 2

Step 1: (The action profile α is a Nash equilibrium.)

Consider an arbitrary agent $i \in \mathcal{N}$.

If $\alpha_i = \emptyset$, then i's expected payoff is zero. If $\alpha_i = a \in \mathcal{A}$, then the symmetry of all agents following a implies that i's expected payoff is

$$\begin{aligned}\mathbb{E}\left[w_i(\alpha, s) - c_a\right] &= \frac{\bar{w}_a}{n_a(\alpha)} \mathbb{P}\left[\exists j \in \mathcal{N}, \text{ s.t. } s_j = 1 \text{ and } \alpha_j = a\right] - c_a \\ &= \varepsilon > 0.\end{aligned}$$

Hence, $\hat{\alpha}_i = \emptyset$, or $\hat{\alpha} = a$ with $n_a(\alpha) = 0$ is never a profitable deviation.

Denote the action profile if all other agents play according to α, and agent i deviates to $\hat{\alpha}_i$ by $\hat{\alpha} \coloneqq (\alpha_1, ..., \alpha_{i-1}, \hat{\alpha}_i, \alpha_{i+1}, ..., \alpha_N)$. The expected payoff from deviating to $\hat{\alpha}_i = a \in \mathcal{A}$ with $n_a(\alpha) > 0$ is, again by symmetry among the then $n_a(\alpha) + 1$ agents following a,

$$\begin{aligned}\mathbb{E}\left[w_i(\hat{\alpha}, s) - c_a\right] &= \frac{\bar{w}_a}{n_a(\alpha)+1} \mathbb{P}\left[\exists j \in \mathcal{N}, \text{ s.t. } s_j = 1 \text{ and } \alpha_j = a\right] - c_a \\ &= \frac{\bar{w}_a}{n_a(\alpha)+1}\left(1 - \prod_{i=1}^{n_a(\alpha)+1}(1-\psi_{a,i})\right) - c_a \\ &= \frac{n_a(\alpha)}{n_a(\alpha)+1} \cdot \frac{1 - \prod_{i=1}^{n_a(\alpha)+1}(1-\psi_{a,i})}{1 - \prod_{i=1}^{n_a(\alpha)}(1-\psi_{a,i})}(c_a + \varepsilon) - c_a \\ &= \underbrace{\frac{n_a(\alpha)}{n_a(\alpha)+1} \cdot \frac{\psi_{a,1} + \sum_{i=2}^{n_a(\alpha)+1}\left[\prod_{j=1}^{i-1}(1-\psi_{a,j})\right]\psi_{a,i}}{\psi_{a,1} + \sum_{i=2}^{n_a(\alpha)}\left[\prod_{j=1}^{i-1}(1-\psi_{a,j})\right]\psi_{a,i}}}_{\coloneqq \gamma_a}(c_a + \varepsilon) - c_a.\end{aligned}$$

Recall that $\psi_{a,i} = \lambda_a \mu_{a,i-1}$. It holds that $\psi_{a,1} = P_a \lambda_a > 0$, and $\psi_{a,i} \geq \psi_{a,i+1}$ for all $a \in \mathcal{A}$ and all $i < n$. This implies that

$$\psi_{a,1} > (1-\psi_{a,1})\psi_{a,2} \geq \prod_{j=1}^{2}(1-\psi_{a,j})\psi_{a,3} \geq ... \geq \prod_{j=1}^{n_a(\alpha)}(1-\psi_{a,j})\psi_{a,n_a(\alpha)+1}.$$

Thus it follows that

$$\frac{\psi_{a,1} + \sum_{i=2}^{n_a(\alpha)}\left[\prod_{j=1}^{i-1}(1-\psi_{a,j})\right]\psi_{a,i}}{n_a(\alpha)} > \frac{\psi_{a,1} + \sum_{i=2}^{n_a+1}\left[\prod_{j=1}^{i-1}(1-\psi_{a,j})\right]\psi_{a,i}}{n_a(\alpha)+1},$$

and hence $\gamma_a < 1$.

Thus for all $a \in \mathcal{A}$ with $n_a > 0$, and all $\varepsilon < \frac{1-\gamma_a}{\gamma_a} c_a$,

$$\mathbb{E}\left[w_i(\hat{\alpha}, s) - c_a\right] = (\gamma_a - 1)c_a + \gamma_a \varepsilon < 0.$$

Therefore, for $\varepsilon < \min_a\{\frac{1-\gamma_a}{\gamma_a} c_a\}$, there is no profitable deviation for any agent and α is a Nash-equilibrium.

Step 2: (There may be no other Nash equilibrium)

Assume $\varepsilon < \min_a \frac{1-\gamma_a}{\gamma_a} c_a$. Then for all $a \in \mathcal{A}$, there exists no Nash equilibrium where more than n_a agents follow a. Suppose to the contrary that there is a Nash equilibrium where $z > n_a$ agents follow a. If $n_a = 0$, then $w_a = 0$, and the payoff of all z agents is $-c_a < 0$ for certain. Thus they have a profitable deviation to abstaining from the contest.

If $n_a > 0$, then the expected payoff of all z agents following a is

$$\underbrace{\frac{n_a}{z} \cdot \frac{\psi_{a,1} + \sum_{i=2}^{n_a} \left[\prod_{j=1}^{i-1}(1-\psi_{a,j})\right] \psi_{a,i}}{\psi_{a,1} + \sum_{i=2}^{z} \left[\prod_{j=1}^{i-1}(1-\psi_{a,j})\right] \psi_{a,i}}}_{:=\gamma_a^z}(c_a + \varepsilon) - c_a$$

Again, it holds that $\psi_{a,1} = \phi_a > 0$, and $\psi_{a,i} \geq \psi_{a,i+1}$ for all $a \in \mathcal{A}$ and all $i < n$, which implies that

$$\psi_{a,1} > (1-\psi_{a,1})\psi_{a,2} \geq \prod_{j=1}^{2}(1-\psi_{a,j})\psi_{a,3} \geq ... \geq \prod_{j=1}^{z}(1-\psi_{a,j})\psi_{a,z}.$$

This implies that $\gamma_a^z \leq \gamma_a$. Thus the expected payoff of all agents following a is

$$\gamma_a^z(c_a - \varepsilon) - c_a \leq \gamma_a(c_a - \varepsilon) - c_a < 0.$$

So each of the z agents following a has a profitable deviation of abstaining. Thus there exists no Nash equilibrium where more than n_a agents follow a.

Similarly, for all $a \in \mathcal{A}$ there exists no Nash equilibrium where less than n_a agents follow a. Suppose to the contrary that there is a Nash equilibrium where $x < n_a$ agents follow a. Since there exists no Nash equilibrium where more agents than under α follow any approach, there must be at least one agent that abstains.

The expected payoff from following a is

$$\underbrace{\frac{n_a}{x} \cdot \frac{\psi_{a,1} + \sum_{i=2}^{n_a} \left[\prod_{j=1}^{i-1}(1-\psi_{a,j})\right] \psi_{a,i}}{\psi_{a,1} + \sum_{i=2}^{x} \left[\prod_{j=1}^{i-1}(1-\psi_{a,j})\right] \psi_{a,i}}}_{:=\gamma_a^x}(c_a + \varepsilon) - c_a.$$

By the same logic as above, it holds that $\gamma_a^x \geq 1$. Hence all agents that abstain have a profitable deviation of following a. Thus there exists no Nash equilibrium where less than n_a agents follow a.

Therefore, all Nash equilibria have the same number n_a on every approach a and are thus permutations of α. $\square$

Proof of Corollary 1 Suppose by contradiction that there exists an action profile α^* that the principal can implement and that makes the principal better off than implementing a socially optimal action profile. Denote by W^{FB} the social surplus if the first best is implemented and by W^* the social surplus when α^* is implemented. It must be that $W^{FB} > W^*$. The principal cannot implement an action profile where an agent has a negative expected payoff because every agent can guarantee a payoff of zero by abstaining. So the expected payoff of the principal cannot be greater than W^* when implementing α^*.

By Proposition 2 the principal can set up an approach-specific contest with $\varepsilon < \frac{W^{FB}-W^*}{n}$ such that the first best is implemented, and in equilibrium the expected payoff of an agent following approach a is

$$\frac{\mathbb{P}\left[\exists j \in \mathcal{N}, \text{ s.t. } s_j = 1 \text{ and } \alpha_j = a\right]}{n_a} \bar{w}_a = \varepsilon < \frac{W^{FB} - W^*}{n}.$$

In addition, agents that abstain receive a payoff of zero. Hence the sum of the expected payoffs of all n agents is strictly less than $W^{FB} - W^*$. It follows that the principal's payoff when implementing the first best is strictly greater than W^*, which is a contradiction. $\square$

Proof of Remark 1

By (1), for $\bar{w} \to \infty$ it must be that $n_A \to \infty$. Suppose not, then there exists $M \in \mathbb{R}_+$, so that for all $\bar{w}$, it holds that $\frac{\bar{w}}{M}(1-(1-\phi_A)^M) \leq c$. For $\bar{w} > \frac{M}{1-(1-\phi_A)^M}$, this is a contradiction.

Rewriting (1) yields

$$\frac{c}{1-(1-\phi_A)^{n_A}} \leq \frac{\bar{w}}{n_A} \leq \frac{n_A+1}{n_A} \frac{c}{1-(1-\phi_A)^{n_A+1}}.$$

For $\bar{w} \to \infty$, both the left-hand and the right-hand side go to c, hence also $\frac{\bar{w}}{n_A}$ goes to c.

Let $\varepsilon > 0$, then there exists $x > 0$ such that for all $\bar{w} > x$, it holds that $\frac{n_A}{\bar{w}} > \frac{1}{c+\varepsilon}$. This implies that for all $\bar{w} > x$,

$$0 \leq (1-\phi_A)^{n_A}\bar{w} \leq (1-\phi_A)^{\bar{w}/(c+\varepsilon)}\bar{w}.$$

For $\bar{w} \to \infty$, the right-hand side converges to zero. It follows that $(1-\phi_A)^{n_A}\bar{w} \to 0$. Thus for $\bar{w}$ large enough it holds that $(1-\phi_A)^{n_A}\bar{w} < \frac{c}{\phi_B}$. Therefore for $\bar{w}$ large enough, no agent will follow approach B in equilibrium. □

Proof of Proposition 3

Step 1: (Eliminating inefficient strategies.)

Consider an arbitrary strategy profile $\sigma = (\sigma_1, ..., \sigma_n)$ and an arbitrary agent i. If $\sigma_i(h^0) \in \mathcal{A}_F$, then i's probability of success in period 2 is always zero, hence i quitting in period 2 is socially optimal. If $\sigma_i(h^0) \in \mathcal{A}_S$, then, for any observable history $h^1 \in \mathcal{H}^1$, if h^1 contains a success in period 1, it is socially optimal for i to quit. Consider the case where h^1 is the observable history that realizes when agents act according to σ in period 1 and that does not contain a success in period 1. Then if $\sigma_i(h^1) = \emptyset$, replacing $\sigma_i(h^0) = a_S \in \mathcal{A}_S$ with $\sigma_i(h^0) = \emptyset$ increases the social surplus by $c_{S,1}$.

Thus for every agent i, attention can be restricted to three types of strategies. First, the strategy where agent i already abstains in period 1, denoted by $\sigma_i = \emptyset$. Second, the strategy where i follows $a_F \in \mathcal{A}_F$ in period 1 and then quits in period 2, denoted by $\sigma_i = a_F$. Third, the strategy where i follows $a_S \in \mathcal{A}_S$ in period 1 and then continues in period 2 if and only if no agent succeeds in period 1, denoted by $\sigma_i = a_S$.

Step 2: (Defining a relaxed optimization problem.)

Because successes are independent between different approaches, the social optimum can be written as

$$\max_{(\sigma_1,\dots,\sigma_n)} \left(1 - \prod_{a\in A}\prod_{i=1}^{n_a}(1-\psi_{a,i})\right)v - c_F \cdot n_F - (c_{S,1} + (1-P_F)c_{S,2})\, n_S.$$

Let $\Psi_F \coloneqq \{\psi_{a_F,i} \colon a_F \in \mathcal{A}_F, i \in \mathcal{N}\}$ and $\Psi_S \coloneqq \{\psi_{a_S,i} \colon a_S \in \mathcal{A}_S, i \in \mathcal{N}\}$, and consider a solution (Σ_F^*, Σ_S^*) to a relaxed optimization problem:

$$(\Sigma_F^*, \Sigma_S^*) \in \arg\max_{(\Sigma_F,\Sigma_S)\subset\Psi_F\times\Psi_S} \left(1 - \prod_{\psi\in\Sigma_F\cup\Sigma_S}(1-\psi)\right)v - |\Sigma_F|c_F - |\Sigma_S|(c_{S,1} + \left(\prod_{\psi\in\Sigma_F}(1-\psi)\right)c_{S,2})$$
$$\text{s.t. } |\Sigma_F| + |\Sigma_S| \leq n.$$

It is a relaxed problem because for every strategy profile $\sigma = (\sigma_1, ..., \sigma_n)$, there is a *corresponding pair of sets*

$$(\Sigma_F^\sigma, \Sigma_S^\sigma) = (\{\psi_{a_F,i} : a_F \in \mathcal{A}_F, i \leq n_{a_F}\}, \{\psi_{a_S,i} : a_S \in \mathcal{A}_S, i \leq n_{a_S}\}) \subset \Psi_F \times \Psi_S,$$

that yields the same value in the relaxed problem, as σ in the original problem.

Step 3: (Solution to the relaxed problem must be a pair of intervals.)

However it must be that for all $\psi_F, \psi'_F \in \Psi_F$, if $\psi_F > \psi'_F$ and $\psi'_F \in \Sigma_F^*$ hold, then also $\psi_F \in \Sigma_F^*$. Otherwise one could improve on (Σ_F^*, Σ_S^*) by removing ψ'_F from Σ_F^* and adding ψ_F instead. In the same way, it must be that for all $\psi_S, \psi'_S \in \Psi_S$, if $\psi_S > \psi'_S$ and $\psi'_S \in \Sigma_S^*$ hold, then also $\psi_S \in \Sigma_S^*$. Otherwise one could improve on (Σ_F^*, Σ_S^*) by removing ψ'_S from Σ_S^* and adding ψ_S instead.

This means that every solution (Σ_F^*, Σ_S^*) is a pair of *intervals*, that is,

$$\text{for all } \psi_F^* \in \Sigma_F^* \text{ and all } \psi_F \in \Psi_F \backslash \Sigma_F^*, \text{ it holds that } \psi_F^* \geq \psi_F, \text{ and}$$
$$\text{for all } \psi_S^* \in \Sigma_S^* \text{ and all } \psi_S \in \Psi_S \backslash \Sigma_S^*, \text{ it holds that } \psi_S^* \geq \psi_S.$$

Step 4: (The iterative greedy algorithm produces a pair of intervals in the relaxed problem.)

The expected net marginal social benefit of an additional agent using strategy $a \in \mathcal{A}$ given some strategy profile σ is

$$MB_a(\sigma) = \begin{cases} \psi_{a,n_a+1}(1-P_F)(1-P_S)v - c_F + \psi_{a,n_a+1}(1-P_F)n_S \cdot c_{S,2} & \text{if } a \in \mathcal{A}_F, \\ \psi_{a,n_a+1}(1-P_F)(1-P_S)v - c_{S,1} - (1-P_F)c_{S,2} & \text{if } a \in \mathcal{A}_S, \end{cases}$$

which is maximized when selecting a such that ψ_{a,n_a+1} is maximal. Hence the iterative greedy algorithm will always produce a strategy profile σ^* such that, for some $\eta, \kappa \in \mathbb{N}$, with $\eta + \kappa \leq n$, the corresponding pair of sets $(\Sigma_F^{\sigma^*}, \Sigma_S^{\sigma^*})$ are such that $\Sigma_F^{\sigma^*}$ contains the largest η elements of Ψ_F, and $\Sigma_S^{\sigma^*}$ contains the largest κ elements of Ψ_S.

Step 5: (No pair of intervals improves on the one the iterative greedy algorithm produces.)

For $i = 0, ..., n-1$, in Step 2 of the algorithm, it iteratively considers strategy profiles $\sigma^{i,*}$ with corresponding pairs of sets (Σ_F^i, Σ_S^i), where the sets $\Sigma_F^0, ..., \Sigma_F^{n-1} \subset \Psi_F$ are intervals that contain zero to $n-1$ elements respectively. Before, for every i and the corresponding Σ_F^i, in Step 1 (a) to (d), the greedy algorithm is used to select σ^{i*}, so that according to

Proposition 1, for Σ_S^i it holds that

$$\Sigma_S^i \in \arg\max_{\Sigma_S \subset \Psi_S} \left(1 - (1 - P_F^i) \prod_{\psi_S \in \Sigma_S} (1 - \psi_S)\right) v - |\Sigma_S|(c_{S,1} + (1 - P_F^i) c_{S,2}) \quad \text{s.t. } |\Sigma_S| \leq n - i.$$

Moreover, in Step 3 of the iterative greedy algorithm, for $i \in \{0, ..., n-1\}$, it selects as σ^* the strategy profile that corresponds to the pair $(\Sigma_S^*, \Sigma_F^*) \in \{(\Sigma_S^0, \Sigma_F^0), ..., (\Sigma_S^i, \Sigma_F^i)\}$ that yields the highest value in the relaxed problem. Finally, for $n-1$, Step 6 is reached, where also the case where Σ_F^n is an interval that contains n elements and Σ_S^n is the empty set is considered. At this point, the iterative greedy algorithm selects σ^* so that the corresponding pair (Σ_S^*, Σ_F^*) maximizes the value of the relaxed problem for all pairs of intervals in $\{(\Sigma_S^0, \Sigma_F^0), ..., (\Sigma_S^n, \Sigma_F^n)\}$, and hence for all pairs of intervals (Σ_F, Σ_S) with $|\Sigma_F| + |\Sigma_S| \leq n$.

Because the solution to the relaxed problem must be a pair of intervals, the pair of intervals (Σ_S^*, Σ_F^*) that corresponds to the strategy profile σ^* selected by the iterated greedy algorithm is a solution to the relaxed problem. Therefore, σ^* itself must solve the original optimization problem of the social planner. $\square$

Proof of Proposition 4

Denote the probability that some approach other than a is feasible by

$$P_{-a} := 1 - \prod_{a' \neq a} (1 - P_{a'}).$$

Since $P_a < 1$ for all $a \in \mathcal{A}$, it holds that $P_{-a} > 0$. Clearly, $(1 - P_{-a})$ constitutes a lower bound for the probability that no success is obtained on any other approach than a.

Step 1: (Never a profitable deviation to $a = \emptyset$ for any v)

Consider an arbitrary agent $i \in \mathcal{N}$.

If $\sigma_i = \emptyset$, then i's expected payoff is zero. If $\sigma_i^* = a \in \mathcal{A}_\mathcal{F}$, then the symmetry of all agents following a implies that i's expected payoff is

$$\bar{w}_a \frac{P_a \left(1 - (1 - \lambda_{a,1})^{n_a^*}\right)}{n_a^*} - c_a = \varepsilon > 0.$$

Likewise, if $\sigma_i^* = a \in \mathcal{A}_\mathcal{S}$, the symmetry of all agents following a implies that i's expected payoff is

$$\bar{w}_a \frac{(1-\pi_F^*)P_a\left(1-(1-\lambda_{a,2})^{n_a^*}\right)}{n_a^*} - c_{a,1} - (1-\pi_F^*)c_{a,2} = \varepsilon > 0.$$

Hence, $\sigma_i = \emptyset$ is never a profitable deviation.

Step 2: (Never a profitable deviation to $a \in A_F$ for any v)

This Step is practically identical to Step 1 in the proof of Proposition 2.

Recall that $\psi_{a,n_a^*} = \lambda_a \frac{P_a(1-\lambda_a)^{n_a-1}}{P_a(1-\lambda_a)^{n_a-1}+1-P_a}$, where $\lambda_a = \lambda_{a,1}$ if $a \in a_F$, and $\lambda_a = \lambda_{a,2}$ if $a \in a_S$.

The expected payoff of an agent from deviating to $\sigma_i = a \in \mathcal{A}_{\mathcal{F}}$ is, again by symmetry among the then $n_a + 1$ agents following a,

$$\begin{aligned}
\mathbb{E}\left[w_i - c_{a,1}\right] &= \bar{w}_a \frac{P_a\left(1-(1-\lambda_{a,1})^{n_a^*} + \lambda_{a,1}(1-\lambda_{a,1})^{n_a^*}\right)}{n_a^*+1} - c_{a,1} \\
&= \frac{\bar{w}_a}{n_a^*+1}\mathbb{P}\left[\exists j \in \mathcal{N},\ \text{s.t. } s_j = 1 \text{ and } \alpha_j = a\right] - c_{a,1} \\
&= \frac{\bar{w}_a}{n_a^*+1}\left(1 - \prod_{i=1}^{n_a^*+1}(1-\psi_{a,i})\right) - c_{a,1} \\
&= \frac{n_a^*}{n_a^*+1} \cdot \frac{1-\prod_{i=1}^{n_a^*+1}(1-\psi_{a,i})}{1-\prod_{i=1}^{n_a^*}(1-\psi_{a,i})}(c_{a,1}+\varepsilon) - c_{a,1} \\
&= \underbrace{\frac{n_a^*}{n_a^*+1} \cdot \frac{\psi_{a,1} + \sum_{i=2}^{n_a^*+1}\left[\prod_{j=1}^{i-1}(1-\psi_{a,j})\right]\psi_{a,i}}{\psi_{a,1} + \sum_{i=2}^{n_a^*}\left[\prod_{j=1}^{i-1}(1-\psi_{a,j})\right]\psi_{a,i}}}_{:=\gamma_a}(c_{a,1}+\varepsilon) - c_{a,1}.
\end{aligned}$$

It holds that $\psi_{a,1} = \phi_a > 0$, and $\psi_{a,i} \geq \psi_{a,i+1}$ for all $a \in \mathcal{A}$ and all $i < n$. This implies that

$$\psi_{a,1} > (1-\psi_{a,1})\psi_{a,2} \geq \prod_{j=1}^{2}(1-\psi_{a,j})\psi_{a,3} \geq \ldots \geq \prod_{j=1}^{n_a^*}(1-\psi_{a,j})\psi_{a,n_a^*+1}.$$

Thus it follows that

$$\frac{\psi_{a,1} + \sum_{i=2}^{n_a^*}\left[\prod_{j=1}^{i-1}(1-\psi_{a,j})\right]\psi_{a,i}}{n_a^*} > \frac{\psi_{a,1} + \sum_{i=2}^{n_a^*+1}\left[\prod_{j=1}^{i-1}(1-\psi_{a,j})\right]\psi_{a,i}}{n_a^*+1},$$

and hence $\gamma_a < 1$.

Thus for all $a \in \mathcal{A}_{\mathcal{F}}$, and all $\varepsilon < \frac{1-\gamma_a}{\gamma_a} c_{a,1}$,

$$\mathbb{E}\left[w_i - c_{a,1}\right] = (\gamma_a - 1)c_{a,1} + \gamma_a \varepsilon < 0.$$

Hence, for $\varepsilon < \gamma_F \coloneqq \min_{a \in A_F}\{\frac{1-\gamma_a}{\gamma_a} c_{a,1}\}$, $\sigma_i = a \in A_F$ is not a profitable deviation.

Step 3: (No profitable deviation to $a \in A_S$, for v large enough)

The expected payoff from deviating to $a \in A_S$ is maximal for an agent following the approach $\hat{a} = \arg\max_{a \in A_F} \psi_{a,n_a^*}$. Consider an agent i with $\sigma_i^* = \hat{a}$. The payoff from deviating to $\sigma_i = a \in \mathcal{A}_{\mathcal{S}}$ is, by symmetry among the then $n_a^* + 1$ agents following a,

$$\mathbb{E}\left[w_i\right] - c_{a,1} - \frac{1-\pi_F^*}{1-\psi_{\hat{a},n_{\hat{a}^*}}} c_{a,2} = \frac{1-\pi_F^*}{1-\psi_{\hat{a},n_{\hat{a}^*}}} \cdot \frac{\bar{w}_a}{n_a^*+1}\left(1 - \prod_{i=1}^{n_a^*+1}(1-\psi_{a,i})\right) - c_{a,1} - \frac{1-\pi_F^*}{1-\psi_{\hat{a},n_{\hat{a}^*}}} c_{a,2}$$
$$= \frac{1}{1-\psi_{\hat{a},n_{\hat{a}^*}}} \cdot \frac{n_a^*}{n_a^*+1} \cdot \frac{1 - \prod_{i=1}^{n_a^*}(1-\psi_{a,i}) + \psi_{a,n_a^*+1}\prod_{i=1}^{n_a^*}(1-\psi_{a,i})}{1 - \prod_{i=1}^{n_a^*}(1-\psi_{a,i})}(c_{a,1} + (1-\pi_F^*)c_{a,2} + \varepsilon) - c_{a,1} - \frac{1-\pi_F^*}{1-\psi_{\hat{a},n_{\hat{a}^*}}} c_{a,2}$$

Since $\psi_{a_F,n_{a_F}}(1-\lambda_{a_F}) < \psi_{a_F,n_{a_F}+1}$, and for all $a_F \in A_F$, and

$$(1-P_{-a_F})(1-P_{a_F} + P_{a_F}(1-\lambda_{a_F})^{n^*_{a_F}})\psi_{a_F,n^*_{a_F}+1} \cdot v \le c_{a_F,1},$$

For all $a_F \in A_F$, it also holds that

$$\psi_{a_F,n^*_{a_F}} v < \frac{c_{a_F,1}}{(1-\lambda_{a_F})(1-P_{a_F})(1-P_{-a_F})}.$$

Define

$$k_F \coloneqq \max_{a_F \in A_F} \frac{c_{a_F,1}}{(1-\lambda_{a_F})(1-P_{a_F})(1-P_{-a_F})},$$

then clearly for $v > k_F$,

$$\frac{1}{1-\psi_{\hat{a},n^*_{\hat{a}}}} < \frac{v}{v-k_F}.$$

Similarly for all $a_S \in A_S$, it holds that $\psi_{a_S,n^*_{a_S}} v < \frac{c_{a_S,1}+c_{a_S,2}}{(1-\lambda_{a_S})(1-P_{a_S})(1-P_{-a_S})}$. Define

$$k_{a_S} \coloneqq \max_{a_S \in A_S} \frac{c_{a_S,1}+c_{a_S,2}}{(1-\lambda_{a_S})(1-P_{a_S})(1-P_{-a_S})}.$$

It follows that

$$\frac{1 - \prod_{i=1}^{n_a^*}(1-\psi_{a,i}) + \psi_{a,n_a^*+1}\prod_{i=1}^{n_a^*}(1-\psi_{a,i})}{1 - \prod_{i=1}^{n_a^*}(1-\psi_{a,i})} < \frac{v + \frac{k_a}{1-P_a}}{v}.$$

Moreover, to derive an upper bound on $n^*_{a_S}$, consider the following inequality which must be fulfilled if $n^*_{a_S} > 0$:

$$P_{a_S}\lambda_{a_S}(1-\lambda_{a_S})^{n^*_{a_S}-1} \cdot v \geq c_{a_S,1}$$
$$\iff n^*_{a_S} \leq \frac{\log\left(\frac{(1-\lambda_{a_S})c_{a_S,1}}{P_{a_S}\lambda_{a_S}\cdot v}\right)}{\log(1-\lambda_{a_S})} = \frac{\log\left(\frac{(1-\lambda_{a_S})c_{a_S,1}}{P_{a_S}\lambda_{a_S}}\right) - \log v}{\log(1-\lambda_{a_S})}.$$

Combining the above observations, the following upper bound for the expected payoff of deviating to a holds:

$$\begin{aligned}
\mathbb{E}\left[w_i\right] &- c_{a,1} - \frac{1-\pi^*_F}{1-\psi_{\hat{a},n_{\hat{a}^*}}}c_{a,2} \\
&< \frac{v}{v-k_F} \cdot \frac{\log v - \log\left(\frac{(1-\lambda_a)c_{a,1}}{P_a\lambda_a}\right)}{\log v - \log\left(\frac{(1-\lambda_a)^2 c_{a,1}}{P_a\lambda_a}\right)} \cdot \frac{v+\frac{k_a}{1-P_a}}{v}(c_{a,1}+(1-\pi^*_F)c_{a,2}+\varepsilon) - c_{a,1} - \frac{1-\pi^*_F}{1-\psi_{\hat{a},n_{\hat{a}^*}}}c_{a,2} \\
&= \underbrace{\frac{v+\frac{k_a}{1-P_a}}{v-k_F} \cdot \frac{\log v - \log\left(\frac{(1-\lambda_a)c_{a,1}}{P_a\lambda_a}\right)}{\log v - \log\left(\frac{(1-\lambda_a)^2 c_{a,1}}{P_a\lambda_a}\right)}}_{:=\gamma_a}(c_{a,1}+(1-\pi^*_F)c_{a,2}+\varepsilon) - c_{a,1} - \frac{1-\pi^*_F}{1-\psi_{\hat{a},n_{\hat{a}^*}}}c_{a,2}.
\end{aligned}$$

For v sufficiently large, it holds that $\frac{v+\frac{k_a}{1-P_a}}{v-k_F} \cdot \frac{\log v - \log\left(\frac{(1-\lambda_a)c_{a,1}}{P_a\lambda_a}\right)}{\log v - \log\left(\frac{(1-\lambda_a)^2 c_{a,1}}{P_a\lambda_a}\right)} < 1$.

Take $\hat{v}$ such that the above inequality is satisfied, and define $\hat{n} := \sum_{a\in\mathcal{A}} \frac{\log\left(\frac{(1-\lambda_a)c_{a,1}}{P_a\lambda_a\cdot\hat{v}}\right)}{\log(1-\lambda_a)}$. Thus, for all $a \in \mathcal{A_S}$ and all $\varepsilon < \frac{1-\gamma_a}{\gamma_a}c_{a,1}$,

$$\mathbb{E}\left[w_i\right] - c_{a,1} - \frac{1-\pi^*_F}{1-\psi_{\hat{a},n_{\hat{a}^*}}}c_{a,2} < (\gamma_a - 1)c_{a,1} + \gamma_a\varepsilon < 0.$$

Hence, for $\varepsilon < \gamma_F := \min_{a\in A_S}\{\frac{1-\gamma_a}{\gamma_a}c_{a,1}\}$, $\sigma_i = a \in A_S$ is not a profitable deviation.

In Conclusion, for all $v > \hat{v}$, and $n > \hat{n}$, there exists $\gamma := \min\{\gamma_F, \gamma_S\}$, such that for $\varepsilon < \gamma$ there is no profitable deviation, and hence σ^* is a Nash equilibrium. $\square$