

Diversification and Information in Contests*

Jorge Lemus[†] and Emil Temnyalov[‡]

March 17, 2022

Abstract

We study contests with technological uncertainty, where contestants can invest in different technologies of uncertain value. The principal can disclose an informative yet noisy public signal about the merit of each technology, which can focus investments into more promising ones or increase diversification. We characterize the principal's optimal disclosure of information about the technologies, which depends on her value of diversification, the informativeness of available signals, and the asymmetry of technologies. We also show that a larger prize and more contestants typically decrease the extent of information disclosure.

JEL codes: O32, C72, D62, D72, D83.

Keywords: contests, diversification, information.

*For helpful comments we thank seminar and conference participants at APIOC 2018, IIOC 2019, EARIE 2019, Jornadas de Economía Industrial (2019), SAET 2021, CEMFI/UC3M, the University of Melbourne, the University of Tokyo and the University of Toronto Rotman School of Business.

[†]University of Illinois Urbana-Champaign, Department of Economics. jalemus@illinois.edu

[‡]University Technology Sydney, Business School. emil.temnyalov@uts.edu.au

1 Introduction

Many contests feature technological uncertainty: there are multiple technologies, approaches, or methods that contestants could pursue, and it is not clear, both for the contest sponsor and the contestants, which of these will ultimately be successful. For instance, in prediction competitions hosted either on digital platforms (e.g., Kaggle) or by a firm (e.g., Netflix), there are multiple algorithms that contestants can use (e.g. machine learning or regression methods). Contestants can use different farming methods in agricultural yield contests, such as those sponsored by the National Corn Growers Association. In these settings, contestants do not know which method will deliver the best results *ex post*. Technological uncertainty is also present in procurement contests, such as those hosted by the U.S. Department of Defense, where contestants may not fully know the procurer’s preferences over different possible attributes or designs. Similarly, in contest-like settings within organizations, competing workers face uncertainty regarding the impact of different tasks on their promotion chances.

Both uncertainty and competitive pressure affect how much contestants (the ‘agents’) invest in each technology. Uncertainty means that some agents will inevitably invest in unsuccessful technologies. For this reason, *diversifying* agents’ investments across multiple technologies can improve the value that the contest sponsor (the ‘principal’) obtains in the contest. However, diversifying investments can also reduce the total investment on the *ex-post* best technology, reducing the principal’s payoff from implementing it. In such a setting, the principal, who is also uncertain about the value of each technology, may want to disclose information that can change the agents’ beliefs about the prospects of these alternatives, as their beliefs in turn determine their investments.

Public information revelation is a feature of many contest settings. For instance, in prediction contests, the principal provides data that contestants use to train and develop their algorithms; these data reveal information about the probability of success of different algorithms. In yield contests, the principal could reveal or hide the performance of farming methods used in the past. The National Wheat Foundation, for example, publicly distributes information on different farming methods. In procurement contests, prototype evaluations publicly revealed to contestants are informative about the procurer’s preferences. In all these cases, the principal can *strategically* choose what data to make public to influence the contestants’ investment decisions. In general, the

principal wants to avoid “over-fitting,” i.e., selecting a method that performs well on one dataset but not in general. Instead, she seeks to incentivize the contestants to invest in valuable technologies beyond the existing data. Revealing all the historical data can bias agents towards an approach that worked in the past but might not work well in the future. Similarly, rating prototypes can affect the agents’ beliefs about the likelihood of success of each design.

We investigate the contest-design question of how much information about the technologies the principal should reveal. Revealing more information could bias agents to invest too much on some technologies and induce too little diversification—i.e., too little investment across technologies that do not seem promising in the current data but could be the best method *ex post*. This information design aspect applies to a number of problems in innovation, procurement, and organizations, where technological uncertainty is important. For instance, in prediction contests the contestants develop algorithms without knowing if their chosen method is the best; but they can infer the performance of different alternatives by testing them on a public dataset determined by the contest designer.¹ This was also the case in Netflix’s recommendation competition in 2009, where Netflix publicly revealed a subset of all its available data on users’ preferences to allow contestants to evaluate the potential of different prediction algorithms. From Netflix’s perspective there is an option value to procuring a number of different algorithms because it is uncertain which one will be the most valuable in the long run.²

We introduce a model to study technological uncertainty in contests and characterize when the principal should reveal information to maximize the value of the contest. Depending on the prior belief regarding the technologies, optimal information disclosure can steer agents to focus on one specific technology or to diversify their investments across technologies.

In our setting, M agents choose how much to invest in each of N different technologies.³

¹Online platforms like Kaggle, DrivenData, and crowdAI, among others, offer firms the possibility to sponsor contests to outsource their data science needs. The contest designer partitions the available data into a “test dataset” and an “evaluation dataset.” The test dataset is publicly available and allows participants to test their algorithms. The prize is allocated based on the performance of the algorithm over the evaluation dataset.

²In addition to whichever algorithm performs best given their *current data*, the best algorithms in the long run will depend on Netflix’s future data on users’ preferences.

³These different technologies may represent different approaches to solve a problem in the case of innovation; different characteristics, features, or designs in the case of procurement; or different tasks or projects in the case of a worker competing within an organization.

Only one of these technologies is valuable for the principal—the “best” technology—and it is the one that she will implement *ex post*. The contest winner is selected probabilistically among the agents who invested in the best technology: the more an agent invests in the best technology, the higher the chances of winning the prize. We model this competition as N parallel Tullock contests, one for each technology.

The principal designs a public “experiment,” in the Bayesian persuasion sense of the term, to reveal information to the agents regarding the likelihood-of-success of each technology. In prediction contests, for instance, this experiment corresponds to the public data made available to the participants, who can then test different classes of algorithms on the data and update their beliefs about the possible value of each approach.

The main trade-off in choosing the information structure that maximizes the principal’s expected payoff is one between diversification and focus: revealing more precise information induces agents to focus on more promising technologies, but if the agents rationally over-react to such information in equilibrium, then this may lead to too little diversification from the principal’s perspective. As a result of this trade-off, we characterize conditions under which the optimal policy is maximally informative, partially informative, or completely uninformative, depending on the features of the environment. Whether the principal wants to reveal or hide information depends critically on three factors: (i) the value of technological diversification; (ii) the quality of the principal’s information; and (iii) the extent of technological uncertainty.

First, investing in multiple technologies generates diversification: even if one technology looks more promising than another *ex ante*, the latter may turn out to be more valuable in the long run. The larger this value of diversification, which turns out to be related to a measure of risk aversion associated with the principal’s objective function, the more likely it is that the principal chooses not to reveal information.

Second, the quality of the principal’s information matters: if she can design a very informative experiment, *i.e.*, an experiment that reveals the best technology with very high probability, there is little justification to withhold revealing that information. In practice, however, the principal may not have access to very informative signals. For instance, in a prediction contest the principal may have a limited amount of data to provide to the contestants.

Third, the extent of technological uncertainty reflects how asymmetric the different approaches are a priori. If the agents' beliefs about the technologies are ex ante very asymmetric, the principal may want to reveal information to either reinforce or weaken the extent of this asymmetry. The more symmetric the technologies are ex ante, the more diversification there is without disclosing information.

We also study the relationship between information design and more traditional contest design aspects, such as the number of competitors and the size of the prize. Under mild conditions, we find that the optimal design features less information disclosure in contests with more competitors or larger prizes. This is because when the prize is larger or more agents compete, the aggregate investment on each technology is larger. Hence, the value of disclosing information to focus investment towards one technology is generally lower than the value of diversification.

Lastly, we explore the interaction between information disclosure and the optimal size of the prize. We characterize conditions under which information design generally adds value above and beyond what can be achieved by choosing the prize optimally and not disclosing information. This is because the size of the prize affects the levels of investments in the contest, whereas information disclosure affects their direction—i.e. the relative investments across different technologies. Hence, the latter can yield investment allocations that the former cannot necessarily achieve. We provide a condition under which, when optimizing over both information disclosure and prizes, the optimal solution entails a smaller prize and more information disclosure. In other words, in this case information and prizes are substitutes: the principal may be better off by offering a smaller prize and optimally revealing information, instead of offering a larger prize and not revealing information.

Related Literature. Our paper contributes to the recent literature on diversification in contests. Preferences for diversification has been studied theoretically (e.g., [Terwiesch and Xu, 2008](#)) and empirically (e.g. [Boudreau et al., 2011](#)). [Letina and Schmutzler \(2019\)](#) characterize the optimal prize structure when the designer wants to induce a variety of approaches. Rather than using an optimal prize structure to induce variety, we focus on optimally disclosing information, which can induce variety, so our results are complementary to theirs.

Most of the existing literature on information disclosure in contest focuses on disclosing

information about the *agents' characteristics* or *actions*. For instance, [Aoyagi \(2010\)](#) compares equilibria in contests where agents' *performance* is fully disclosed and contests where no information is provided. [Ederer \(2010\)](#) adds private information to this setting, and [Klein and Schmutzler \(2016\)](#) add allocation of prizes and alternative performance evaluations. Performance feedback in dynamic contests has been theoretically studied by [Bimpikis et al. \(2014\)](#) and [Benkert and Letina \(2016\)](#), and empirically by [Gross \(2015, 2017\)](#), [Huang et al. \(2014\)](#), [Kireyev \(2016\)](#), [Bockstedt et al. \(2016\)](#), and [Lemus and Marshall \(2021\)](#). [Fu et al. \(2016\)](#) and [Xin and Lu \(2016\)](#) study optimal information disclosure regarding agents' *entry decisions* in contests. [Zhang and Zhou \(2016\)](#) and [Serena \(2021\)](#) study information disclosure regarding players' *effort costs*. [Mihm and Schlapp \(2018\)](#) study the optimal information disclosure when players are uncertain about the *principal's preferences*. [Kovenock et al. \(2015\)](#) study the effect of players sharing information throughout the contest.

None of the papers above study a setting where the contestants and the contest designer are uncertain about the value of the underlying *technologies*, which we focus on in this paper. Thus we contribute to the literature by studying a novel setting where information design can improve contest outcomes.

Also related are articles studying diversification in settings other than contests. For instance, [Letina \(2016\)](#) studies the effect of market competition and mergers on variety and finds conditions such that the research portfolio under market competition features too much (or too little) variety. [Toh and Kim \(2013\)](#) study how aggregate uncertainty affects technological diversification within a firm, finding that a firm's technology specializes more under greater uncertainty. Relatedly, [Krishnan and Bhattacharya \(2002\)](#) study how a firm should design a product when there are several uncertain alternatives for the product's underlying technology.

Our paper also relates to R&D models with multiple risky technologies. [Dasgupta and Maskin \(1987\)](#) show that, in a winner-takes-all competition, the equilibrium allocation of research on correlated projects is too high relative to the socially efficient allocation, so there is less diversification in equilibrium. [Bhattacharya and Mookherjee \(1986\)](#) present a similar framework, but they study the level of risk taken by the firms, finding that the optimal research strategy may feature excessive or insufficient risk taking, depending on the level of risk aversion and the shape of the distribution of research outcomes. [Cabral \(1994\)](#) shows that, when the competition is not winner-takes-all,

the level of risk taking is lower than the socially optimal level. [Cabral \(2003\)](#) explores the same question in a dynamic environment, showing that a follower firm takes more risk than the leader. Our paper contributes to the literature by studying information disclosure by the contest designer in the framework of [Kamenica and Gentzkow \(2011\)](#).

A recent literature has explored information design in games more generally. This work includes [Mathevet et al. \(2017\)](#), [Laclau and Renou \(2016\)](#), [Alonso and Câmara \(2016\)](#), and [Ederer et al. \(2018\)](#), among others, as well as models of contests for experimentation, such as [Halac et al. \(2017\)](#).

Our paper mostly focuses on a single element of contest design, information disclosure, which relates broadly to the literature of contest design, where the goal is to study the effect of alternative designs on players' incentives. This literature includes the work of [Taylor \(1995\)](#) and [Fullerton and McAfee \(1999\)](#) on restricting the number of competitors in winner-takes-all tournaments, [Moldovanu and Sela \(2001\)](#) on the optimal number of prizes, and [Che and Gale \(2003\)](#) on both the number of prizes and of participants.

2 Model

A principal organizes a contest to solve a problem. There are M agents, indexed by $i \in \{1, \dots, M\}$, who compete in the contest. There are N technologies, indexed by $t \in \{1, \dots, N\}$, that the agents can use to tackle the problem that needs to be solved.

It is common knowledge that only one of these technologies will be valuable to the principal, but it is uncertain ex ante which technology that is: we call this technology the “best” one ex post. The agents and principal hold a common prior, $\Theta \in \Delta^N$, where

$$\Delta^N \equiv \{(\theta_1, \dots, \theta_N) : \theta_t \in [0, 1], \sum_{t=1}^N \theta_t = 1\}$$

and θ_t is the common prior belief that technology t is the best one, with $t = 1, \dots, N$. Agents choose how much to invest into each technology, and the contest awards a prize, V , to one of the agents who invested in the best technology.

We model competition within each technology as a standard Tullock contest (see, e.g., [Pérez-Castrillo and Verdier, 1992](#); [Konrad, 2009](#)). Agent i wins the contest for technol-

ogy t with probability

$$p_i(x_{i,t}, x_{-i,t}) = \frac{x_{i,t}}{\sum_{j=1}^M x_{j,t}},$$

where $x_{i,t}$ is agent i 's investment in technology t and $x_{-i,t}$ is the investment profile of agent i 's rivals on technology t .⁴

The principal's payoff depends on the aggregate investment, $x_t = \sum_{i=1}^M x_{i,t}$, on the ex post best technology, t . Specifically, the principal's payoff is

$$v(x) = \sum_{t=1}^N f(x_t) \cdot I(t \text{ is the best technology}),$$

where $f(\cdot)$ is a strictly increasing, twice differentiable, and concave production function and $I(\cdot)$ is the indicator function. In words, the principal derives no benefit from investments on technologies other than the best one. A larger aggregate investment on the best technology increases the principal's payoff, but there are decreasing marginal returns.⁵

For any possible belief $\Theta = (\theta_1, \dots, \theta_N) \in \Delta^N$, where $\theta_t = \Pr(t \text{ is the best technology})$, and any aggregate investment profile (x_1, \dots, x_N) , the principal's expected payoff is

$$E_{\Theta}[v(x)] = \sum_{t=1}^N \theta_t f(x_t). \quad (1)$$

Note that, whenever there is residual uncertainty, i.e., $\theta_t < 1$ for all $t = 1, \dots, N$, the concavity of $f(\cdot)$ means that the principal values *diversification* among the technologies.

Modeling Assumptions. The principal is risk neutral over output profiles, $(f(x_1), \dots, f(x_N))$. However, when we consider expected payoffs over the profile of total investment on each technology, (x_1, \dots, x_N) , the principal's preferences over such profiles are risk averse due to the concavity of $f(\cdot)$. The latter interpretation—with risk aversion over investment profiles—will turn out to be insightful because the Arrow-Pratt coefficient of relative risk aversion associated with $f(\cdot)$ plays an important role in our findings. Hence, we will refer to the principal as being risk averse throughout the paper, and we will treat

⁴We make the standard assumption in the contest literature: $p_i(0, \dots, 0) = 1/N$.

⁵In Appendix B, we provide two models which induce such increasing and concave payoffs for the principal that are a function of aggregate investment, and which are also consistent with the winning probabilities of a Tullock contest. In the main text we work directly with the reduced-form representation given by the Tullock winning probabilities and the principal's payoff function above.

the principal’s payoffs as defined over investment profiles rather than over the outputs that they correspond to.

We assume agents compete in a contest with a simple prize structure: a single prize of size V . This prize structure is commonly used in practice and is optimal in many contest settings (see, e.g., [Clark and Riis, 1998](#)). Furthermore, we model the prize as technology-neutral. In practice it may be difficult to contract over technology-specific prizes. In fact, virtually no contest in practice uses technology-contingent prizes. In theory one could imagine a contest where the prize depends on which technology the winner uses. This is weakly better than a single prize, but the analysis of this case is less tractable. We study this as an extension in the Appendix; for the sake of tractability, our baseline model uses a prize that is technology-neutral.

2.1 Preliminary Analysis

Consider agent i ’s problem of choosing how much to invest in technology t . Given any common belief $\Theta = (\theta_1, \dots, \theta_N)$ and the aggregate investment of agent i ’s rivals on technology t , agent i ’s optimal investment profile solves the following problem:

$$\max_{x_{i,t} \geq 0} \sum_{t=1}^N V \cdot \theta_t \cdot p_i(x_{i,t}, x_{-i,t}) - x_{i,t},$$

Agent i receives the prize, V , if and only if: (1) technology t is the best approach—under the agent’s belief, this event happens with probability θ_t —and (2) agent i wins the contest for technology t , which happens with probability $p_i(x_{i,t}, x_{-i,t})$.

In contrast, in the first-best allocation—if the principal could control the agent’s investments—the investment on technology t is found by solving:

$$\max_{y_t \geq 0} \sum_{t=1}^N \theta_t f(y_t) - y_t.$$

Proposition 1. *We have:*

1. In the contest's unique equilibrium, agent i 's investment on technology t is

$$x_{i,t}^* = \theta_t \left(\frac{M-1}{M^2} \right) V.$$

2. The aggregate equilibrium investment on technology t is

$$x_t^* \equiv \sum_{i=1}^M x_{i,t}^* = \theta_t \Omega, \quad \text{where} \quad \Omega = \left(\frac{M-1}{M} \right) V. \quad (2)$$

3. In the first-best allocation, the investment on technology t is

$$y_t^* = \begin{cases} (f')^{-1} \left(\frac{1}{\theta_t} \right) & \text{if } \theta_t f'(0) > 1 \\ 0 & \text{otherwise} \end{cases}$$

From [Proposition 1](#), the principal's expected equilibrium payoff from a contest in which agents compete under belief $(\theta_1, \dots, \theta_N)$ is

$$E_{\Theta}[v(x^*)] = \sum_{t=1}^N \theta_t f(x_t^*) - V = \sum_{t=1}^N \theta_t f(\theta_t \Omega) - V. \quad (3)$$

The agents' and the principal's incentives are generally misaligned. The equilibrium and first-best allocations differ in terms of both levels and directions. In terms of absolute levels, it is easy to see from [Proposition 1](#) that the equilibrium can generally feature both over-investment and under-investment. For example, if V is large enough, we can have that $x_t^* > y_t^*$ for all t , i.e., agents invest too much across all technologies relative to the first-best. The opposite happens when V is small enough. More generally, there can also simultaneously be under-investment and over-investment across different technologies, for other levels of V .⁶

Another question is what *proportion* of the total investment is allocated to different technologies. That is, which technologies will agents invest too much or too little in, as a function of their beliefs Θ , compared to the relative investment shares in the first-best allocation? This comparison turns out to be more relevant for our subsequent results, when we consider how the principal can strategically manipulate beliefs through information disclosure, as the responsiveness of agents' equilibrium investments with

⁶We consider the prize that maximizes the principal's payoff later, in sections [2.3](#) and [2.4](#).

respect to their beliefs plays a key role in that analysis. Specifically, the proportion of total investment allocated towards technology t in equilibrium is

$$\frac{x_t^*}{\sum_{t=1}^N x_t^*} = \theta_t.$$

This is generically different than the first-best proportion of total investment allocated towards technology t , which is

$$\frac{y_t^*}{\sum_{t=1}^N y_t^*} = \frac{(f')^{-1}\left(\frac{1}{\theta_t}\right)}{\sum_{t=1}^N (f')^{-1}\left(\frac{1}{\theta_t}\right)}.$$

Figure 1 compares the proportion of total investment allocated towards technology 1 in equilibrium and in the first best, with two technologies ($N = 2$). Figure 1 (left panel) shows that, when $f(x) = \sqrt{x}$, agents *under-react* to their beliefs regarding the technologies relative to the first-best: when $\theta > 0.5$, the proportion of total investment allocated in equilibrium towards technology 1 (dashed line) is *lower* relative to the first best (solid line). Analogously, when $\theta < 0.5$, the proportion of total investment allocated to technology 1 in equilibrium (dashed line) is *larger* relative to the first best (solid line). Therefore, in equilibrium the proportion of total investment allocated towards the most promising technology is lower than the first-best proportion.

Figure 1 (right panel) considers $f(x) = 1 - \exp(-8x)$ and shows that, in this case, agents can both *over-react* and *under-react* to their beliefs regarding the technologies, depending on the range of beliefs considered. When the belief is around 0.5, agents over-react to beliefs, in contrast to the left panel of the figure. In particular, when $\theta > 0.5$ (respectively, $\theta < 0.5$) and not too far from 0.5, the proportion of total investment allocated in equilibrium towards technology 1 is larger (respectively, lower) relative to the first best. Hence the equilibrium proportion of total investment allocated to the more promising technology is too large relative to first-best when the beliefs are not too extreme, i.e., close enough to 0.5. On the other hand, when beliefs are far enough from 0.5, agents in equilibrium under-react to their beliefs regarding the technologies relative to the first best allocation, as in the left panel of the figure.




Figure 1: Comparison of the proportion of total investment allocated towards technology 1 as a function of the belief θ . In the left panel, the principal’s preference is $f(x) = \sqrt{x}$, whereas in the right panel it is $f(x) = 1 - \exp(-8x)$.

The over- and under-reaction illustrated in Figure 1 can be counteracted through information design. In the next section, we explore if and when the principal can increase her payoff by disclosing some information about the merits of each technology through the results of an experiment. Such an experiment can improve the allocation of investments across the technologies. On the other hand, if agents over-react to such information, there may be too little diversification from the principal’s perspective. Thus the principal may prefer to disclose less information, potentially inducing investment misallocation in equilibrium, in exchange for more diversification.

2.2 Information Design

We now turn to the analysis of optimal information disclosure. In general, the principal can reveal information to the agents regarding the feasibility or likelihood of success of the different technologies. Importantly, it may be *impossible* for the principal to fully reveal which technology is the best: if this were possible, the principal would always choose to do so. One important feature of our analysis is that, even after information is publicly revealed, there is residual uncertainty. For example, in prediction contests there is usually a public “test dataset” that reveals the performance of different algorithms. However, there is residual uncertainty: the “best” algorithm, which determines the winner of the contest, is decided by evaluating the performance of the algorithms

using a different dataset.⁷ In that sense, the test dataset is a public signal that agents use to update their beliefs before choosing which algorithm to submit to the contest. In such settings with residual uncertainty, we ask a number of questions. Can the principal disclose information to improve the equilibrium investments across technologies? Should the principal release information at all? Would a more risk-averse principal reveal more or less information? How does competition impact the extent of information disclosure? What is the optimal size of the prize if information disclosure is taken into account?

The principal commits to an information policy *ex ante*, in a Bayesian persuasion framework, which will inform the agents about the feasibility or value of the different technologies. After observing the information revealed in this experiment, each agent chooses how much to invest in the technologies. Once each agent has made their choice, the contest resolves all remaining uncertainty and prizes are allocated. We model public rather than private information disclosure for two reasons. First, information disclosure is typically public in many of the applications of our model—in prediction contests the principal reveals the same dataset to all contestants, who then use the data to form beliefs over the value of different approaches and to develop their prediction algorithms; in procurement contests the principal uses the same procurement tender to solicit proposals from the contestants. Second, disclosing information privately would entail differential treatment of contest participants, which could raise fairness and corruption concerns in some settings, such as procurement.

Information structures. The principal designs a public signal structure $s = (\mathcal{M}, \tilde{G}(\cdot|t))$, where \mathcal{M} is a set of messages and $\tilde{G}(m|t)$ is the probability that message $m \in \mathcal{M}$ is sent when the state of nature is t , i.e., when technology t is the best choice *ex post*. Let S be the set of all such signal structures available to the principal. Importantly, as noted before, we do not assume that the principal must have access to a perfectly informative signal. Instead, we solve for the optimal information disclosure policy for *any possible* set S of available signal structures.

Each signal structure, $s \in S$, induces some distribution over posterior beliefs on the technologies, $G_s(\Theta)$, and we denote the set of posterior beliefs in the support of that

⁷In many instances of prediction contests, the best solution is one that gives the best out-of-sample predictions. Thus, to determine the winner, each algorithm could be evaluated on datasets that do not even exist when agents choose which technology to use.

signal as \mathcal{P}_s , with generic elements $\Theta \in \mathcal{P}_s$. Let \mathcal{P}_S denote the set of all feasible posterior beliefs that can be induced by some signal, including signals that are compositions and combinations of signals available to the principal.

Lemma 1. *For any set of signal structures S , the set of posteriors that can be induced by the principal, \mathcal{P}_S , is convex.*

Intuitively, for any two posterior beliefs $\tilde{\theta}, \tilde{\theta}' \in \mathcal{P}_S$ that can be induced with some signal structures, s' and s'' , the principal can also induce any belief that is a convex combination of the two posteriors, $\alpha\tilde{\theta} + (1 - \alpha)\tilde{\theta}'$. This can be achieved using an appropriate mix between the signal structures, s' and s'' , and generally requires defining a new message space and constructing a new signal structure. Therefore, the set of feasible posteriors is a convex subset of the $N - 1$ simplex. In particular, when there are only two technologies, the set of feasible posteriors is an interval.

The following example illustrates how a standard and simple signal structure translates into a set of feasible posteriors and how the latter can be derived from a given prior and a set of available signal structures.

Example 1. *There are two technologies. The state of the world, $t \in \{1, 2\}$, describes which of the two technologies is the best. Let the most informative experiment that the principal can design be one that sends signal $s \in \{1, 2\}$ according to $\Pr(s = t|t) = \alpha$, with $\alpha \in (0.5, 1)$. For a prior $\theta \in [0, 1]$, the posterior beliefs conditional on signal s are:*

$$\Pr(t = 1|s = 1) = \frac{\alpha\theta}{\alpha\theta + (1 - \alpha)(1 - \theta)} \quad \text{and} \quad \Pr(t = 1|s = 2) = \frac{(1 - \alpha)\theta}{(1 - \alpha)\theta + \alpha(1 - \theta)}. \quad (4)$$

Equation 4 generates two extreme posteriors for each prior: the boundary of \mathcal{P}_S . The entire set of feasible posteriors, \mathcal{P}_S , is convex, and is obtained by appropriately mixing the principal's most informative signal with an uninformative signal (see Lemma 1).

Figure 2 shows the range of feasible posteriors for three values of α , a parameter that represents the informativeness of the signals available to the principal. In each panel, the x-axis corresponds to the prior belief, and the y-axis corresponds to the posteriors generated by the experiment defined in Equation 4. In all panels, the vertical line corresponds to the range of feasible posteriors for the prior belief $\theta = 0.4$. With a less informative signal (e.g. $\alpha = 0.55$, Figure 2, left panel), the posteriors are very close to the prior, resulting in a posterior set narrowly concentrated around the prior. With




Figure 2: All panels show the range of posteriors for a prior of $\theta = 0.4$. Left panel: $\alpha = 0.55$, posterior in $[0.35, 0.45]$. Middle panel: $\alpha = 0.75$, posterior in $[0.18, 0.67]$. Right panel: $\alpha = 0.95$, posterior in $[0.07, 0.86]$.

more informative signals (e.g. $\alpha = 0.75$, [Figure 2](#), middle panel), the set of feasible beliefs expands around the prior. As the signal becomes almost perfectly informative (e.g. $\alpha = 0.95$, [Figure 2](#), right panel), the set of attainable posteriors expands even further. Note, however, that while the principal’s signal is informative, with some probability it will steer agents to invest in the “wrong” technology.

In our analysis we take a common prior, Θ_0 , and the set of feasible posteriors for each prior, $\mathcal{P}_S(\Theta_0) \subseteq \Delta^N$, as the primitives of the model, assuming that $\mathcal{P}_S(\Theta_0)$ is closed and convex, following [Lemma 1](#). In particular, $\Theta_0 \in \mathcal{P}_S(\Theta_0)$ because the principal can always decide not to reveal anything, so the posterior equals the prior. For the sake of notation, we denote $\mathcal{P}_S(\Theta_0)$ simply by \mathcal{P}_S , and we denote by $\partial\mathcal{P}_S$ the boundary of \mathcal{P}_S . Any signal structure can be translated into this framework. Our results then characterize the optimal information disclosure policy for arbitrary primitives.

Optimal Information Design. We analyze the value of information disclosure in terms of the posterior beliefs that a signal structure induces. Recall that when agents hold the belief Θ we have $x_t^* = \theta_t \Omega$ in equilibrium. Therefore, the principal’s expected payoff from inducing the posterior Θ is

$$\nu(\Theta) \equiv \sum_{t=1}^N \theta_t f(\theta_t \Omega). \quad (5)$$

As in [Kamenica and Gentzkow \(2011\)](#), the value of information disclosure is described by the convexity of $\nu(\Theta)$.⁸ Let $\hat{\nu}(\Theta, \mathcal{P}_S)$ be the concave closure of $\nu(\Theta)$ over \mathcal{P}_S .

⁸Note, however, that $\nu(\Theta)$ is non-linear. When posterior beliefs are distributed according to some

The principal strictly benefits from persuasion whenever $\hat{\nu}(\Theta, \mathcal{P}_S) > \nu(\Theta)$. Next, we characterize whether $\nu(\cdot)$ is concave or convex.

Lemma 2. $\nu(\Theta)$ is concave at $\Theta = (\theta_1, \dots, \theta_N) \in \Delta^N$ if and only if

$$\Omega\theta_t f'(\theta_t\Omega)[2 - r_f(\theta_t\Omega)] + \Omega\theta_N f'(\theta_N\Omega)[2 - r_f(\theta_N\Omega)] < 0, \text{ for all } t = 1, \dots, N - 1, \quad (6)$$

where $r_f(x) \equiv -\frac{xf''(x)}{f'(x)}$ is the relative risk aversion coefficient associated with $f(\cdot)$.

The necessary and sufficient condition for $\nu(\cdot)$ to be locally concave at Θ in [Lemma 2](#) depends on the Arrow-Pratt *relative risk aversion coefficient* associated with $f(\cdot)$. For instance, when $r_f(\theta\Omega) > 2$, the principal does not benefit from information disclosure; when $r_f(\theta\Omega) < 2$, the principal benefits from information disclosure.

We illustrate the relationship between the value of information disclosure and the concavity of $f(\cdot)$ for the case of two technologies. Consider $f(x) = x^a$ for any $a \in (0, 1)$; we have $r_f(x) = 1 - a < 2$, for any x . From [Lemma 2](#), $\nu(\cdot)$ is globally *convex*, so there are always gains from information disclosure for any prior and for any \mathcal{P}_S . Furthermore, the principal designs an experiment that generates the most extreme possible posteriors in \mathcal{P}_S : for any prior belief, agents receive (with some probability) the most favorable signal for one of the technologies. In other words, this signal structure induces agents to *focus* on one technology in equilibrium, reducing diversification across technologies. Next, consider $f(x) = 1 - e^{-\lambda x}$; we have $r_f(x) = \lambda x$. Depending on λ , [Lemma 2](#) shows that $\nu(\cdot)$ might not be globally convex. Hence for some beliefs and for some set of posteriors, full information revelation is not optimal, resulting in *less* focus, as the agents' posteriors would not be as extreme as under full information revelation.

distribution G , this non-linearity implies that we cannot write the principal's expected payoff as a function of one of G 's moments (e.g., posterior means). By contrast, recent Bayesian persuasion papers have studied special settings where the principal's payoff depends on *expected* posteriors—see e.g. [Kolotilin \(2018\)](#); [Hwang et al. \(2019\)](#).




Figure 3: Function $\nu(\theta)$, the value for the principal of inducing posterior θ , for different functions $f(\cdot)$. Left Panel: $f(x) = \sqrt{x}$; Right Panel: $f(x) = 1 - \exp(-8x)$.

Why would the principal want to use a more or less informative signal structure? First, consider the case of $f(x) = \sqrt{x}$. Figure 1 (left panel) shows that there is too little investment towards the most promising technology, in proportional terms, compared to the first-best allocation, i.e., there is too much diversification in equilibrium. Figure 3 (left panel) shows $\nu(\cdot)$ corresponding to $f(x) = \sqrt{x}$, which is strictly convex, so the principal’s optimal information disclosure policy always induces the most extreme posterior possible, which reduces the diversification and induces more focus in equilibrium.

Second, consider the case with $f(x) = 1 - e^{-8x}$. Figure 1 (right panel) shows that, for priors around 0.5, there is too much investment towards the most promising technology, in proportional terms, compared to the first-best allocation; there is too little diversification in equilibrium. Figure 3 (right panel) shows $\nu(\cdot)$ corresponding to $f(x) = 1 - e^{-8x}$, which is locally concave near $\theta = 0.5$ and convex towards the extremes. As we will show, the principal in this case may not disclose any information, either because the signals available are not too informative (i.e. posteriors are narrowly concentrated around the 0.5) or because generating extreme posteriors induces even less diversification. On the other hand, if the priors heavily favor one of the technologies, competition pushes agents to invest relatively more towards the less promising technology and there is too much diversification in equilibrium. To induce more focus, the principal designs an experiment that generates more extreme posteriors.

We formally show that, in general with $N \geq 2$ technologies, the gain from information disclosure (i.e. whether $\nu(\cdot)$ is locally convex or concave) depends on the value the prin-

principal assigns to diversification. When $r_f(\cdot)$ is relatively large, i.e. $f(\cdot)$ is very concave, diversification is more valuable to the principal than to the agents. This is because the equilibrium allocation for any belief Θ is too responsive to differences among the technologies relative to the principal's first-best investment across technologies. In this case, agents over-react to differences among the technologies. In contrast, when $r_f(\cdot)$ is relatively small, i.e. $f(\cdot)$ is relatively less concave, then the value of diversification to the principal is smaller so revealing information that induces more extreme beliefs is more valuable and produces an allocation closer to the first-best. Therefore, the principal may prefer to reveal information that induces more extreme beliefs.

Proposition 2. *The function $\nu(\Theta)$ has the following properties:*

- (i) *All of its global maxima are at the vertices of the $N - 1$ simplex;*
- (ii) *The center of the simplex, $\Theta = (1/N, \dots, 1/N)$ is a local maximum if and only if*

$$2 < r_f \left(\frac{\Omega}{N} \right)$$

- (iii) *If there exists a local maximum, it must be $\Theta = (1/N, \dots, 1/N)$.*

[Proposition 2](#) shows that for any concave function $f(\cdot)$, the function $\nu(\cdot)$ is either strictly convex or it has a unique local maximum at the center of the simplex. Additionally, $\nu(\cdot)$ has global maxima with value $f(\Omega)$ at the extreme points of the simplex because at those extreme points the principal and agents know with certainty which technology is the best choice. Hence, if the principal has access to a perfectly informative signal, revealing that signal is optimal.

Importantly, while [Figure 3](#) is an illustration for $N = 2$, [Proposition 2](#) shows that this figure captures all the relevant features of the problem for $N \geq 2$. In more dimensions, we will have the same two cases: either the function will be globally convex, or it will have a local maximum at the center of the simplex. Hence [Figure 3](#) is more than just an example—it qualitatively depicts *all* of the possible cases for the shape of $\nu(\cdot)$.

Whether $\nu(\cdot)$ is locally concave at the center of the simplex depends on $r_f \left(\frac{\Omega}{N} \right)$, the coefficient of relative risk aversion associated with f , evaluated at the most diversified possible equilibrium investment profile—where, in aggregate, agents invest $\frac{\Omega}{N}$ in each

technology. In particular, following [Lemma 2](#), if the relative risk aversion at the most diversified equilibrium profile is not too large, i.e., $r_f\left(\frac{\Omega}{N}\right) \leq 2$, then diversification is not an issue, reflected in $\nu(\cdot)$ being convex, and the principal always uses the most informative experiment, disclosing as much information as possible. On the contrary, if the value of relative risk aversion in the most diversified equilibrium profile is large, i.e., $r_f\left(\frac{\Omega}{N}\right) > 2$, then diversification is a problem for the principal, reflected in $\nu(\cdot)$ being locally concave at the center of the simplex. In this case, the principals' information design problem is more involved, and we analyze it next.

Characterization of the Optimal Information Design. We now characterize the optimal signal structure in our setting. Let $\Theta_0 \in \Delta^N$ be the common prior belief profile over the N technologies. Conditional on prior Θ_0 , let $\mathcal{P}_S(\Theta_0) \subseteq \Delta^N$ (simply denoted \mathcal{P}_S) be the convex set of posterior beliefs that the principal can induce that are consistent with Bayes' rule (e.g., see [Example 1](#)). From these two primitives, we now define key objects to characterize the principal's optimal information design.

Let $\Theta_C \equiv \{\Theta \in \Delta^N : \hat{\nu}(\Theta, \mathcal{P}_S) = \nu(\Theta)\}$ be the set of all posteriors where the value function, ν , agrees with its concave closure, $\hat{\nu}$, over \mathcal{P}_S . Let $\tilde{v} \equiv \sup\{\nu(\Theta) : \nu''(\Theta) \leq 0, \Theta \in \mathcal{P}_S\}$ be the largest value of the value function over the region of \mathcal{P}_S where it is concave.⁹ We can now characterize the optimal disclosure policy.

Proposition 3. *The optimal disclosure policy, s^* , with distribution G_{s^*} such that $E_{G_{s^*}}[\Theta] = \Theta_0$, is*

1. **maximally informative** if $\nu(\bar{\theta}) \geq \tilde{v}$ for all $\bar{\theta} \in \partial\mathcal{P}_S$; then s^* induces a distribution over posterior beliefs with support consisting only of points in the boundary of the feasible set of posteriors.
2. **partially informative** if $\nu(\bar{\theta}) < \tilde{v}$ for some $\bar{\theta} \in \partial\mathcal{P}_S$ and $\mathcal{P}_S \not\subseteq \Theta_C$; then s^* induces a distribution with support consisting of boundary points in $\partial\mathcal{P}_S$ or in $\mathcal{P}_S \cap \Theta_C$.
3. **uninformative** if $\mathcal{P}_S \subseteq \Theta_C$; then s^* induces the belief $\Theta = \Theta_0$.

We now consider an illustrative example with two technologies to explain the three cases in [Proposition 3](#). We take the common prior, θ_0 , and set of feasible posteriors, \mathcal{P}_S , as

⁹If this region is empty, $\tilde{v} = -\infty$.

primitives of the model and illustrate the optimal information policy in [Proposition 3](#) for different priors and feasible posteriors.¹⁰ For the sake of this example, we assume $\mathcal{P}_S = (\underline{\theta}_1, \bar{\theta}_1)$ with $0 \leq \underline{\theta}_1 \leq \bar{\theta}_1 = 1 - \underline{\theta}_1 \leq 1$, although this symmetry is not required for our results.

[Figure 4](#) shows the principal's payoff with and without information disclosure in three different scenarios when $f(x) = 1 - \exp(-\lambda x)$ with $\lambda = 8$ and $\Omega = 0.9$. In each panel, the black solid line represents the principal's payoff without disclosure, $\nu(\theta_1) \equiv \nu(\theta_1, 1 - \theta_1)$, and the red dashed line represents the concavification of the principal's payoff over the region of feasible posteriors, $\hat{\nu}(\cdot)$ over \mathcal{P}_S . Each panel highlights three features: (1) the inflection points of $\nu(\cdot)$, denoted by C and D , at $\theta_1 = 0.26$ and $\theta_1 = 0.74$, respectively; (2) the unique local maximum of $\nu(\cdot)$ inside the region of feasible posteriors, which in each case corresponds to $\theta_1 = \frac{1}{2}$; and (3) the points B and E , where the value of $\nu(\cdot)$ equals its value at the interior local maximum. The function ν is concave in the region between C and D and convex otherwise. In each case, we have that $\tilde{v} = \sup\{\nu(\theta) : \nu''(\theta) < 0, \theta \in \mathcal{P}_S\} = f\left(\frac{1}{2}\right) = 1 - \exp(-\lambda/2)$, and we have some point $\hat{\theta}_1$ such that $\hat{\theta}_1 = \sup\{\theta : \hat{\nu}(\theta, [\underline{\theta}_1, \bar{\theta}_1]) = \nu(\theta)\}$.




Figure 4: The solid black line corresponds to the principal's payoff without disclosure with $f(x) = 1 - \exp(-\lambda x)$, $\lambda = 8$, and $\Omega = 0.9$. The dashed red line corresponds to the function $\hat{\nu}(\cdot) = \text{concavification}(\nu)$ over the region of feasible posteriors \mathcal{P}_S . Each panel shows a different scenario for the set of feasible posteriors $\mathcal{P}_S = [\underline{\theta}_1, \bar{\theta}_1]$ and the prior is θ_0 .

Left Panel: $\underline{\theta}_1 = 0.02$ and $\bar{\theta}_1 = 0.98$. Prior $\theta_0 = 0.5$.

Middle Panel: $\underline{\theta}_1 = 0.1$ and $\bar{\theta}_1 = 0.9$. Prior $\theta_0 = 0.8$.

Right Panel: $\underline{\theta}_1 = 0.3$ and $\bar{\theta}_1 = 0.7$. Prior $\theta_0 = 0.5$.

¹⁰For any prior and any set of feasible posteriors, there is a signaling technology that generates the latter from the former under Bayesian updating. When we compare the optimal policy with various priors and different sets of feasible posteriors, the signaling technology itself is changing to keep the priors and posteriors Bayes-consistent (see [Figure 6](#) for an example). Our goal here is to illustrate how different priors and feasible posteriors give rise to different optimal policies.

The principal’s optimal signal structure depends on the a priori asymmetry of the technologies (i.e. the prior Θ_0), the quality or informativeness of the feasible signals the principal can use for a given prior (i.e., $\mathcal{P}_S(\Theta_0)$), and the value of diversification to the principal (i.e. \tilde{v}).

First, if the principal can generate highly informative signals, then \mathcal{P}_S includes posterior beliefs close to 0 and 1. When $\min\{\nu(\underline{\theta}_1), \nu(\bar{\theta}_1)\} > \tilde{v}$, a maximally-informative disclosure is optimal. Graphically, in [Figure 4](#) (left panel), $\underline{\theta}_1$ lies somewhere between points A and B, and $\bar{\theta}_1$ lies somewhere between points E and F, and the concave closure of $\nu(\cdot)$ is the line that connects $\nu(\underline{\theta}_1)$ and $\nu(\bar{\theta}_1)$. The optimal signal reveals $\underline{\theta}_1$ with some probability q and $\bar{\theta}_1$ with the remaining probability $1 - q$, where q is such that the expected posterior is equal to the prior, θ_0 . The principal benefits from inducing extreme posteriors. The reason is that, when the value of diversification to the principal is low (i.e. the $r_f(x)$ coefficient is low enough), revealing information to the agents increases the principal’s expected value because, in equilibrium, agents under-react to asymmetries in the technologies when beliefs are extreme. This leads to an optimal disclosure rule that mixes between the two most extreme posteriors possible within \mathcal{P}_S .

Second, suppose the available signals are not as informative and the set of feasible posteriors, \mathcal{P}_S , is narrow enough so that $\max\{\nu(\underline{\theta}_1), \nu(\bar{\theta}_1)\} < \tilde{v}$. Moreover, suppose that the technologies are ex-ante asymmetric, with technology 1 more likely to be successful ex ante, i.e., θ_0 is to the right of point D . Graphically, in [Figure 4](#) (middle panel), $\underline{\theta}_1$ lies somewhere between points B and C, and $\bar{\theta}_1$ lies somewhere between points D and E. This requires a large enough value of diversification so that \tilde{v} is large. In this case, the optimal signal is partially informative: it reveals the posterior $\hat{\theta}_1$ with some probability q and $\bar{\theta}_1$ with the remaining probability $1 - q$, where q is such that the expected posterior is equal to the prior.

Third, suppose that the principal has access to relatively uninformative signals, so she can induce posterior beliefs in \mathcal{P}_S close to 0.5 in the region where $\nu(\cdot)$ is concave. Graphically, in [Figure 4](#) (right panel), $\underline{\theta}_1$ and $\bar{\theta}_1$ lie somewhere between points C and D. In this case, $\hat{\nu}(\cdot, \mathcal{P}_S) = \nu(\cdot)$ over \mathcal{P}_S —there is no value from information disclosure, and the optimal signal is perfectly uninformative, inducing a posterior equal to the prior. The value of diversification is large enough, and the technologies are symmetric enough that revealing information would lead to more extreme posteriors, which agents would over-react to in equilibrium compared to the first-best.

2.3 Degree of Information Disclosure: Size of the Prize and Number of Competitors

We now study how the size of the prize and the number of competitors affect the optimal information disclosure policy. Both the size of the prize, V , and the number of competitors, M , impact $\nu(\cdot)$ only indirectly, through $\Omega = \left(\frac{M-1}{M}\right)V$. Thus, to study the effects of V and M , it is enough to examine the comparative statics on Ω .

Define the *degree of information disclosure* to be the type of information policy that is optimal for the principal. From [Proposition 3](#): for any given prior Θ_0 and any given set of attainable posteriors $\mathcal{P}_S(\Theta_0)$, we have *no* disclosure, *partial* disclosure, or *maximal* disclosure, depending on the parameters of the contest, including the size of the prize and the number of competing agents.

If the region of beliefs over which the principal's value function ν is concave, denoted by Θ_C , expands in terms of set inclusion, then the degree of information disclosure weakly decreases, i.e. it decreases whenever the optimal policy switches from maximal or partial disclosure towards partial or no disclosure. We characterize the effect of V and M on the degree of information disclosure by examining their effect on Θ_C .

For tractability, we focus on the case of two technologies, so $\Theta = (\theta, 1 - \theta)$. The condition for concavity of ν at Ω when $N = 2$ ([Lemma 2](#)) is equivalent to $g(\Omega, \theta) < 0$, where

$$g(\Omega, \theta) \equiv h(\Omega\theta) + h(\Omega(1 - \theta))$$

and $h(x) = xf'(x)[2 - r_f(x)]$. $g(\Omega, \theta^*) = 0$ at an inflection point, θ^* , of $\nu(\cdot)$. Thus, for each Ω , the inflection point of $\nu(\cdot)$ in $[0, 0.5]$ is the solution to the implicit equation

$$g(\Omega, \theta^*(\Omega)) = 0. \tag{7}$$

Lemma 3. *The solution to equation (7), $\theta^* \equiv \theta^*(\Omega)$, is decreasing in Ω if and only if*

$$\frac{h'(\Omega(1 - \theta^*))}{h'(\Omega\theta^*) - h'(\Omega(1 - \theta^*))} + \theta^* > 0 \tag{8}$$

A sufficient condition for inequality (8) to hold is that h is increasing and concave.

Directly from this lemma we can describe how the size of the prize and the number of

contestants affect the optimal degree of information disclosure in the contest.

Proposition 4. *The optimal degree of information disclosure weakly decreases when the number of competitors increases or when the size of the prize increases, if and only if inequality (8) holds.*

Increasing the number of competing agents or increasing the size of the prize increases the total investment on each technology. Because $f(\cdot)$ is concave, the marginal return to investment diminishes, so the value of focus becomes less important to the principal compared to the value of diversification. When inequality (8) holds, the objective becomes more concave around the center of the simplex. Hence, for any given prior Θ_0 and set of attainable posteriors \mathcal{P}_S , as the region of concavity expands, the concavification of the objective function changes according to each case of [Proposition 3](#). Therefore, the optimal information disclosure policy weakly changes from more disclosure (maximal or partial disclosure) to less disclosure (partial or no disclosure).

Example 2. *To illustrate [Proposition 4](#), we consider $f(x) = 1 - \exp^{-8x}$, for which inequality (8) holds. This means that as Ω increases the region of concavity of ν increases. [Figure 5](#) shows how the function $\nu(\cdot)$ changes for prizes of different sizes, $V = 0.6, 0.8, 1$, with $M = 5$ competitors. The function ν becomes concave around the center and the concave region expands as the prize increases: ν is convex when $V = 0.6$, but it is locally concave for $V = 0.8$, and the region of concavity expands for $V = 1$.*

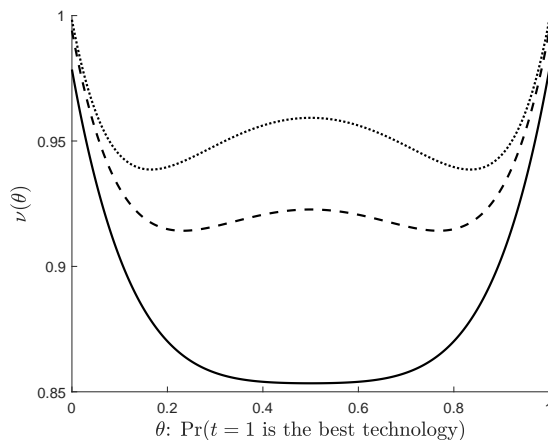


Figure 5: Function ν when $f(x) = 1 - e^{-8x}$ for different values of the prize $V = 0.6$ (solid line), $V = 0.8$ (dashed line), $V = 1$ (dotted line). The function ν is convex when $V = 0.6$; it is concave in the region $[0.29, 0.70]$ when $V = 0.8$; and it is concave in the region $[0.25, 0.74]$ when $V = 1$.

Additionally, suppose the prior is $\theta = 0.4$ and the most informative experiment is generated by the same signal specification as in [Example 1](#), with a parameter $\alpha = 0.75$. We can now characterize the optimal information disclosure policy. Note that the set of attainable posteriors is $\mathcal{P}_S = [0.18, 0.67]$.¹¹ [Figure 6](#) shows both $\nu(\cdot)$ (dashed line) and its concavification, $\hat{\nu}(\cdot)$ (solid red line). The optimal information disclosure policy depends on V : the higher the prize values, the lower the extent of information disclosure (see [Proposition 4](#)). Thus, with lower prize values (e.g. $V = 0.6$, [Figure 6](#), left panel), the principal uses maximal disclosure. At higher prize values (e.g. $V = 0.8$, [Figure 6](#), middle panel), the principal uses partial disclosure. At even higher prize values (e.g. $V = 1$, [Figure 6](#), right panel), the principal does not disclose information.




Figure 6: Function ν (dashed line) and its concavification $\hat{\nu}$ (solid line) for $f(x) = 1 - e^{-8x}$, for different prize size. The vertical line indicates the fixed prior 0.4. The set of attainable posteriors with the experiment technology in [Example 1](#) is $\mathcal{P}_S = [0.18, 0.67]$.

We can generalize [Example 2](#) and carry out the same analysis for any prior belief $\theta \in [0, 1]$. The vertical axis in [Figure 7](#) shows the set of posteriors used by the optimal information disclosure policy. In the left panel of [Figure 7](#), when $\Omega = 0.48$, the function ν is convex and, conditional on the set of feasible posteriors for each prior, the optimal policy is maximal information disclosure. In other words, for each prior, the optimal information structure induces the same posteriors induced by the most informative experiment. In the middle panel, when $\Omega = 0.64$, there is no information revelation for any prior in $[0.41, 0.59]$, partial information disclosure for priors in $[0.26, 0.41] \cup [0.59, 0.74]$ and full disclosure for priors in $[0, 0.26] \cup [0.74, 1]$. In the right panel, when $\Omega = 0.8$, there is no information revelation for any prior in $[0.35, 0.65]$, partial information disclosure for priors in $[0.21, 0.35] \cup [0.65, 0.79]$ and full disclosure

¹¹The boundary of the set of posteriors for each prior is illustrated in [Figure 2](#) (middle panel).




Figure 7: As Ω increases (from left panel to right panel), the extent of information disclosure decreases. For each prior, the solid (red) line shows the posteriors induced by the optimal information structure. In the “Full” region, there is full information disclosure, i.e., signals induce most extreme posteriors possible. In the “Partial” region, there is partial information disclosure, i.e., posteriors are not the most extreme ones. In the “None” region, there is no information disclosure, i.e., the posterior equals the prior.

for priors in $[0, 0.21] \cup [0.79, 1]$. As the number of competitors increases, the inflection point of $\nu(\cdot)$ decreases, meaning that the extent of information disclosure decreases. In other words, in a contest with more competitors, the principal may be less willing to disclose information relative to a contest with the same characteristics but with fewer competitors.

2.4 Optimal Size of the Prize and Information Disclosure

We now examine the interaction between information design and the size of the prize. In the previous section we characterized when increasing the prize reduces the extent of information disclosure. We now ask: For a given prior Θ_0 and set of posteriors $\mathcal{P}(\Theta_0)$, what is the optimal combination of prize and information disclosure? In particular, does the principal strictly benefit by disclosing information or is an optimal prize sufficient? We show that information disclosure has value for the principal even after the prize is chosen optimally, by studying the optimal size of the prize in the setting of [Example 2](#).

Without information disclosure, the principal’s expected payoff after choosing the optimal prize, for any prior Θ , is

$$\max_{V \geq 0} \nu(V, \Theta) - V, \quad (9)$$

where we now make explicit that ν depends on V and Θ . The principal’s expected payoff

from using an optimal information structure and then choosing the optimal prize is

$$\max_{V \geq 0} \hat{\nu}(V, \Theta) - V, \quad (10)$$

Since, in general, $\hat{\nu}(V, \Theta) \neq \nu(V, \Theta)$, we would not expect the maximum values in (9) and (10) to be the same. The principal gets a weakly larger payoff by optimising over the prize and the information structure. In fact, by definition,

$$\hat{\nu}(V, \Theta) = \nu(V, \Theta) + \gamma(\Theta, V),$$

where $\gamma(\cdot)$ is a positive function, since $\hat{\nu} \geq \nu$. Therefore, for any prior Θ , and any prize V , we have

$$\hat{\nu}(V, \Theta) - V \geq \nu(V, \Theta) - V,$$

so the maximum value in (10) is always larger than the maximum value in (9). Using these observations, we have the following result.

Proposition 5. *Let V_0^* be the solution to (9), i.e., the optimal prize when the principal does not disclose information. The principal strictly benefits from using information disclosure in combination with an optimal prize when $\gamma(\Theta, V_0^*) > 0$. Furthermore, when $\gamma(\Theta, V)$ is submodular in V and Θ , the optimal prize is smaller when the principal discloses information optimally than without disclosure.*

This proposition means that information disclosure and prizes can be substitutes: using a smaller prize and disclosing information can make the principal strictly better off relative to the case of no information disclosure (which requires a larger optimal prize).

As an example, we find the jointly optimal prize and information disclosure for any prior θ for $f(x) = 1 - \exp^{-8x}$ with two technologies, five competitors, and we use the experiment technology in [Example 1](#) and [Example 2](#), with $\alpha = 0.75$.

[Figure 8](#) (left panel) plots the principal's payoff as a function of the prize, V , when the prior is $\theta = 0.4$ and the set of posteriors is $\mathcal{P} = [0.182, 0.667]$. In the figure, the solid line corresponds to $\hat{\nu}(V; \theta) - V$, the principal's expected payoff when the principal uses the optimal information disclosure, whereas the dashed line corresponds to $\nu(V; \theta) - V$, the principal's expected payoff when she does not disclose information.




Figure 8: Left Panel: Functions $\hat{\nu}(\theta)$ and $\nu(\theta)$, the value for the principal of inducing posterior θ with and without information disclosure, for $f(x) = 1 - e^{-8x}$ for different values of the prize V . Right Panel: The optimal size of V as a function of the prior belief, θ , assuming the principal uses the optimal information disclosure policy.

Figure 8 (left panel) shows that information disclosure is valuable to the principal as a design tool, as it increases the principal's payoff beyond that from a contest with an optimal prize. In this example, absent any disclosure the optimal prize is $V = 0.356$ and the principal's expected payoff is 0.33. In contrast, with the optimal disclosure policy the optimal prize is $V = 0.321$, and the principal obtains an expected payoff of 0.37. This shows that the optimal information disclosure policy and the prize interact with each other: By using a combination of optimal information disclosure and prize, the principal increases her revenue by 13.4 percent while reducing the size of the prize by 9.8 percent.

Figure 8 (right panel) shows the optimal prize when we repeat the exercise of finding the optimal prize for any prior $\theta \in [0, 0.5]$. The figure shows that the optimal prize is increasing in the prior, θ_1 , for $\theta_1 \leq 0.5$. That is, as the two technologies become more symmetric in terms of their probabilities of being the best one, the principal's optimal prize increases. The intuition for this is clearer from Figure 5 and Figure 6: for priors closer to 0.5, as the prize increases, the principal's objective tends to increase more compared to how much it increases for priors further away from 0.5. Hence the optimal prize, which implicitly also depends on the optimal information policy, tends to increase when the technologies are more symmetric and the prior is closer to 0.5.

3 Conclusion

When there are different approaches to tackle a problem and agents compete in a contest to find the best solution, we ask whether it is beneficial for the contest sponsor (the principal) to disclose information regarding the different approaches. We find that it is not always beneficial to reveal that one technology is more promising than the rest when the principal cares about diversification: revealing information can induce too many agents to work on the most promising technology, which reduces diversification.

We present a tractable framework to study contests with technological uncertainty and to analyze the trade-off between information revelation and diversification. In our setting, each agent chooses one out of N available technologies to compete in the contest, and only one of these technologies is the best one *ex post*. The principal can commit to reveal to the agents the results of an experiment that signals the success of each technology. We fully characterize the optimal signal structure that maximizes the principal's expected payoff from the contest as a function of the set of all signals available to the designer. We show that the informativeness of the optimal signal structure crucially depends on three main features of the environment: (i) the value of technological diversity; (ii) the quality of the principal's information; and (iii) the extent of technological uncertainty. Each of these factors affects the principal's choice of information structure, as it affects the key trade-off between diversification and focus.

Revealing more precise information about the technologies induces more extreme posteriors, which incentivizes agents to focus on more promising technologies in equilibrium. However, the equilibrium allocation of agents' investments may over-react to such asymmetries in their beliefs regarding the different technologies compared to the principal's first-best allocation. Because the technologies are uncertain, the principal's payoff includes the option value of developing less promising technologies, so diversification is also valuable and conflicts with the incentive to focus on more promising technologies. The optimal signal structure balances these considerations and can be maximally informative, partially informative, or completely uninformative in different cases.

These results apply to contest setting where agents can pursue different approaches, such as in procurement, contests for innovation, promotions within organizations, and others. In all of these settings, the agents and the principal may be unsure about which technology, idea, or project will be most valuable or feasible *ex post*.

4 References

- Alonso, Ricardo and Odilon Câmara (2016) “Persuading voters,” *American Economic Review*, Vol. 106, pp. 3590–3605.
- Aoyagi, Masaki (2010) “Information feedback in a dynamic tournament,” *Games and Economic Behavior*, Vol. 70, pp. 242–260.
- Benkert, Jean-Michel and Igor Letina (2016) “Designing dynamic research contests,” *University of Zurich, Department of Economics, Working Paper*.
- Bhattacharya, Sudipto and Dilip Mookherjee (1986) “Portfolio choice in research and development,” *The RAND Journal of Economics*, pp. 594–605.
- Bimpikis, Kostas, Shayan Ehsani, and Mohamed Mostagir (2014) “Designing dynamic contests,” *Working paper, Stanford University*.
- Bockstedt, Jesse, Cheryl Druehl, and Anant Mishra (2016) “Heterogeneous Submission Behavior and its Implications for Success in Innovation Contests with Public Submissions,” *Production and Operations Management*.
- Boudreau, Kevin J, Nicola Lacetera, and Karim R Lakhani (2011) “Incentives and problem uncertainty in innovation contests: An empirical analysis,” *Management Science*, Vol. 57, pp. 843–863.
- Cabral, Luis (1994) “Bias in market R&D portfolios,” *International Journal of Industrial Organization*, Vol. 12, pp. 533–547.
- Cabral, Luis MB (2003) “R&D competition when firms choose variance,” *Journal of Economics & Management Strategy*, Vol. 12, pp. 139–150.
- Che, Yeon-Koo and Ian Gale (2003) “Optimal design of research contests,” *The American Economic Review*, Vol. 93, pp. 646–671.
- Clark, Derek J. and Christian Riis (1998) “Influence and the discretionary allocation of several prizes,” *European Journal of Political Economy*, Vol. 14, pp. 605–625.
- Corchón, Luis C and Marco Serena (2018) “Contest theory,” in *Handbook of Game Theory and Industrial Organization, Volume II*: Edward Elgar Publishing.

- Dasgupta, Partha and Eric Maskin (1987) “The simple economics of research portfolios,” *The Economic Journal*, Vol. 97, pp. 581–595.
- Ederer, Florian (2010) “Feedback and motivation in dynamic tournaments,” *Journal of Economics & Management Strategy*, Vol. 19, pp. 733–769.
- Ederer, Florian, Richard Holden, and Margaret Meyer (2018) “Gaming and strategic opacity in incentive provision,” *The RAND Journal of Economics*, Vol. 49, pp. 819–854.
- Fu, Qiang, Jingfeng Lu, and Jun Zhang (2016) “On Disclosure Policy in Tullock Contests with Asymmetric Entry,” *Canadian Journal of Economics*, Vol. 49.
- Fullerton, Richard L and R Preston McAfee (1999) “Auctioning entry into tournaments,” *Journal of Political Economy*, Vol. 107, pp. 573–605.
- Gross, Daniel P (2015) “Creativity Under Fire: The Effects of Competition on Creative Production,” *Available at SSRN 2520123*.
- (2017) “Performance feedback in competitive product development,” *The RAND Journal of Economics*, Vol. 48, pp. 438–466.
- Halac, Marina, Navin Kartik, and Qingmin Liu (2017) “Contests for experimentation,” *Journal of Political Economy*, Vol. 125.
- Huang, Yan, Param Vir Singh, and Kannan Srinivasan (2014) “Crowdsourcing new product ideas under consumer learning,” *Management science*, Vol. 60, pp. 2138–2159.
- Hwang, Ilwoo, Kyungmin Kim, and Raphael Boleslavsky (2019) “Competitive Advertising and Pricing,” *Working Paper*.
- Kamenica, Emir and Matthew Gentzkow (2011) “Bayesian persuasion,” *American Economic Review*, Vol. 101, pp. 2590–2615.
- Kireyev, Pavel (2016) “Markets for Ideas: Prize Structure, Entry Limits, and the Design of Ideation Contests,” *HBS, Working Paper*.
- Klein, Arnd Heinrich and Armin Schmutzler (2016) “Optimal effort incentives in dynamic tournaments,” *Games and Economic Behavior*.

- Kolotilin, Anton (2018) “Optimal information disclosure: A linear programming approach,” *Theoretical Economics*, Vol. 13, pp. 607–635.
- Konrad, Kai (2009) *Strategy and Dynamics in Contests*: Oxford University Press.
- Kovenock, Dan, Florian Morath, and Johannes Münster (2015) “Information sharing in contests,” *Journal of Economics & Management Strategy*, Vol. 24, pp. 570–596.
- Krishnan, Vish and Shantanu Bhattacharya (2002) “Technology selection and commitment in new product development: The role of uncertainty and design flexibility,” *Management Science*, Vol. 48, pp. 313–327.
- Laclau, Marie and Ludovic Renou (2016) “Public persuasion.”
- Lemus, Jorge and Guillermo Marshall (2021) “Dynamic tournament design: Evidence from prediction contests,” *Journal of Political Economy*, Vol. 129, pp. 383–420.
- Letina, Igor (2016) “The road not taken: competition and the R&D portfolio,” *The RAND Journal of Economics*, Vol. 47, pp. 433–460.
- Letina, Igor and Armin Schmutzler (2019) “Inducing variety: A theory of innovation contests,” *International economic review*, Vol. 60, pp. 1757–1780.
- Mathevet, Laurent, Jacopo Perego, and Ina Taneva (2017) “On information design in games,” Technical report, mimeo.
- Mihm, Jürgen and Jochen Schlapp (2018) “Sourcing innovation: On feedback in contests,” *Management Science*, Vol. 65, pp. 559–576.
- Moldovanu, Benny and Aner Sela (2001) “The optimal allocation of prizes in contests,” *American Economic Review*, pp. 542–558.
- Pérez-Castrillo, J David and Thierry Verdier (1992) “A general analysis of rent-seeking games,” *Public choice*, Vol. 73, pp. 335–350.
- Serena, Marco (2021) “Harnessing beliefs to optimally disclose contestants’ types,” *Economic Theory*, pp. 1–30.
- Taylor, Curtis R (1995) “Digging for golden carrots: an analysis of research tournaments,” *The American Economic Review*, pp. 872–890.

Terwiesch, Christian and Yi Xu (2008) “Innovation contests, open innovation, and multiagent problem solving,” *Management science*, Vol. 54, pp. 1529–1543.

Toh, Puay Khoon and Taekyu Kim (2013) “Why put all your eggs in one basket? A competition-based view of how technological uncertainty affects a firm’s technological specialization,” *Organization Science*, Vol. 24, pp. 1214–1236.

Xin, Feng and Jingfeng Lu (2016) “The Optimal Disclosure Policy in Contests with Stochastic Entry: A Bayesian Persuasion Perspective,” *Economics Letters*, Vol. 147.

Zhang, Jun and Junjie Zhou (2016) “Information disclosure in contests: A Bayesian persuasion approach,” *The Economic Journal*, Vol. 126, pp. 2197–2217.

A Appendix: Proofs

Proof of Proposition 1

Proof. These results are standard in the literature (see, e.g., [Corchón and Serena, 2018](#)). Agent i solves:

$$\max_{x_{i,t} \geq 0} \sum_{t=1}^N \left[V \theta_t \frac{x_{i,t}}{\sum_{j=1}^M x_{j,t}} - x_{i,t} \right].$$

This problem is separable in t , and so the first-order optimization conditions yield the following equilibrium level of investment on technology t :

$$x_{i,t}^* = \theta_t \left(\frac{M-1}{M^2} \right) V. \tag{11}$$

Thus, the aggregate equilibrium investment on technology t is

$$x_t^* \equiv \sum_{i=1}^M x_{i,t}^* = \theta_t \Omega, \tag{12}$$

where $\Omega = \left(\frac{M-1}{M} \right) V$.

If the principal could control the agent's investments, the principal's would choose the aggregate investment on technology t according to:

$$\max_{y_t} \sum_{t=1}^N \theta_t f(y_t) - y_t$$

which implies that, at an interior solution, $\theta_t f'(y_t^*) = 1$ or

$$y_t^* = h(\theta_t) \equiv (f')^{-1} \left(\frac{1}{\theta_t} \right). \quad (13)$$

Note that y_t^* increases with θ_t because $h(\cdot)$ is increasing (by concavity of $f(\cdot)$).

For technology t both the equilibrium investment, $x_t^*(\theta_t)$, and the first-best, $y_t^*(\theta_t)$, only depend on the belief that technology t is the best one, rather than on the entire profile of beliefs. The aggregate investment for different technologies, however, are not arbitrary because they relate to each other through the condition $\sum_{t=1}^N \theta_t = 1$.

The aggregate equilibrium investment across all technologies is

$$\sum_{t=1}^N x_t^*(\theta_t) = \Omega,$$

which is *independent* of the agents' beliefs and only depends on the size of the prize and the number of agents. In contrast, the aggregate first-best investment is

$$\sum_{t=1}^N y_t^*(\theta_t) = \sum_{t=1}^N h(\theta_t),$$

which depends on the profile of beliefs over all the technologies. The equilibrium investment on technology t depends only on θ_t and Ω , while the first-best level depends on the *shape* of the function f and on the belief θ_t . \square

Proof of Lemma 1

Proof. Consider any two posteriors $\tilde{\theta}, \tilde{\theta}' \in \mathcal{P}_S$ induced by some messages m' and m'' , from (possibly different) signal structures s' and s'' , respectively. For any $\alpha \in (0, 1)$, the posterior $\alpha\tilde{\theta} + (1-\alpha)\tilde{\theta}'$ can be induced with a signal structure s^* that is an appropriately

chosen mixture of s' and s'' . First, define a new message m^* , and replace m' and m'' with m^* in s' and s'' , respectively. Thus, if the receiver knows whether a message is generated by s' or s'' , their posterior conditional on m^* would be $\tilde{\theta}$ or $\tilde{\theta}'$, respectively. Now define the mixture s^* so that with probability $\Pr(s') = \frac{\alpha \cdot \Pr(m^*)}{\Pr(m^*|s')}$ a message is generated according to s' , and with corresponding probability $\Pr(s'') = \frac{(1-\alpha) \cdot \Pr(m^*)}{\Pr(m^*|s'')}$ a message is generated according to s'' .

The mixture probabilities $\Pr(s')$ and $\Pr(s'')$ are chosen such that conditional on observing the message m^* , the receiver believes that with probability α this message was generated by s' , and with probability $1 - \alpha$ it was generated by s'' . The receiver's posterior is $\alpha\tilde{\theta} + (1 - \alpha)\tilde{\theta}'$, so the set \mathcal{P}_S is convex. \square

Proof of Lemma 2

Proof. If $\Theta = (\theta_1, \dots, \theta_N) \in \Delta^{N-1}$ we can write $\theta_N = 1 - \sum_{t=1}^N \theta_t$. Then, restricted to the $(N - 1)$ -simplex, ν is a function of $N - 1$ variables. To show that this function is concave, let $H(\Theta)$ be the $(N - 1) \times (N - 1)$ matrix of second derivatives of $\nu(\Theta)$. Given the separability of the function $\nu(\cdot)$ the matrix $H(\Theta)$ is diagonal. The condition in the lemma corresponds to imposing that the element in row t is negative, which is a necessary and sufficient condition for the Hessian to be negative definite. \square

Proof of Proposition 2

Proof. For part (i), note that at each vertex $\theta_i = 1$ for some $i \in N$ and $\theta_j = 0 \forall j \neq i$. Hence the values at the vertices are $\nu(0, \dots, 0, 1, 0, \dots, 0) = 1 \cdot f(\Omega)$. The only points that can obtain the global maximum of $f(\Omega)$ are the vertices of the simplex. To see this, consider any point $\tilde{\theta} = (\theta'_1, \dots, \theta'_N)$ such that $\nu(\tilde{\theta}) = \sum_j \theta'_j f(\theta'_j \Omega) = 1$. Then $f(\theta'_j \Omega) \geq f(\Omega)$ for some $j \in N$. Since f is increasing, this requires $\theta'_j = 1$, hence the point $\tilde{\theta}$ is a vertex.

For part (ii), note

$$\nu(\Theta) \equiv \sum_t \theta_t f(\theta_t \Omega).$$

We can replace $\theta_N = 1 - \sum_{i=1}^{N-1} \theta_i$. Then we have

$$\nu(\Theta | \Theta \in \Delta N) \equiv \theta_N f(\theta_N \Omega) + \sum_{t=1}^{N-1} \theta_t f(\theta_t \Omega).$$

Taking the derivative w.r.t θ_t , for $t = 1, \dots, N - 1$ we get

$$\frac{\partial \nu}{\partial \theta_t} = -f(\theta_N \Omega) - \theta_N \Omega f'(\theta_N \Omega) + f(\theta_t \Omega) + \theta_t \Omega f'(\theta_t \Omega).$$

For an interior maximum, we need $\nabla \nu = 0$, so we require

$$f(\theta_N \Omega) + \theta_N \Omega f'(\theta_N \Omega) = f(\theta_t \Omega) + \theta_t \Omega f'(\theta_t \Omega), \quad \text{for all } t = 1, \dots, N - 1.$$

Clearly $\theta_t = \frac{1}{N}$ is a solution to this equation.

Furthermore, the second order condition is negative if and only if $2 < r_f\left(\frac{\Omega}{N}\right)$, which yields part (ii).

For part (iii), suppose for the sake of a contradiction that there exists a local maximum $\Theta = (\theta_1, \dots, \theta_N) \neq (1/N, \dots, 1/N)$. Then there exist some indices i, j such that $\theta_i < \theta_j$. W.l.o.g. we can relabel $\theta_t = \theta_i$ and $\theta_N = \theta_j$. From the FOC, and using that f is increasing, we get:

$$0 < f(\theta_N \Omega) - f(\theta_t \Omega) = -\theta_N \Omega f'(\theta_N \Omega) + \theta_t \Omega f'(\theta_t \Omega).$$

Hence $\theta_t \Omega f'(\theta_t \Omega) - \theta_N \Omega f'(\theta_N \Omega) > 0$. Plugging this in the second order condition we get:

$$\Omega \{2[\theta_t f'(\theta_t \Omega) - \theta_N f'(\theta_N \Omega)] - \theta_t f'(\theta_t \Omega) r_f(\theta_t \Omega) - \theta_N f'(\theta_N \Omega) r_f(\theta_N \Omega)\} > 0,$$

which is a contradiction. So any belief Θ that satisfies the FOC but where $\theta_t < \theta_N$ cannot be a maximum, which yields part (iii). \square

Proof of Proposition 3

Proof. The proof follows from the convexity of \mathcal{P}_S in Lemma 1, the characterization of the objective function in Proposition 2, and the standard concavification argument.

1. A maximally informative signal obtains when the set of posteriors is rich enough, so the global maxima of $\nu(\Theta)$ over \mathcal{P}_S is below the concave closure of $\nu(\Theta)$ over the region \mathcal{P}_S , because $\partial\mathcal{P}_S$ includes points close to the vertices of the simplex, so the concave closure of $\nu(\Theta)$ over \mathcal{P}_S corresponds to the plane that connects the boundary of \mathcal{P}_S .

Then the optimal signal s^* only induces posteriors in $\partial\mathcal{P}_S$, so for any arbitrary prior $p \in \mathcal{P}_S$, Bayesian consistency of the posteriors determines the distribution over posteriors on $\partial\mathcal{P}_S$.

2. A partially informative signal obtains when the set of posteriors is limited, so the concave closure of $\nu(\Theta)$ over \mathcal{P}_S coincides with $\nu(\Theta)$ for some values (near the center of the simplex Δ^N).

3. An uninformative signal obtains when the set of posteriors is concentrated towards the center of the simplex Δ^N , where ν is concave, so the concavification of $\nu(\cdot)$ over \mathcal{P}_S and $\nu(\cdot)$ itself coincide. \square

Proof of Lemma 3

Proof. First, recall that $\nu(\cdot)$ satisfy $\nu(\theta) = \nu(1 - \theta)$, it is convex at $\theta = 0$, and $\theta = 0.5$ is the unique local maximum. Therefore, there exists a unique $\tilde{\theta} \in [0, 0.5]$ such that $\nu(\cdot)$ is convex in $[0, \tilde{\theta}]$ and concave in $[\tilde{\theta}, 0.5]$. For any Ω , the inflection point as an implicit function of Ω defined by

$$g(\Omega|\tilde{\theta}(\Omega)) = 0.$$

Taking derivative of this expression with respect to Ω we get:

$$\frac{\partial g(\Omega|\tilde{\theta}(\Omega))}{\partial \Omega} + \frac{\partial g(\Omega|\tilde{\theta}(\Omega))}{\tilde{\theta}} \frac{d\tilde{\theta}(\Omega)}{d\Omega} = 0.$$

Thus,

$$\frac{d\tilde{\theta}(\Omega)}{d\Omega} = \frac{-\frac{\partial g(\Omega|\tilde{\theta}(\Omega))}{\partial \Omega}}{\frac{\partial g(\Omega|\tilde{\theta}(\Omega))}{\tilde{\theta}}}.$$

We can write $g = h(\Omega\theta) + h(\Omega(1 - \theta))$. Then, the derivative of g with respect to Ω is $\theta h'(\Omega\theta) + (1 - \theta)h'(\Omega(1 - \theta))$ and the derivative with respect to θ is $\Omega[h'(\Omega\theta) - h'(\Omega(1 - \theta))]$.

$\theta))$]. Therefore,

$$\frac{d\tilde{\theta}(\Omega)}{d\Omega} = \frac{-[\theta h'(\Omega\theta) + (1-\theta)h'(\Omega(1-\theta))]}{\Omega[h'(\Omega\theta) - h'(\Omega(1-\theta))]} = \frac{-h'(\Omega(1-\theta))}{\Omega[h'(\Omega\theta) - h'(\Omega(1-\theta))]} - \frac{\theta}{\Omega}.$$

Hence, the derivative is, evaluating at the inflection point we have, negative if

$$\frac{h'(\Omega(1-\theta))}{h'(\Omega\theta) - h'(\Omega(1-\theta))} + \theta > 0$$

which is the condition in the lemma. If h is increasing and concave, this condition clearly holds. \square

Proof of Proposition 4

The optimal degree of information disclosure is characterized by the inflection point θ^* . The further this point is from 0.5, the less the degree of information disclosure. The condition in Lemma 3 implies that the inflection point moves further from 0.5. Since Ω is increasing both in the size of the prize and the number of competitors, the result obtains.

B Microfoundations for Tullock and $f(\cdot)$

In this section, we provide two possible micro-foundations that deliver particular cases of the two primitives of our model: (1) agents compete in a Tullock contest on each technology; and (2) the principal's payoff is given by a strictly increasing and concave function, $f(x_t)$.

B.1 Achieving maximum accuracy

Agents compete in a contest to achieve the highest accuracy, such as in a prediction contest where the goal is to minimise prediction errors. Their investments translate into a *score*, and we assume that the player with the lowest score wins. In a prediction contest, for example, the lowest score reflects the most accurate prediction, which has

the lowest mean-squared error.

Let us consider one of the technologies separately. Suppose that agent i obtains score $s_i \sim \exp(-x_i)$ when investing x_i . That is,

$$\Pr(s_i \leq s) = 1 - \exp(-x_i s).$$

Agent i obtains the most accurate predictions when $s_i < \min_{j \neq i} s_j$. Note that $\min_{j \neq i} s_j \sim \exp(-\sum_{j \neq i} x_j)$. Therefore,

$$\begin{aligned} P(s_i < \min_{j \neq i} s_j) &= \int_0^\infty P(s_i < \min_{j \neq i} s_j | s_i = s) P(s_i = s) ds \\ &= \int_0^\infty \exp\left(-s \cdot \sum_{j \neq i} x_j\right) x_i \exp(-x_i s) ds \\ &= \frac{x_i}{\sum_{k=1}^M x_k} \end{aligned}$$

Therefore, agent i 's probability of winning is given by the ratio of his investment over the total investment on a particular technology.

Next, suppose that the principal's payoff from obtaining a minimum score of s for the best technology is $B(s)$. Assume that $B(s) = 0$ for $s > \bar{s}$, for some threshold \bar{s} . In other words, a principal that seeks to achieve accuracy gains nothing when the most accurate prediction is, in fact, sufficiently inaccurate. Let $X \equiv \sum_{k=1}^M x_k$ be the agents' aggregate investment on the best technology. Then the principal's expected payoff is

$$f(X) = \int_0^\infty B(s) X \exp(-X \cdot s) ds,$$

which is a function of the agents' total effort. This function is strictly increasing and concave as long as $sX < 1$. Note that the agents' equilibrium aggregate investment is strictly lower than the prize V , so $sX < 1$ for any $X \in [0, V]$ as long as $\bar{s}V \leq 1$.

Hence this simple foundation induces the reduced-form representation in our main model, where the contestants win according to the Tullock contest probabilities, while the principal's payoff is strictly increasing and concave over the relevant range of the aggregate investment on each technology.

B.2 Noisy Investments and Exponential Cost

Suppose that an agent’s actual investment on a given technology depends on a level of investment that they choose, y_i , plus a random shock, ε_i , which the agent cannot control. A negative shock may represent that the agent cannot commit the planned investment, and must invest less. A positive shock may represent a “donation” made by an external party, which adds to the agent’s investment. Thus, the agent’s actual investment in a technology is

$$s_i = y_i + \varepsilon_i.$$

The principal awards the prize to the agent with the largest actual investment. Assuming that ε_k is i.i.d. from a type-I extreme value distribution, the probability that agent i wins the contest is

$$\Pr(s_i > \max_{j \neq i} s_j) = \frac{\exp(y_i)}{\sum_{k=1}^M \exp(y_k)}.$$

Next, suppose that the principal’s payoff is equal to the highest actual investment by some agent on the best technology. We have that

$$E[\max_{j=1, \dots, M} s_j] = \ln \left(\sum_{j=1}^M \exp(y_j) \right)$$

Furthermore, assuming that the investment cost is exponential, and making the change of variables $x_i = \exp(y_i)$, we obtain both the Tullock contest winning probabilities and the appropriate principal’s payoff, $f(x) = \ln(x)$, which is strictly increasing and concave.

Hence this foundation also induces the reduced-form representation in our main model, in this case using noisy investments and an exponential cost function.

C Technology-specific Prizes

In this section we study the possibility of choosing technology-specific prizes, $\{V_t\}_{t \in N}$, which naturally changes the allocation of investment across technologies. There are two important differences with our baseline analysis. First, the parameter Ω in the baseline model is now technology-specific, $\Omega_t = \left(\frac{M-1}{M}\right) V_t$. This means that the equilibrium

investment towards technology t is $x_t^* = \theta_t \Omega_t$. Second, the prize is awarded only to the ex-post best technology, which means that V_t is paid with probability θ_t .

Suppose the technology-specific prizes are announced before the realization of the principal's experiment. That is, the prizes are chosen ex-ante and are independent of the posterior beliefs. Then, the principal's payoff is

$$\nu(\Theta; V_1, \dots, V_N) \equiv \sum_{t=1}^N \theta_t [f(\theta_t \Omega_t) - V_t]$$

Figure 9 shows this function for $f(x) = \sqrt{x}$ and for $f(x) = 1 - \exp^{-8x}$. Note that the asymmetry of technology-specific prizes invalidates some properties of ν used in the main analysis. For instance, the global maximum is not attained at all the extrema of the simplex: it is attained at the technology t that maximizes $f(\Omega_t) - V_t$. We also lose symmetry with respect to the center of the simplex.




Figure 9: Left Panel: $f(x) = \sqrt{x}$, technology-specific prizes $V_1 = 0.5$ and $V_2 = 0.3$. Right Panel: $f(x) = 1 - \exp^{-8x}$, technology-specific prizes $V_1 = 0.8$ and $V_2 = 0.85$.

Our analysis readily extends to the case with technology-specific prizes: we can compute the function ν and find its concavification, to study whether or not there are gains from information disclosure.

Lemma 4. $\nu(\Theta)$ is concave at $\Theta = (\theta_1, \dots, \theta_N) \in \Delta^N$ if and only if

$$\Omega_t \theta_t f'(\theta_t \Omega_t) [2 - r_f(\theta_t \Omega_t)] + \Omega_N \theta_N f'(\theta_N \Omega_N) [2 - r_f(\theta_N \Omega_N)] < 0, \text{ for all } t = 1, \dots, N - 1. \quad (14)$$

where $r_f(x) \equiv -\frac{x f''(x)}{f'(x)}$ is the relative risk aversion coefficient associated with $f(\cdot)$.

These result is analogous to Lemma 2, except that different prizes in combination with beliefs determine the concavity of $\nu(\cdot)$. However the asymmetry introduced by

technology-specific prizes makes the taxonomy of information policies more cumbersome to describe. Figure 10 shows three cases in which the principal benefits from information disclosure when there are technology-specific prizes.




Figure 10: Value of information disclosure for three particular combinations of technology-specific prizes for $f(x) = 1 - \exp^{-8x}$ when the prior is $\theta = 0.4$ and the most informative experiment is generated by $\alpha = 0.75$ in Example 1. Left Panel: $V_1 = 0.6$ and $V_2 = 0.65$. Middle Panel: $V_1 = 0.7$ and $V_2 = 0.75$. Right Panel: $V_1 = 0.8$ and $V_2 = 0.85$.

Next we ask: what is the optimal combination of prizes and information disclosure?




Figure 11: The value of information disclosure for different combinations of technology-specific prizes for $f(x) = 1 - \exp^{-8x}$, when the prior is $\theta = 0.4$ and the most informative experiment is generated by $\alpha = 0.75$ in Example 1.

Figure 11 shows the principal’s payoff for different combinations of technology-specific prizes (V_1, V_2) , when $f(x) = 1 - \exp^{-8x}$, when the prior is $\theta = 0.4$ and the most informative experiment is generated by $\alpha = 0.75$ in Example 1. In this case, the

optimal prize using the optimal information disclosure is $V_1 = V_2 = 0.32$, whereas the optimal prize without information disclosure is $V_1 = 0.35 < V_2 = 0.37$.