Diversification and Information in Contests*

Jorge Lemus[†] and Emil Temnyalov[‡]

April 21, 2021

Abstract

We study contests with technological uncertainty—where agents can develop different technologies to compete, but it is uncertain which technology is the best. An experiment that reveals information about the merit of each technology can focus efforts onto better technologies, but it can also reduce technological diversification. We characterize the optimal experiment, which depends on the value of diversification, the informativeness of available experiments, and the asymmetry of technologies. We also study the interaction between prize and information design: the extent of information disclosure decreases with the size of the prize and the number of competitors.

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), 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 to compete in the contest, and it is not clear which of these will ultimately be successful. In prediction competitions hosted either on digital platforms (e.g. Kaggle or Innocentive) or by a firm (e.g. Netflix), there are multiple algorithms that agents can use. In agricultural yield contests, such as those sponsored by the National Corn Growers Association, farmers compete to produce the highest yield and can use different methods whose performance is subject to uncertainty (e.g. the yield may be affected by the weather or the soil). 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 preferences of the procurer over different possible attributes or designs that could be developed. Similarly, technological uncertainty is present in contest-like settings within organizations: competing workers face uncertainty regarding the impact of different tasks on their promotion chances.

Both uncertainty and competitive pressures affect the agents' decisions of how much effort to invest into each technology. Uncertainty means that some agents will inevitably work on technologies that will turn out to be unsuccessful. For this reason, incentivizing agents to develop multiple approaches increases the principal's chances to ultimately implement a successful technology. However, diversification across technologies also reduces the total amount of effort devoted to developing the best technology, which reduces the principal's payoff from implementing the best technology. In such a setting, the principal may be able to reveal some information regarding the different technologies to affect the agents' beliefs about the prospects of the different alternatives, which in turn changes the equilibrium allocation of effort across technologies in the contest.

Public information revelation is a feature of many contest settings. In prediction contests, the contest sponsor (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, and hence the contestants can use these data to choose which approach to develop. Furthermore, the principal wants to avoid "over-fitting," i.e., algorithms that perform well on one dataset but not in general, and seeks to incentivize the contestants to develop algorithms that are valuable beyond the existing data. In yield contests, the principal can reveal the performance of agricultural approaches used in the past, although this is not a guarantee that these approaches will work well in the future because, for example, weather conditions could be different. In procurement contests, participants submit designs and prototypes that can be evaluated by the procurer at some preliminary stage. Revealing these evaluations will affect the contestants' beliefs about the likelihood of success of each design, which in turn will affect the contestants' choices of what to develop subsequently.

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

The principal's optimal policy balances the trade-off between providing the agents with better information about the technologies and exploiting the option value of diversification. In our model agents compete in a contest and choose how much effort to exert in developing each of multiple possible technologies.[1] Only one of these technologies works—this is the "best" technology—and it is the method or technology that will be implemented ex post by the principal. Neither the agents nor the principal are informed ex ante about which technology is the best one. The principal designs a public experiment and commits to its result in order to reveal information to the agents regarding the value or likelihood-of-success of each technology. The information revealed by the experiment is useful to assess the merits of each technology, but it may induce an equilibrium where too many agents choose the more promising technology and there is too little technological diversification from the perspective of the principal.

We characterize the optimal information structure that maximizes the principal's expected payoff from the contest. The main trade-off 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 find that the optimal policy can be either maximally informative, partially informative, or completely uninformative, depending on the features

[1] These 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.

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, developing multiple different approaches has an option value: 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 with very high probability which technology is the right one, then she is more likely to want to reveal such information. In practice, however, the principal may not have access to very informative signals, in which case she may prefer to reveal less information.

Third, the extent of technological uncertainty reflects how similar or 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 likely it is that the principal will choose to not reveal 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 in the contest. We find that, in contests with more competitors or larger prizes, the optimal design features less information disclosure. This is because when the prize is larger or there are more agents competing, the value of diversification across technologies tends to be larger for the principal. Hence the value of disclosing information and focusing efforts towards one technology or another is generally lower than the value of diversification, as there is more effort invested in each technology.

We then explore the interaction between the amount of information disclosure and the optimal size of the prize. We find that information design generally adds value above and beyond what can be achieved by just choosing the prize optimally and not disclosing information. This is because the size of the prize affects the levels of effort invested in the contest, whereas information disclosure also affects the direction of effort—the relative allocation of effort across different technologies. Hence the latter can yield effort allocations that the former cannot necessarily achieve. We find that, when optimizing over both information disclosure and prizes, the optimal solution entails a smaller prize and more information disclosure. In other words, 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.

Our results apply to a number of problems in innovation, procurement, and organizations, where technological uncertainty is important. The model sheds light on the information design aspect of online platforms like Kaggle, DrivenData, and crowdAI, among others, which offer firms the possibility to sponsor contests to outsource their data science needs. In these contests, when contestants choose among multiple technologies to develop their algorithms (e.g. machine learning, regression methods, or other prediction algorithms), they do not know which method will turn out to be the best. However, contestants can infer the performance of different alternatives by testing them on a dataset provided by the principal, who chooses how much data to disclose.[2] 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.[3] Thus an important contest design question is how much data to reveal in the contest. Revealing more allows contestants to choose a method that performs well in the current data, but it may also induce too little diversification, i.e. too little effort across methods that do not work well but could turn out to be the best method ex post.

Related Literature. Our paper contributes to the recent literature on diversification in contests. Letina and Schmutzler (2017) characterize the optimal prize structure when the designer wants to induce a variety of approaches. Our paper offers a new model of variety, and our results are complementary to theirs in that we focus on the problem of information design, which can also be used to induce variety, rather than on the optimal prize structure. Terwiesch and Xu (2008) incorporate diversity into the preferences of the contest designer and show that more participants may be preferred.

[2]In these contests, the contest designer partitions the available data into a "test dataset" and an "evaluation dataset." The test dataset allows participants to learn how their algorithms perform, but the final prize allocation is based on the performance of the algorithm over the evaluation dataset.

[3]In addition to whichever algorithm performs best given their *current data*, the best algorithms will be determined once Netflix has more data on users' preferences.

Boudreau et al. (2011) empirically test the effect of the number of participants on diversity. 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. They find that a firm's technology becomes more specialized under greater uncertainty. Related to this, Krishnan and Bhattacharya (2002) study how a firm should design a product when there are several uncertain alternatives for the product's underlying technology.

Because we focus on information design, our paper also contributes to the literature on disclosure and feedback in contests. The existing literature focuses on information regarding the agents' characteristics or actions rather than information about the underlying technologies, which we focus on in our paper. For instance, Aoyagi (2010) studies a dynamic tournament and compares the effort provision by agents under full disclosure of their performance (i.e., players observe their relative position) and no information disclosure. Ederer (2010) adds private information to this setting, and Klein and Schmutzler (2016) add other decisions regarding the allocation of prizes and alternative performance evaluations. Fu et al. (2016) and Xin and Lu (2016) study optimal information disclosure regarding agents' entry decisions in contests. Zhang and Zhou (2016) study information disclosure regarding one player's effort costs, whereas Mihm and Schlapp (2018) study the optimal information disclosure to maximize the provision of effort when players are uncertain about the principal's preferences. Kovenock et al. (2015) study the effect of players sharing information throughout the contest. Feedback in dynamic contests has been recently studied by Bimpikis et al. (2014) and Benkert and Letina (2016). Recent empirical work assessing the effect of performance feedback on competition outcomes includes Gross (2015, 2017), Huang et al. (2014), Kireyev (2016), Bockstedt et al. (2016), and Lemus and Marshall (2017).

This 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 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 number of prizes and the number of participants.

2 Model

A principal organizes a contest to solve a problem. The contest awards a prize, V, to the best solution to the problem. There are M agents, indexed by $i \in \{1, ..., M\}$, who compete in the contest. There are N different methods or technological approaches (henceforth, "technologies"), indexed by $t \in \{1, ..., N\}$, that the agents can use to tackle the problem that needs to be solved in the contest.

It is common knowledge that only one of these technologies is appropriate to solve the problem, 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, $\Theta = (\theta_1, ..., \theta_N) \in \Delta^N$, where

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

and θ_t is the common prior belief that technology t is the best one, with $t = 1, ..., N$.

Agents choose how much effort to invest into each technology. We model competition within each technology as a standard Tullock contest (see, e.g., Pérez-Castrillo and Verdier, 1992). Conditional on technology $t \in \{1, ..., N\}$ being the best one ex post, the prize is awarded among all agents who exerted effort in that technology. Agent i wins the contest for technology 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 effort in technology t and $x_{-i,t}$ is the profile of efforts that agent i's rivals exert on technology t.[4]

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

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

where $f(\cdot)$ is a strictly increasing and concave production function and $\mathbb{1}(\cdot)$ is the indicator function. Note that the principal derives no benefit from the effort exerted on technologies other than the best one. More effort on the best technology increases the principal's payoff, but there are decreasing marginal returns. The function $f(\cdot)$ can also be interpreted as a reduced-form representation of a model where agents' efforts lead them to draw a stochastic number of i.i.d. draws out of some output distribution—hence investing more effort implies that the first order statistic increases in expectation as a function of effort, with decreasing marginal returns.

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

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

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

Modeling Assumptions. In our model, the principal is risk neutral over output

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

profiles, $(f(x_1), ..., f(x_N))$. However, when we consider expected payoffs over the profile of total efforts on each technology, $(x_1, ..., x_N)$, the principal's preferences over such profiles are risk averse due to the concavity of $f(\cdot)$. The latter interpretation—with risk aversion over effort profiles—will turn out to be insightful for our results, 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 the principal's payoffs as defined over effort profiles rather than over the outputs that they correspond to.

In our model 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)). Later in the analysis, we will also study the optimal size of the prize in conjunction with optimal information disclosure. 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 and show that there are still gains from information design. For the sake of tractability, our baseline model uses a prize that is technology-neutral.

Finally, 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.

2.1 Preliminary Analysis

Consider agent i's problem of choosing how much effort to exert on technology t. Given the common belief $\Theta = (\theta_1, ..., \theta_N)$ and the aggregate effort that agent i's rivals exert on technology t, agent i's optimal effort solves the following problem:

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

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 the first-best allocation of effort—if the principal could control the agent's efforts—the amount of effort on technology t is found by solving:

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

Proposition 1. *We have:*

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

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

2. *The aggregate equilibrium effort on technology t is*

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

3. *The first-best allocation of effort on technology t is*

$$y_t^* = (f')^{-1} \left(\frac{1}{\theta_t} \right) \cdot 1(\theta_t f'(0) > 1).$$

From Proposition 1, the principal's expected payoff from holding a contest in which agents compete under belief $(\theta_1, ..., \theta_N)$ is

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

Note that this function is non-linear in θ.[5]

[5] When posterior beliefs are distributed according to some distribution G, this non-linearity implies

The agents' and the principal's incentives are generally misaligned. Therefore, the equilibrium and first-best allocations differ in both the level of effort and the relative allocation of effort across technologies. The proportion of total effort allocated towards technology t in equilibrium is

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

and it is generically different than the proportion of effort allocated towards technology t in the first-best, 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 illustrates the proportion of total effort allocated towards technology 1 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 effort 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 effort allocated to technology 1 in equilibrium (dashed line) is *larger* relative to the first best (solid line). Therefore, in equilibrium the proportion of effort 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 effort allocated in equilibrium towards technology 1 is larger (respectively, lower) relative to the first best. Hence the equilibrium proportion of effort 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.

---

that we cannot write the principal's expected payoff as a function of posterior means (or more generally as a function of one of G's moments). This is in contrast to much of the Bayesian persuasion literature, which has studied special settings where the principal's payoff ultimately only depends on expected posteriors—see e.g. Kolotilin (2018); Hwang et al. (2019).

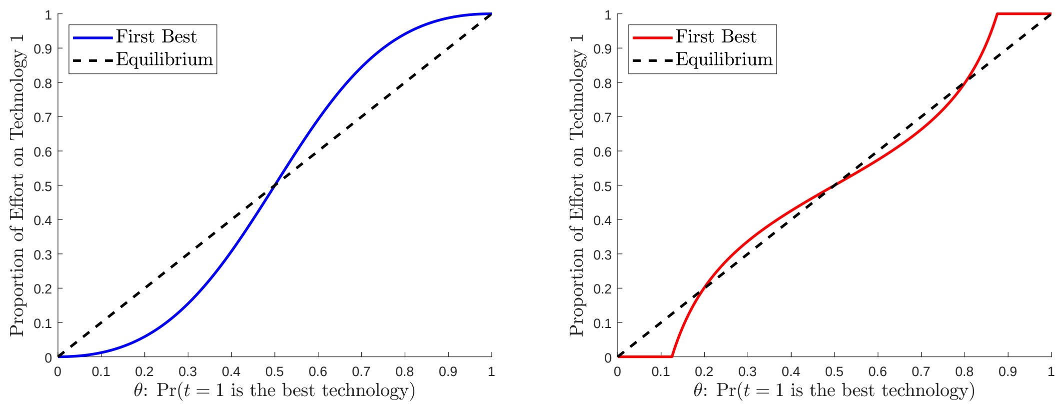

Figure 1: Comparison of the proportion of effort allocated towards technology 1 as a function of the belief θ. In the left panel, the principal's preference is $f(x) = \sqrt{x}$, whereas in the right panel it is $f(x) = 1 - \exp(-8x)$.

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 creates useful information: it can improve the allocation of efforts 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 more misallocation of effort 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 may be able design an experiment that reveals information to the agents regarding the feasibility or likelihood of success of the different technologies. Importantly, it may be *impossible* for the principal to design perfectly revealing experiments: if it was possible, the principal would always choose a perfectly revealing experiment. One important feature of our analysis is that, even after the result of an experiment is publicly revealed, there may be some 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.[6] In that sense, the test dataset is a public signal that agents can 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 contest designer use information disclosure to improve the equilibrium allocation of agents across technologies? Should the principal release information? 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?

We assume that the designer can commit to reveal the result of a public experiment to the agents, informing them about the feasibility or value of the different technologies. We focus on public information disclosure rather than private disclosure because, in practice this is the case in the applications we discuss in the introduction: in prediction contests, for example, all contestants have access to the same test data, which means that they can, in principle, generate the same beliefs regarding the performance of different algorithms. After observing the information revealed in this experiment, each agent chooses how much effort to invest in the technologies. Once each agent has made their choice, the contest resolves all remaining uncertainty and prizes are allocated.

Information structures. The principal designs an experiment that reveals public information about the technologies to all agents. An experiment is a signal structure $s = (\mathcal{M}, \widetilde{G}(\cdot|t))$, where $\mathcal{M}$ is a set of messages and $\widetilde{G}(m|t)$ 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, we do not assume that the principal has access to *every* signal structure. This is motivated by the observation that, in many practical applications, a perfectly informative signal is impossible to generate: in many applications it is unfeasible to design an experiment that eliminates all the uncertainty. Additionally, the problem is trivial if a perfectly informative signal exists: the principal reveals to the agents which technology is the best one, and all agents work on this technology. Instead, we solve for the optimal information disclosure policy for *any arbitrary* set of available signals. The set S allows for discrete or continuous posteriors, and it allows for both partitional or noisy signal structures.

[6] In 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.

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 signal as $\mathcal{P}_s$, with generic elements $\Theta \in \mathcal{P}_s$. Let $\mathcal{P}_S \equiv \cup_{s \in S} \mathcal{P}_s$ denote the set of all posterior beliefs that can be induced by some signal that is available to the principal.

Lemma 1. *For any set of signal structures S, the set of posteriors that can be induced by the principal, $\mathcal{P}_S = \cup_{s \in S} \mathcal{P}_s$, 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''. 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:*

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

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 illustrates the range of feasible posterior beliefs for three values of α. 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 set of feasible posteriors $\mathcal{P}_S$ that is narrow and concentrated around the prior. With more informative signals (e.g. $\alpha = 0.75$,

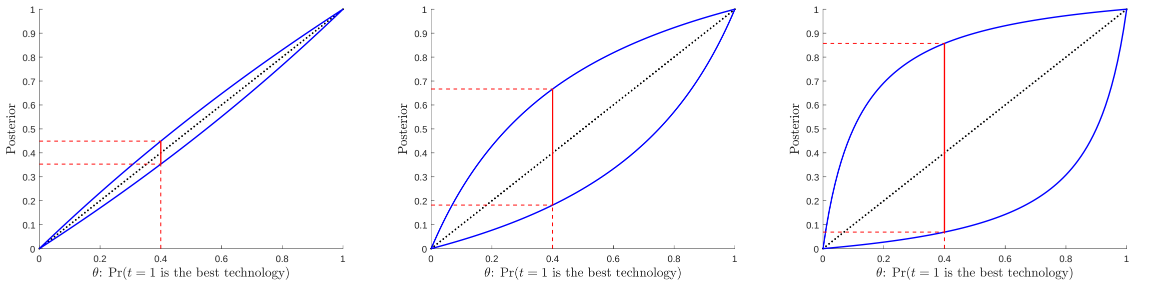

Figure 2: All panels show the range of posteriors for a prior of $\theta = 0.4$. Left panel: $\alpha = 0.55$, posterior in $[0.353, 0.449]$. Middle panel: $\alpha = 0.75$, posterior in $[0.182, 0.667]$. Right panel: $\alpha = 0.95$, posterior in $[0.069, 0.857]$.

Figure 2, middle panel), the set of feasible beliefs expands around the prior. As the signal becomes perfectly informative (e.g. $\alpha = 0.95$, Figure 2, right panel), the set of attainable posteriors expands even further.

Different experiment technologies map into different sets of posteriors. In our analysis, we take the common prior, Θ, and the set of feasible posteriors, $\mathcal{P}_S$, as the primitives of our model. Any signal structure can be translated into this framework. Our results then characterize the optimal information disclosure policy for any arbitrary primitives. In this sense our approach is more general than the standard Bayesian persuasion model because we allow for the principal to have access to an arbitrary set of signal structures, including ones that cannot perfectly resolve all uncertainty in the environment.

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

$$\nu(\Theta) \equiv \sum_t \theta_t f(\theta_t \Omega). \tag{5}$$

As in Kamenica and Gentzkow (2011), the value of information disclosure is described by the convexity of $\nu(\Theta)$. Let $\hat{\nu}(\Theta, \mathcal{P}_S)$ be the concave closure of $\nu(\Theta)$ over $\mathcal{P}_S$. 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 $\Theta = (\theta_1, ..., \theta_N) \in \Delta^N$ if and only if*

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

where $r_f(x) \equiv -\dfrac{xf''(x)}{f'(x)}$ 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 $r_f(\Omega) > 2$, the principal does not benefit from information disclosure; when $r_f(\Omega) < 2$, the principal always 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. Figure 3 (left panel) shows $\nu(\cdot)$ when $f(x) = x^{0.5}$. When $f(x) = x^a$ for any $a \in (0, 1)$, we have $r_f(x) = 1 - a < 2$ for any x. According to 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. These posterior beliefs induce more agents to allocate effort towards the most promising technologies relative to the allocation based on the prior belief. In other words, this signal structure will induce agents to *focus* on one technology in equilibrium, reducing diversification of effort across technologies.

When $f(x) = 1 - e^{-\lambda x}$, we have $r_f(x) = \lambda x$, and Lemma 2 shows that, depending on λ, $\nu(\cdot)$ might not be globally convex. Figure 3 (right panel) shows $\nu(\cdot)$ when $f(x) = 1 - \exp(-8x)$; ν is locally concave near $\theta = 0.5$ and convex towards the extremes. 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.

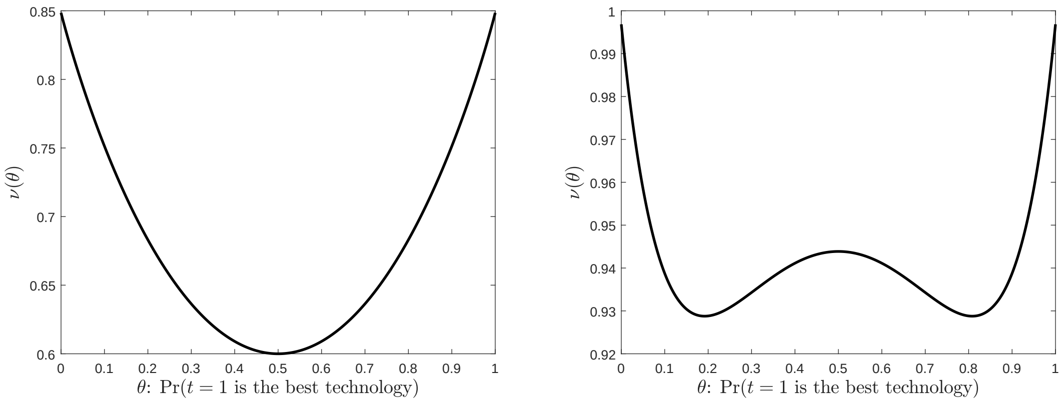

Figure 3: Function $\nu(\theta)$, the value for the principal of inducing posterior θ, for different functions $f(\cdot)$. Left Panel: $f(x) = \sqrt{x}$; Right Panel: $f(x) = 1 - \exp(-8x)$.

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 too little effort is allocated towards the most promising technology in proportional terms compared to the first-best effort allocation. This implies that in equilibrium there is too much diversification, 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, too much effort is allocated towards the most promising technology in proportional terms compared to the first-best effort allocation, so there is too little diversification in equilibrium. 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. it is not possible to generate posteriors far enough from the prior) 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 allocate relatively more effort 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 the gain from information disclosure (i.e. whether $\nu(\cdot)$ is locally convex or concave) depends on the value the 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 allocation of effort 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.

The next lemma shows that the shape of $\nu(\cdot)$ with $N > 2$ technologies is analogous to the two cases shown in Figure 3.

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, $\Theta = (1/N, ...1/N)$ 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 $\Theta = (1/N, ...1/N)$.

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. This, in particular, implies that all of the qualitatively different classes of cases for the shape of $\nu(\cdot)$ when there are two technologies are illustrated in Figure 3. 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 $\frac{\Omega}{N}$. In particular, if $r_f\left(\frac{\Omega}{N}\right) < 2$, then $\nu(\cdot)$ is convex and the solution is obvious: the principal always uses the most informative experiment, disclosing as much information as possible. If $r_f\left(\frac{\Omega}{N}\right) > 2$, then $\nu(\cdot)$ is locally concave at the center of the simplex and the principals' information design problem, which we analyze next, is more involved.

Additionally, Proposition 2 shows that $\nu(\cdot)$ 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. Thus every agent works on that technology in equilibrium, which is the optimal allocation for the principal. This immediately implies that, if the principal has access to a perfectly informative signal, revealing that signal is optimal. However, such a signal need not be available to the principal in general—in many settings it is unrealistic to assume that the principal has a perfect signal that can eliminate all uncertainty in the environment. In the Netflix Prize, for example, to implement such a signal Netflix would need to have infinite amounts of data on consumers' preferences.

Characterization of the Optimal Information Design. We now characterize the optimal signal structure in our setting. Let Θ_0 be the common prior belief profile over the N technologies. Conditional on that prior belief, there exists a convex set of posteriors $\mathcal{P}_S(\Theta_0) \subseteq \Delta^N$ that the principal can induce that are consistent with Bayes' rule (for an example, see Example 1). In particular, we always have that $\Theta_0 \in \mathcal{P}_S(\Theta_0)$ 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 $\mathcal{P}_S$.

Let $\Theta_C \equiv \{\Theta \in \Delta^N : \hat{\nu}(\Theta, \mathcal{P}_S) = \nu(\Theta)\}$ be the set of all posteriors $\Theta \in \Delta^N$ where the value function, ν, agrees with its concave closure, $\hat{\nu}$, over $\mathcal{P}_S$. Let $\widetilde{v} \equiv \sup\{\nu(\Theta) : \nu''(\Theta) \leq 0, \quad \Theta \in \mathcal{P}_S\}$ be the largest value of the value function over the region of $\mathcal{P}_S$ where it is concave.[7] We can now characterize the optimal disclosure policy.

Proposition 3. *The optimal disclosure policy, s^*, is*

1. ***maximally informative*** *if $\nu(\overline{\theta}) \geq \widetilde{v}$ for all $\overline{\theta} \in \partial\mathcal{P}_S$; then s^* induces a distribution over posterior beliefs with support consisting only of points in the boundary of the feasible set of posteriors, with distribution G_s s.t. $\mathrm{E}_{G_s}[\Theta] = \Theta_0$.*
2. ***partially informative*** *if $\nu(\overline{\theta}) < \widetilde{v}$ for some $\overline{\theta} \in \partial\mathcal{P}_S$ 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 $\mathcal{P}_S \cap \Theta_C$, with distribution G_s s.t. $\mathrm{E}_{G_s}[\Theta] = \Theta_0$.*
3. ***uninformative*** *if $\mathcal{P}_S \subseteq \Theta_C$; then s^* induces the prior $\Theta = \Theta_0$.*

The optimal signal structure is characterized with 3 different cases. These cases depend on three key features of the environment: (i) the value of diversification for the

[7] If this region is empty, $\widetilde{v} = -\infty$.

principal; (ii) how informative the principal's signals are for a given prior; and (iii) the extent of technological uncertainty. We illustrate the intuition for each of these features in the case of two technologies. First, if ν is globally convex, then the optimal signal structure is always maximally informative and the optimal persuasion experiment reveals results that lead to extreme posterior beliefs. On the other hand, in the case where the principal is sufficiently risk averse, i.e., when ν is not convex (see Figure 3), the optimal information design problem is more subtle, so we focus on this case in the example below.

Illustrative Example. Consider the case of two technologies and an arbitrary set of posteriors $\mathcal{P}_S = (\underline{\theta}_1, \overline{\theta}_1)$. For the sake of this example, we assume symmetry, so $0 \leq \underline{\theta}_1 \leq \overline{\theta}_1 = 1 - \underline{\theta}_1 \leq 1$, although this is not required for our results. We take the common prior, θ_0, and set of feasible posteriors, $\mathcal{P}_S$, as primitives of the model and illustrate the optimal information policy in Proposition 3 for different priors and feasible posteriors. Note that, for any prior and any set of feasible posteriors, there is implicitly some signaling technology that generates the latter from the former with Bayesian updating of the beliefs after the realization of signals. When we compare the optimal policy with various priors and different sets of feasible posteriors, the implicit signaling technology itself is also changing to keep the priors and posteriors Bayes-consistent. Our goal here is to illustrate how different priors and feasible posteriors give rise to different optimal policies.

Figure 4 shows the principal's payoff with/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 $\theta_1 = 0.2627$ and $\theta_1 = 0.7373$, respectively; (2) the unique interior 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 $\widetilde{v} = \sup\{\nu(\theta) : \nu''(\theta) < 0, \theta \in \mathcal{P}_S\} = f\left(\frac{1}{2}\right) = 1 - \exp(-\lambda/2)$, and we have some point $\hat{\theta}_1$ such that $\hat{\theta}_1 = \sup\{\theta : \hat{\nu}(\theta, [\underline{\theta}_1, \overline{\theta}_1]) = \nu(\theta)\}$.

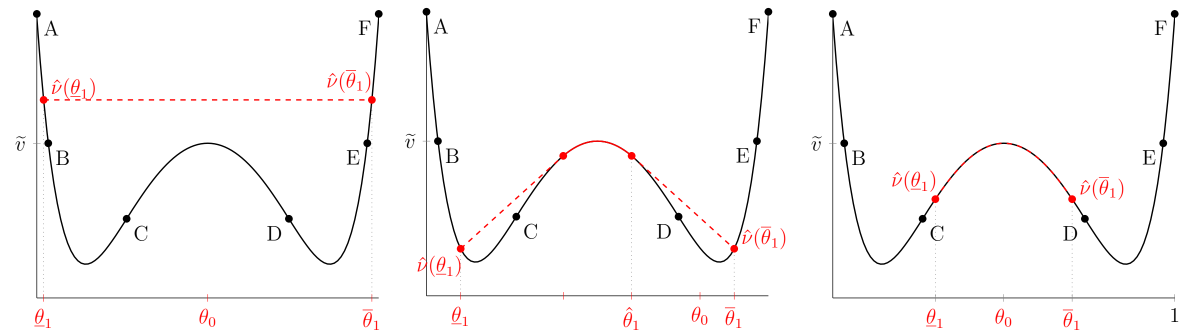

Figure 4: The solid black line corresponds to the principal's payoff without disclosure with $f(x) = 1 - \exp(-\lambda x)$, $\lambda = 8$, and $\Omega = 0.9$. The dashed red line corresponds to the function $\hat{\nu}(\cdot) = \text{concavification}(\nu)$ over the region of feasible posteriors $\mathcal{P}_S$. Each panel shows a different scenario for the set of feasible posteriors $\mathcal{P}_S = [\underline{\theta}_1, \overline{\theta}_1]$ and the prior is θ_0.
Left Panel: $\underline{\theta}_1 = 0.02$ and $\overline{\theta}_1 = 0.98$. Prior $\theta_0 = 0.5$.
Middle Panel: $\underline{\theta}_1 = 0.1$ and $\overline{\theta}_1 = 0.9$. Prior $\theta_0 = 0.8$.
Right Panel:$\underline{\theta}_1 = 0.3$ and $\overline{\theta}_1 = 0.7$. Prior $\theta_0 = 0.5$.

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. $\widetilde{v}$).

First, if the principal can generate highly informative signals, then $\mathcal{P}_S$ includes posterior beliefs close enough to 0 and 1. When $\min\{\nu(\underline{\theta}_1), \nu(\overline{\theta}_1)\} > \widetilde{v}$, a maximally-informative disclosure is optimal. Graphically, in Figure 4 (left panel), $\underline{\theta}_1$ lies somewhere between points A and B, and $\overline{\theta}_1$ lies somewhere between points E and F. In this case, the concave closure of $\nu(\cdot)$ is the line that connects $\nu(\underline{\theta}_1)$ and $\nu(\overline{\theta}_1)$. The optimal signal is one that reveals $\underline{\theta}_1$ with some probability q and $\overline{\theta}_1$ with the remaining probability $1-q$, where q is such that the expected posterior is equal to the prior, θ_0. 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. Thus, the principal benefits from inducing extreme posteriors. 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 $\nu(\underline{\theta}_1) < \widetilde{v}$ and $\nu(\overline{\theta}_1) < \widetilde{v}$. Moreover, suppose that the technologies are ex-ante asymmetric, with technology 1 more likely to be successful ex ante and θ_0 to the right of point D. Graphically, in Figure 4 (middle panel), $\underline{\theta}_1$ lies somewhere between points B and C, and $\overline{\theta}_1$ 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 $\overline{\theta}_1$ with the remaining probability $1 - q$, where q is such that the expected posterior is equal to the prior.

Third, suppose the set S only contains less informative signals, so the principal can induce beliefs in $\mathcal{P}_S$ only around 0.5 in the region where $\nu(\cdot)$ is concave. Graphically, in Figure 4 (right panel), $\underline{\theta}_1$ and $\overline{\theta}_1$ lie somewhere between points C and D. In this case the concave closure of $\nu(\theta_1)$ over $\mathcal{P}_S$ is equal to $\nu(\theta_1)$—there is no value from information disclosure, and the optimal signal is perfectly uninformative, inducing a posterior equal to the prior. The principal is constrained by the set of signals she can use to persuade the agents, which are relatively uninformative signals. 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.

To simplify the exposition, we focus on the case of two technologies, so $\Theta = (\theta, 1-\theta)$. From Lemma 2, the condition for concavity of ν at Ω when $N = 2$ is equivalent to $g(\Omega, \theta) < 0$, where

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

and $h(x) = xf'(x)[2 - r_f(x)]$. If $\theta^* \in [0, 0.5]$ is an inflection point of $\nu(\cdot)$, then $g(\Omega, \theta^*) = 0$. 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 level of effort in each technology. Because $f(\cdot)$ is concave, the marginal return to effort diminishes, so the value of focus becomes less important to the principal compared to the value of diversification. Hence the objective becomes more concave around the center of the simplex. 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).

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 $f(x) = 1 - \exp^{-8x}$ and $M = 5$ competitors. The function becomes concave around the center and the concave region expands as the prize or the number of agents (and hence Ω) increases: ν is convex when $V = 0.6$, but it is locally concave for $V = 0.8$ or $V = 1$, and the region of concavity expands.

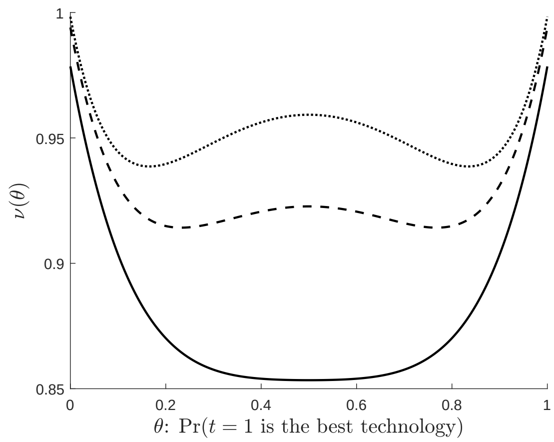

Figure 5: Function ν when $f(x) = 1 - e^{-8x}$ for different values of the prize $V = 0.6$ (solid line), $V = 0.8$ (dashed line), $V = 1$ (dotted line). The function ν is convex when $V = 0.6$; it is concave in the region $[0.2979, 0.7021]$ when $V = 0.8$; and it is concave in the region $[0.2570, 0.7430]$ when $V = 1$.

Next, suppose that the set of most informative experiments is generated by the specification in Example 1 with a parameter $\alpha = 0.75$. In this case, the boundary of the set of posteriors for each prior is illustrated in Figure 2 (middle panel). For instance, when the prior is $\theta = 0.4$, the set of attainable posteriors is $[0.182, 0.667]$. Figure 6 shows both $\nu(\cdot)$ and its concavification, $\hat{\nu}(\cdot)$, over the range of feasible posteriors $[0.182, 0.667]$ for the prior $\theta = 0.4$ under this experiment technology. 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.

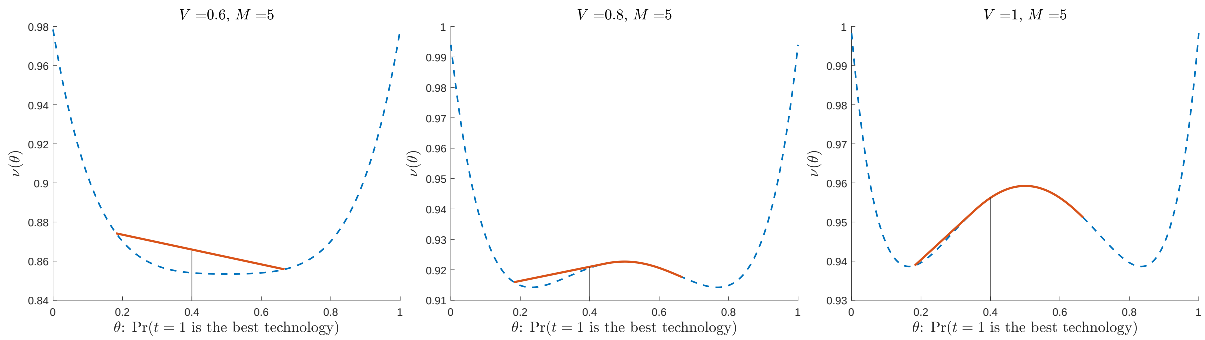

Figure 6: Function $\nu(\theta)$ when $f(x) = 1 - e^{-8x}$ for different values of the prize. The vertical line indicates the fixed prior 0.4. The set of attainable posteriors with the experiment technology in Example 1 is $[0.182, 0.667]$.

We can carry out this exercise for any prior belief $\theta \in [0, 1]$.

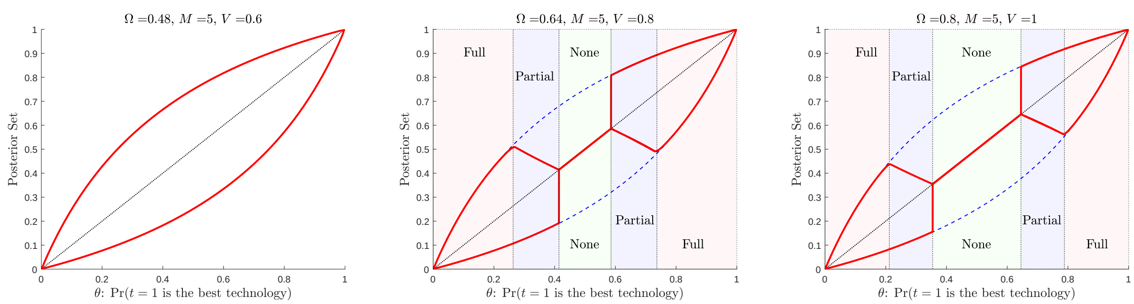

Figure 7: As Ω increases (from left panel to right panel), the extent of information disclosure decreases. For each prior, the solid (red) line shows the posteriors used in the optimal information structure. In the region denoted by "Full," there is full information disclosure, meaning that posteriors are the most extreme ones possible. In the region denoted by "Partial," there is partial information disclosure, meaning that posteriors are not the most extreme ones. In the region denoted by "None," there is no information disclosure, meaning that the posterior equals the prior.

Figure 7 shows the set of posteriors used in the optimal information disclosure policy, for each prior. In the left panel of Figure 7, when $\Omega = 0.6$, 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 $[0.21, 0.35] \cup [0.65, 0.790]$ and full disclosure 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 provided a condition that establishes when increasing the prize reduces the extent of information disclosure. Thus, 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, will the principal disclose information once the prize has been chosen optimally? Understanding the tradeoff between information disclosure and the optimal prize will show the principal's benefit of information disclosure once the prize is chosen endogenously. In principle, information disclosure could have no value once the prize is chosen optimally.

For this exercise, we illustrate this tradeoff for $f(x) = 1 - \exp^{-8x}$ with two technologies and five competitors, i.e., $N = 2$ and $M = 5$. Furthermore, we use the experiment technology in Example 1 with $\alpha = 0.75$. Our goal is to jointly optimize over both the information disclosure policy and the prize, V, for any prior θ and feasible posteriors $\mathcal{P}_S$.

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.

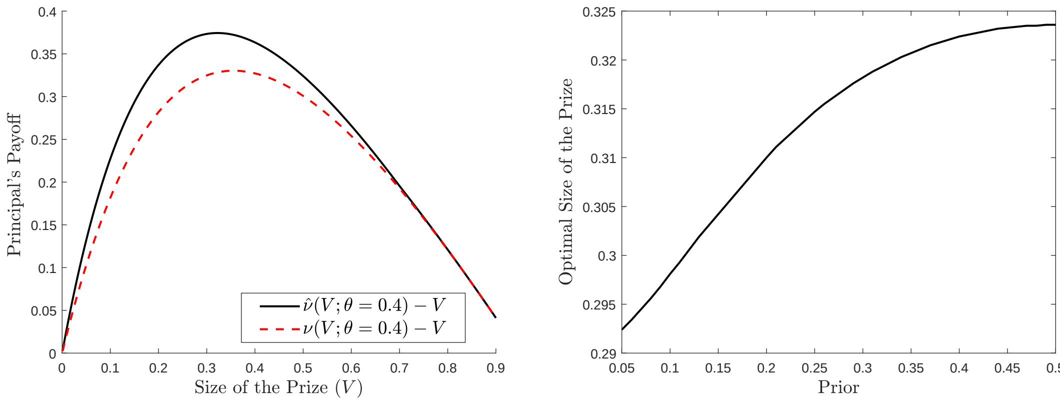

Figure 8: Left Panel: Functions $\hat{\nu}(\theta)$ and $\nu(\theta)$, the value for the principal of inducing posterior θ with and without information disclosure, for $f(x) = 1 - e^{-8x}$ for different values of the prize V.
Right Panel: The optimal size of V as a function of the prior belief, θ, assuming 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 addition, jointly maximizing over the prize and information disclosure actually results in a different optimal prize. In this example, absent any disclosure the optimal prize is $V = 0.356$ and the principal's expected payoff is 0.3303. In contrast, with the optimal disclosure policy the optimal prize is $V = 0.321$, and the principal obtains an expected payoff of 0.3744. 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 13.35 percent while reducing the size of the prize by 9.83 percent.

Figure 8 (right panel) shows the optimal prize when we repeat the exercise of finding the optimal prize for any prior $\theta \in [0, 0.5]$ and set of posteriors from Example 1, i.e., when the most informative experiment reveals the true technology with probability $\alpha = 0.75$. 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' efforts 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 any contest setting where agents can pursue different approaches, such as in procurement, contests for innovation, promotions within organizations, and others. All of these settings have in common the feature that 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.

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.

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 (2017) "Dynamic Tournament Design: An Application to Prediction Contests."

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 (2017) "Inducing Variety: A Theory of Innovation Contests."

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.

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. Agent i solves:

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

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

$$x^*_{i,t} = \theta_t \left(\frac{M-1}{M^2} \right) V. \tag{9}$$

Thus, the aggregate equilibrium effort on technology t is

$$x^*_t \equiv \sum_{i=1}^{M} x^*_{i,t} = \theta_t \Omega, \tag{10}$$

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

If the principal could control the agent's efforts, the principal's would choose the aggregate effort 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

$$y^*_t = h(\theta_t) \equiv (f')^{-1} \left(\frac{1}{\theta_t} \right). \tag{11}$$

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

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

The total level of effort in equilibrium 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 level of effort 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 effort in 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. The proof follows directly from the analysis in section 2.1. $\square$

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 a 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 messages was generated by s', and with probability $1-\alpha$ it was generated by s''. The receiver's posterior is the $\alpha\tilde{\theta} + (1-\alpha)\tilde{\theta}'$, so the set $\mathcal{P}_S$ is convex. $\square$

Proof of Lemma 2

Proof. If $\Theta = (\theta_1, ..., \theta_N) \in \Delta^{N-1}$ 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 $\nu(0, ..., 0, 1, 0, ..., 0) = 1 \cdot f(\Omega)$. The only points that can obtain the global maximum of $f(\Omega)$ are the vertices of the simplex. To see this, consider any point $\tilde{\theta} = (\theta'_1, ..., \theta'_N)$ 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 $\theta_N = 1 - \sum_{i=1}^{N-1} \theta_t$. 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 $t = 1, ..., N-1$ 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

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

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 $\Theta = (\theta_1, ..., \theta_N) \neq (1/N, ..., 1/N)$. Then there exist some indices i, j such that $\theta_{<}\theta_j$. 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 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 $\Omega[h'(\Omega\theta) - h'(\Omega(1-$

$\theta))]$. 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 if, 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 3

Proof. The proof follows from the convexity of $\mathcal{P}_S$ in autorefL-posteriors, 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 then 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 $\delta\mathcal{P}_S$ includes points towards 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. $\square$

B Appendix: 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 effort 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 total effort in equilibrium 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

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

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.

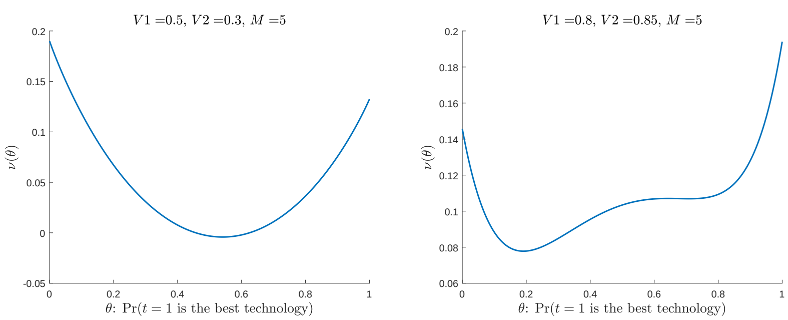

Figure 9: Left Panel: $f(x) = \sqrt{x}$, technology-specific prizes $V_1 = 0.5$ and $V_2 = 0.3$. Right Panel: $f(x) = 1 - \exp -8x$, technology-specific prizes $V_1 = 0.8$ and $V_2 = 0.85$.

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 $\Theta = (\theta_1, ..., \theta_N) \in \Delta^N$ if and only if*

$$\Omega_t \theta_t f'(\theta_t \Omega_t)[2 - r_f(\theta_t \Omega_t)] + \Omega_N \theta_N f'(\theta_N \Omega_N)[2 - r_f(\theta_N \Omega_N)] < 0, \text{ for all } t = 1, ..., N-1. \tag{12}$$

where $r_f(x) \equiv -\dfrac{xf''(x)}{f'(x)}$ *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.

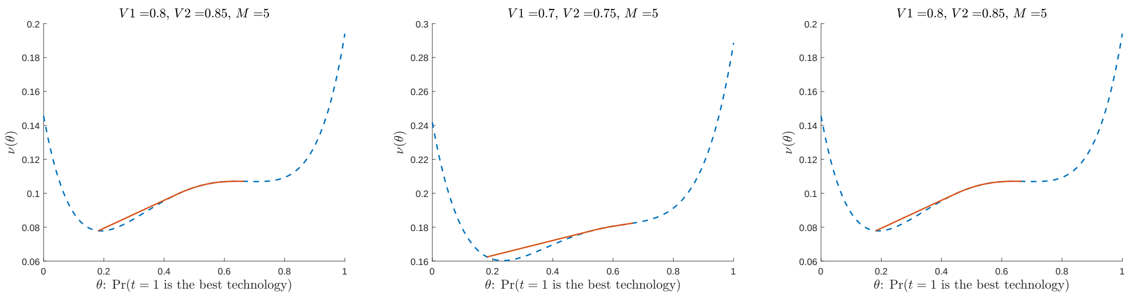

Figure 10: Value of information disclosure for three particular combinations of technology-specific prizes for $f(x) = 1 - \exp -8x$ when the prior is $\theta = 0.4$ and the most informative experiment is generated by $\alpha = 0.75$ in Example 1. Left Panel: $V_1 = 0.6$ and $V_2 = 0.65$. Left Panel: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

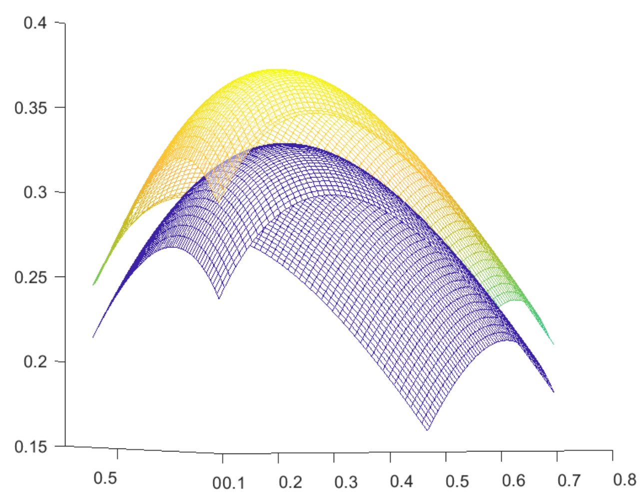

Figure 11: The value of information disclosure for different combinations of technology-specific prizes for $f(x) = 1 - \exp -8x$, when the prior is $\theta = 0.4$ and the most informative experiment is generated by $\alpha = 0.75$ in Example 1.

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 $V_1 = 0.35 < V_2 = 0.37$.