Optimal Feedback Dynamics Against Free-Riding in Collective Experimentation*

Chia-Hui Chen† Hülya Eraslan ‡ Junichiro Ishida§
Takuro Yamashita¶

September 28, 2024

Abstract

We consider a dynamic model in which a principal decides what information to release about a product of unknown quality (e.g., a vaccine) to incentivize agents to experiment with the product. Assuming that the agents are long-lived and forward-looking, their incentive to wait and see other agents' experiences poses a significant obstacle to social learning. We show that the optimal feedback mechanism to mitigate information free-riding takes a strikingly simple form: the principal recommends adoption as long as she observes no bad news, but only with some probability; once she does not recommend at some point, she stays silent forever after that. Our analysis suggests the optimality of premature termination, which in turn implies that: (i) false positives (termination in the good state) are more acceptable than false negatives (continuation in the bad state); (ii) overly cautious mechanisms that are biased toward termination can be welfare-enhancing.

JEL Codes: D82, D83

Keywords: strategic experimentation, information design, free-rider problem, social learning, COVID-19 pandemic.

---

*We thank Nejat Anbarcı, Kim-Sau Chung, Daisuke Hirata, Fahad Khalil, Gilat Levy, Meg Meyer, Yosuke Yasuda and participants at Workshop on Contracts, Incentives and Information (3rd Edition) at Turin, International Workshop on Theory and Experiments at Osaka, Kobe, Hong Kong Baptist, CTW, OEIO, Academia Sinica, and the 2024 Asia Meeting of the Econometric Society for helpful feedback. We acknowledge financial support from National Science Foundation (Grant SES-1730636 for Eraslan, JSPS Kakenhi (23K01309 for Chen, 21H04398 and 23K01347 for Ishida, and 23K01311 for Yamashita), and the Joint Usage/Research Center for Behavioral Economics at ISER.

†Institute of Economic Research, Kyoto University; chchen@kier.kyoto-u.ac.jp.

‡Department of Economics, Rice University, NBER and Institute of Social and Economics Research, Osaka University; eraslan@rice.edu.

§Institute of Social and Economic Research, Osaka University; jishida@iser.osaka-u.ac.jp.

¶Osaka School of International Public Policy, Osaka University; yamashita.takuro.osipp@osaka-u.ac.jp.

1 Introduction

The COVID-19 pandemic presented an unprecedented challenge, compelling governments around the world to navigate through uncertainties and introduce many new, often untested, measures to cope with it. Consider the uptake of mRNA vaccines as an example. While this radically new technology had been known for years and considered safe among medical scientists, it had never been widely used before the pandemic, and politicians and citizens alike were not entirely sure about its possible adverse consequences. Amid this uncertainty, a potential impediment to widespread adoption is citizens' unwillingness to be among the first to try out the new technology: Hamel, Lopes and Brodie (2021) document that "31% of the public say that when an FDA-approved vaccine for COVID-19 is available to them for free, they will 'wait until it has been available for a while to see how it is working for other people' before getting vaccinated themselves." Those who stand to gain less from the vaccine may choose to wait and see others' experiences, but this kind of wait-and-see attitude slows down the adoption process, resulting in a loss of surplus that can be substantial.

The presence of information free-riding is in fact ubiquitous in many instances of policy experimentation that require the voluntary participation of citizens; examples include technology adoption in a variety of sectors such as healthcare, education, and manufacturing. In close inspection, two aspects of this experimentation process are particularly noteworthy. First, experimentation of this kind is often "feedback-based" in that uncertainty regarding the net value of a new policy is revealed gradually through the feedback of the citizens who have embraced the experimentation. Second, it is also often "large-scale," involving a large number of citizens who interact anonymously, such that they do not directly observe what others have experienced on the whole, e.g., the frequencies of severe side effects, and must instead rely on the information released by central authorities. The first aspect makes incentivizing the citizens to experiment a central problem for the government, whereas the second aspect provides a potentially viable tool to achieve this goal. Along with the fact that contingent monetary transfers are often not feasible in this type of environment, these two aspects give rise to a new class of information-design problems in which the government determines what information to disseminate to the public, with the aim of overcoming information free-riding and facilitating social learning.

To address this issue, we consider a collective experimentation environment that consists of a principal (e.g., a government) and a continuum of long-lived and forward-looking agents (e.g., citizens). The principal has a product to be consumed (or a technology to be adopted) by the agents who are heterogeneous with respect to their valuations of the product. The quality of the product is initially unknown to anyone and gradually revealed via the arrival of news whose arrival rate is proportional to the measure of agents who consume the product. However, reflecting the large-scale nature of the experimentation, news is observable only to the principal, who therefore faces a problem of what information to disseminate to the agents. Formally, the principal designs and commits to a feedback mechanism that specifies the distribution of the sequence of beliefs subject to the Bayes plausibility. Given some feedback mechanism, each agent independently decides when, if ever, to adopt the product based on his expectation of how the belief evolves as dictated by the mechanism in place.

We start by analyzing a bad-news case with a benevolent principal, i.e., the case where bad news (a "breakdown") arrives only when the state is bad, and the principal maximizes the discounted sum of the payoffs of all agents. We first observe that in the environment described above, the amount of adoption under full disclosure is generally insufficient compared to its first-best level. This observation stems from the free-rider problem that has been well documented in the strategic experimentation literature (Bolton and Harris, 1999; Keller, Rady and Cripps, 2005): because information is a public good whose benefits are not fully internalized by the agents, each agent has an excessive incentive to wait and see the experimentation outcomes of other agents. However, in large-scale experimentation where the agents have no direct access to the outcomes, the extent of the free-rider problem depends crucially on what information the principal disseminates to the agents. This point can be seen most clearly if we consider an extreme measure of no disclosure: if the principal commits to revealing no information, there is no point in waiting for news, and the free-rider problem dissipates completely. Of course, revealing no information is never optimal because that would imply entirely giving up the gains from additional information gleaned from the experimentation. The principal thus needs to strike the right balance between exploiting the gains from the experimentation on one hand and managing the free-rider incentive on the other.

The principal's problem is potentially complicated because the set of mechanisms is large in our dynamic context, with a myriad of choices available for the principal. Specifically, the principal can choose not only what information to disclose in a given period but also when to do so: for instance, she may withhold favorable information by garbling signals in some period and deliver it later with a delay. Despite this potential complication, we show that the optimal feedback mechanism to mitigate information free-riding takes a strikingly simple form: the principal recommends adoption as long as she observes no news, but only with some probability; once she stops recommending, she will never recommend again even if she has observed no news. Owing to this result, the principal's problem effectively reduces to an optimal-stopping problem, and the resulting optimal mechanism can be characterized thoroughly by a sequence of the continuation probabilities conditional on having observed no news.

This characterization result is a consequence of two observations, which we label as *termination* and *caution* for clarity. The termination property suggests that if the principal suspends the adoption process for a period, she will never resume the process again, even if she has received no bad news. This property alternatively implies that there is no strategic gain from withholding information to slow down adoption, and the optimal mechanism entails no delayed information release, which enables us to substantially reduce the set of mechanisms we need to consider. The caution property suggests that the optimal mechanism should exercise caution by being fully revealing when the principal is relatively pessimistic while randomizing when she is relatively optimistic. This latter result pertains to the fact that the value of information varies across agents with different valuations: information that distinguishes small risk and no risk is crucial for agents with lower valuations but irrelevant for those with higher valuations. The optimal mechanism can mitigate the free-rider problem and achieve the optimal outcome because it provides no useful information for agents with higher valuations, similar to the no-disclosure policy, but still provides enough information to improve the welfare of those with lower valuations. Both of these properties are largely detail-free and robust to various alterations to the underlying setup, as a consequence of which our main characterization result holds under a more general objective function that accommodates a wider range of social goals (Section 6.1) as well as in an alternative good-news scenario where news (a "breakthrough") arrives only when the state is good (Section 6.2).

An important economic insight of our analysis is the optimality of premature termination, i.e., it is welfare-enhancing for the principal to occasionally terminate the adoption process even when she has observed no bad news and is still relatively optimistic. This finding provides some practical implications for how the government should structure its decision-making process—the process in which it collects and analyzes data to produce evidence and transforms it into a recommendation. First, the optimality of premature termination suggests that false positives (termination in the good state) are more acceptable than false negatives (continuation in the bad state). This finding in turn implies that if the government can design its experiments to determine the precision of the information-generating process, it should direct more resources to minimizing false negatives while accepting false positives to some extent. Second, the optimality of premature termination also points to a virtue of overly cautious mechanisms that are excessively biased toward termination. It is worth noting that we obtain this conclusion even though all the concerned parties are assumed to be risk-neutral.

Literature. Our model lies at the intersection of strategic experimentation (Bolton and Harris, 1999) and information design (Kamenica and Gentzkow, 2011). Our learning environment builds on the exponential-bandits model of Keller et al. (2005).[1] A crucial departure from this strand of literature is that we consider large-scale collective experimentation that involves a large number of agents, where each individual agent cannot directly observe the experimentation outcomes of other agents and must instead rely on the information released by the principal. This aspect of our model draws a clear contrast to canonical models of strategic experimentation that typically focus on a relatively small group of agents and assume that each agent can directly observe the experimentation outcomes of other agents (Keller et al., 2005; Bonatti and Hörner, 2011). The fact that only the principal can observe the experimentation outcomes gives her some leeway to control the agents' belief formation process, amounting to a new class of information-design problems as noted above.

Our analysis is in spirit most closely related to Kremer, Mansour and Perry (2014), Che and Hörner (2018), and Knoepfle and Salmi (2024) in that they all explore ways to facilitate social learning, with emphasis on agents as both consumers and generators of information.[2] Kremer et al. (2014) consider a social-learning environment à la Bikhchandani, Hirshleifer and Welch (1992) and characterize the optimal disclosure policy to mitigate information herding. Che and Hörner (2018) consider a similar feedback-based information structure to ours in which information arrives at a rate proportional to the amount of adoption and study how a recommender system can improve the incentives for early exploration. In both of these models, however, agents are assumed to be myopic and make once-and-for-all adoption decisions. By contrast, in our setting, agents are forward-looking and strategically choose their timing of adoption.[3] Within this framework, we focus on a different incentive issue—free-riding stemming from the option to wait and see others' experiences—and explore the optimal feedback mechanism to alleviate this problem. Knoepfle and Salmi (2024) consider forward-looking agents as we do but focus their attention on the optimal timing of disclosure, while we study optimal information design. Also, agents are heterogeneous with respect to discount rates in their model, while with respect to valuations in our model. The differences are crucial, as the two papers deliver different insights focusing on different aspects of an otherwise similar

[1] More precisely, our baseline model follows the bad-news ("breakdown") specification of Keller and Rady (2015), although we later extend our analysis to the good-news ("breakthough") case.

[2] Baccara, Levy and Razin (2024) consider an environment in which there are two risky "fields" to choose from, and each agent can irreversibly join either field or wait for more information. In this framework, they also assume that the rate of information arrival in each field depends on the measure of agents who joined that field.

[3] Frick and Ishii (2024) also study social learning by forward-looking agents but without information design.

environment.[4] In this sense, their analysis and ours are complementary to each other.

Several works explore how to incentivize experimentation via monetary transfers. To name a few, Manso (2011) analyzes a moral-hazard problem with unknown success probability and shows that the optimal incentive contract exhibits substantial tolerance for early failure. Halac, Kartik and Liu (2016) consider a model of long-term contracting with both moral hazard and adverse selection and obtain a characterization of optimal menu contracts. Halac, Kartik and Liu (2017) consider both monetary rewards and information disclosure (about whether success had occurred or not) in a contest environment where multiple agents compete for a prize. We view our analysis as complementary to this strand of literature, where we focus exclusively on the use of information to motivate agents to take a risky alternative without monetary rewards. Our analysis is more applicable to situations where contingent monetary transfers are either unconventional or infeasible due to institutional factors, as is often the case in the context of policy experimentation.

Finally, we also aim to contribute more broadly to the literature on dynamic information design. There are several recent works that analyze information disclosure in dynamic environments where an agent makes stopping decisions (Au, 2015; Ely, 2017; Nikandrova and Pancs, 2018; Che and Mierendorff, 2019; Ely and Szydlowski, 2020; Orlov, Skrzypacz and Zryumov, 2020).[5] There are two notable differences from this strand of literature. First, while those previous works generally focus on single-agent cases, we consider an environment where the principal attempts to persuade a continuum of agents who are heterogeneous with respect to their valuations. Second, and more importantly, the principal's information structure in our setting is endogenous in that the precision of the information she acquires depends on the agents' adoption decisions. This aspect of our model stands in sharp contrast to the standard setting where the information structure is exogenously given and fixed over time.

[4] They find that it is optimal to promptly disclose conclusive bad news and terminate the adoption process (but delay disclosure of conclusive good news); unless the principal observes bad news, the adoption process continues. On the contrary, we find that, when the agents are sufficiently forward-looking, it is optimal to terminate the adoption process with some probability even if no bad news has arrived. This *caution* property is key to incentivizing agents to contribute to social learning in earlier periods who otherwise may free-ride on others, and this cannot be implemented simply by timing the disclosure of bad news as in Knoepfle and Salmi (2024). Aside from the fact that they restrict attention to the timing of disclosure, the difference in the nature of agent heterogeneity also plays a major role: our optimal mechanism exploits differences in adoption thresholds stemming from the heterogeneity in valuations, which is not considered in Knoepfle and Salmi (2024) where the agents are assumed to possess the same valuation.

[5] Ball (2023) and Correia da Silva and Yamashita (2024) consider information disclosure in repeated interactions.

2 Illustrative example

We first provide a simple example with two periods and two agent types to highlight the main ideas. Suppose a government attempts to introduce a new vaccine to a continuum of citizens with unit measure. Each citizen can take the vaccine in either period or can choose to abstain from it altogether. The quality of the vaccine, which is initially unknown to anyone, can be either good ($\omega = 0$) or bad ($\omega = 1$), where the common prior probability that the vaccine is bad is $m \in (0, 1)$. The citizens are divided into two preference types, where a half of them have low valuation ($v = v_\ell$) and the other half have high valuation ($v = v_h > v_\ell$). If a citizen with valuation v chooses to take the vaccine, his payoff is $v - \omega$; if not, his payoff is normalized to 0. We assume $v_h < 1$ so that even the high-valuation citizens would not take the vaccine if they knew that the vaccine is bad. For simplicity, we assume that there is no time discounting between the two periods.

The quality of the vaccine is partially revealed via the arrival of news at the end of period 1. Here, we consider a bad-news scenario where news (e.g., the occurrence of severe side effects) arrives only when the state is bad. We make two crucial assumptions regarding the information structure. First, as news reflects the aggregate outcome, it is observable only to the government but not to each citizen. Second, conditional on the state being bad, the likelihood of news arriving depends on and is increasing in the measure of citizens who choose to take the vaccine. Specifically, after a measure n of citizens take the vaccine in period 1, the government observes bad news with probability $p(n)\omega$, where $p : [0, 1] \to [0, 1]$ is some increasing function. Let μ_2^p denote the government's belief (that the state is bad) in period 2. If bad news is observed, the government learns that the vaccine is bad for sure, and hence $\mu_2^p = 1$. If not, taking n as given, the government's belief is updated to

$$\mu_2^p = \underline{\mu}(n) := \frac{m(1 - p(n))}{m(1 - p(n)) + 1 - m},$$

which is lower than the initial prior for any $n > 0$, i.e., "no news is good news."

The government's problem is to determine the distribution of period-2 beliefs of the citizens subject to the Bayes plausibility, so as to maximize the sum of the payoffs of all citizens. Let μ_2 denote the citizens' period-2 belief that the vaccine quality is bad, i.e. $\omega = 1$, based on the information released by the government. Since period 2 is the last period, the problem faced by the citizens in that period is straightforward: given some belief μ_2, a citizen with valuation v takes the vaccine if $v \geq \mu_2$. Taking this as given, in period 1, a citizen with valuation v takes the vaccine if

$$v - m \geq \mathbb{E}_\mu[\max\{v - \mu, 0\}]. \tag{1}$$

Each citizen's decision in period 1 depends on the continuation payoff, which in turn depends on the distribution of period-2 beliefs. For illustration, assume $1 > v_h > m > v_\ell > \underline{\mu}(0.5)$, so that the low-valuation citizens would never take the vaccine in period 1 but could be persuaded to do so in period 2 if enough information were generated in period 1.

To illustrate the value of information design in this context, we begin with two extreme cases as benchmarks. We first consider the full-disclosure policy where the government mechanically reveals everything it observes (so that $\mu_2 = \mu_2^p$ with probability 1). In this case, if a high-valuation citizen waits until period 2, he takes the vaccine in period 2 if and only if no news is observed, which occurs with probability $m(1 - p(n^*)) + 1 - m$. The expected payoff of taking the vaccine in this contingency is $v_h - \underline{\mu}(n^*)$, where n^* is the equilibrium measure of agents who take the vaccine in period 1. Since the expected payoff of taking the vaccine in period 1 is $v_h - m$, the IC constraint for the high-valuation citizens to take the vaccine in period 1 can be written as

$$\begin{aligned} v_h - m &\geq [m(1 - p(n^*)) + 1 - m](v_h - \underline{\mu}(n^*)) \\ &= [m(1 - p(n^*)) + 1 - m]v_h - m(1 - p(n^*)), \end{aligned}$$

taking n^* as given.[6] It is easy to verify that the IC constraint does not hold for any $n^* > 0$ for any high-valuation citizen, and therefore, it does not hold for any low-valuation citizen either. As such, no one chooses to take the vaccine in period 1 under full disclosure. This is a manifestation of the free-rider problem, and as a consequence, the opportunity for social learning is entirely lost, leading to a suboptimal outcome.[7]

In a sense, the full-disclosure policy fails because it is too informative, giving the high-valuation citizens an excessive incentive to wait for news. This implies that the optimal feedback mechanism must be somewhat more obscure than the full-disclosure policy. To illustrate this point, we now turn to the opposite case of no disclosure where the government commits to revealing no additional information (so that $\mu_2 = m$ with probability 1). In this case, there is clearly no point in waiting for news, and it is (weakly) optimal for the high-valuation citizens

[6] Throughout the analysis, we assume that an agent chooses to adopt if he is indifferent.

[7] Although we consider a simple binary-type distribution and no time discounting, the observation that the amount of adoption is insufficient can be extended to any (continuous) type distribution and discount rate. See Appendix B for a more formal argument.

to take the vaccine in period 1,[8] resulting in the maximum amount of information that can be generated in this environment. However, this no-disclosure policy is also suboptimal because none of the low-valuation citizens could be induced to take the vaccine in period 2, given $\mu_2 = m$. This argument suggests that the government must find a middle ground between the two extreme policies.

So, what is the optimal way to disclose information in this environment? To achieve the optimal outcome, the government must induce (i) the high-valuation citizens to take the vaccine in period 1 and (ii) the low-valuation citizens to do so in period 2 occasionally (in case no news is observed). We have already observed that the no-disclosure policy achieves (i) but not (ii). Note that the no-disclosure policy is effective for (i) because it provides no valuable information to the high-valuation citizens in that their behavior does not depend on the information released in period 2. This reasoning suggests, however, that any mechanism that has this feature can achieve (i) as well. Specifically, we modify the no-disclosure mechanism by applying a mean-preserving spread to m and splitting it into $\underline{\mu}(0.5)$ and v_h. This modified mechanism gives the high-valuation citizens the same continuation payoff and hence continues to achieve (i) because the support of the belief distribution lies entirely to the left of their indifference point v_h; in other words, the extra information provided by the mean-preserving spread is irrelevant for the high-valuation citizens. It is, however, beneficial for the low-valuation citizens because the support spans over their indifference point $\mu_2 = v_\ell$, allowing them to make more informed decisions in period 2. Figure 1 graphically illustrates this situation.

This argument suggests that the government must provide the most accurate information, i.e., $\mu_2 = \underline{\mu}(0.5)$, whenever it chooses to continue (recommending the low-valuation citizens to adopt) in period 2. This in turn implies that the government must surely terminate if it has observed bad news. If it has observed no news, it may continue with some probability less than 1 (given that full disclosure is not optimal in this example). To sum up, the government can alleviate information free-riding and achieve the optimal outcome by the following mechanism:

- If the government observes bad news, it terminates with probability 1;
- If the government observes no news, it continues with some probability $\beta \in (0, 1)$ and terminates with the remaining probability.

This proposed mechanism is characterized by the continuation probability β, which measures the degree of transparency with $\beta = 1$ ($\beta = 0$) corresponding to full (no) disclosure. It is fully

[8] The high-valuation citizens are indifferent between taking the vaccine and waiting for news because of no time discounting. The incentive can be made strict if they discount future payoffs even slightly.

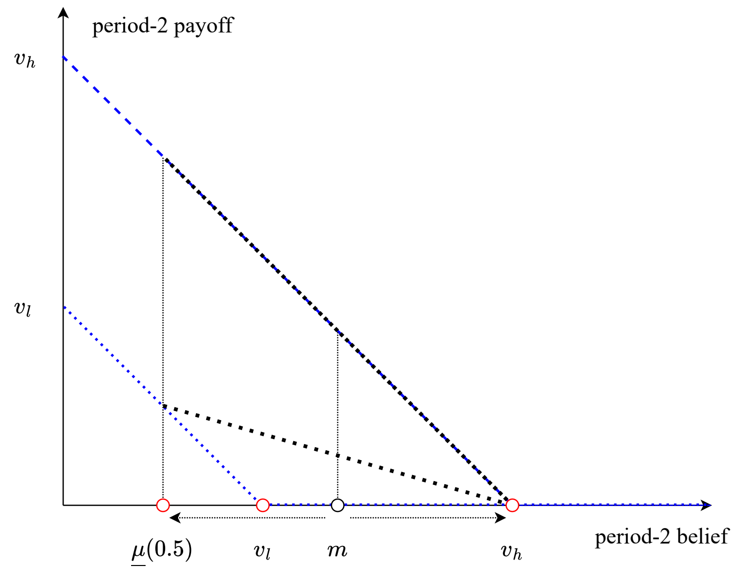

Figure 1: The dashed line at the top represents the value function of the high-valuation citizens, and the dotted line at the bottom represents the value function of the low-valuation citizens. A mean-preserving spread of m into $\underline{\mu}(0.5)$ and v_h raises the expected payoff of the low-valuation citizens while keeping the expected payoff of the high-valuation citizens constant.

revealing when the government is pessimistic (has observed bad news) but randomizes when it is optimistic (no news). The reason why such a mechanism works pertains to the fact that the value of information differs across different preference types: more precise information at the low end of the belief distribution is valuable for the low-valuation citizens but is irrelevant for the high-valuation agents. Since the government must persuade the high-valuation citizens to adopt early to generate the information necessary to persuade the low-valuation citizens, a mechanism that is precise at the low end but obscure at the high end works to achieve the optimal outcome. Remarkably, this observation can be extended to any (continuous) type distribution and any number of periods as we will detail below, even though the problem becomes substantially more complicated.

This illustrative example is simple enough to explicitly obtain the optimal mechanism. We now consider an arbitrary discount factor $\delta \in (0, 1]$ and let $p(n) = 1 - e^{-\lambda n}$ for concreteness. Since the low-valuation citizens always benefit from more information, the optimal mechanism is such that it provides the most accurate information to the extent that it does not affect the incentive of the high-valuation citizens. This implies that the high-valuation citizens must be held indifferent between adopting and waiting. If we let β^* be the optimal continuation probability, it must solve

$$v_h - m = \delta[\beta^*(me^{-0.5\lambda} + 1 - m)v_h - me^{-0.5\lambda}].$$

Thus, the optimal continuation probability is given by

$$\beta^* = \min\left\{\frac{v_h - m(1 - \delta e^{-0.5\lambda})}{\delta(me^{-0.5\lambda} + 1 - m)v_h}, 1\right\}.$$

Some useful observations can be drawn from this.

1. A higher δ (more patience) leads to less transparency.

2. A lower v_h leads to less transparency.

3. An increase in λ (more efficient feedback) leads to less transparency.

4. A higher m (more pessimistic outlook) leads to less transparency.

As is intuitively clear, anything that aggravates the free-rider problem—more patience, low valuations, more efficient feedback, and a more pessimistic outlook—makes the optimal mechanism less transparent (a lower β^*). Whenever we have $\beta^* < 1$, although lowering the continuation probability in the optimistic state necessarily entails a loss of surplus, the gain from alleviating information free-riding is more than enough to compensate for the loss.

3 Model

Environment. Consider the same experimentation environment in which a principal (e.g., a government) has a product to be consumed by a continuum of agents (e.g., citizens). Time is discrete and denoted by $t = 1, \ldots, T$ where T is the terminal period. In each period t, each agent decides whether to "adopt" the product or "wait" for news; if an agent chooses to adopt, he immediately leaves the game. The cost of adopting the product is determined by the state of nature $\omega \in \{0, 1\}$, which is initially unknown to anyone, where $\omega = 0$ denotes the good (low-cost) state while $\omega = 1$ denotes the bad (high-cost) state. The common prior of the state being bad is $m \in (0, 1)$.

Information. The true state is gradually revealed via the arrival of a signal (news). We primarily consider a bad-news ("breakdown") case where a signal arrives only when the state is bad; in Section 6.2, we extend the analysis and consider the opposite case where a signal (or a "breakthrough") arrives only when the state is good.[9] Let $s_t \in \{b, \varnothing\}$ be the signal realized at the end of period t, where $s_t = b$ indicates bad news while $s_t = \varnothing$ indicates no news. The probability of receiving signal b at the end of period when the state is ω and the measure of agents who adopt in period t is n_t is given by

$$\mathbb{P}[s_t = b \mid \omega, n_t] = 1 - e^{-\lambda \omega n_t}.$$

This suggests "no news is good news." We assume that only the principal can observe signals while the agents do not. At the beginning of each period t, the principal updates her belief (that the state is bad) upon observing s_{t-1} which we denote by μ_t^p with $\mu_1^p = m$.

Mechanism. The principal designs and commits to a feedback mechanism P at the outset of the game. The mechanism specifies the distribution of the sequence of beliefs μ_t subject to the Bayes plausibility and her own information structure, where μ_t denotes the agents' period-t belief. The public history of the game at the beginning of period t can be summarized by the sequence of realized beliefs up to that point, which we denote by $\mu^{t-1} := (\mu_1, \ldots, \mu_{t-1})$. The mechanism can then be written as $P = (P_t(\cdot \mid \mu^{t-1}))_{t=1}^T$, where $P_t(\cdot \mid \mu^{t-1})$ is the probability mass function of period-t beliefs conditional on the history of the game. Let $M_t(\mu^{t-1})$ denote the support of $P_t(\cdot \mid \mu^{t-1})$, i.e., $M_t(\mu^{t-1})$ is the smallest set M such that $\sum_{\mu_t \in M} P(\mu_t \mid \mu^{t-1}) = 1$.

Payoffs. Each agent is characterized by his valuation $v \in [0, 1]$, which is distributed according to some distribution F with full support over $[0, 1]$. Let f denote the corresponding density. If an agent with valuation v (hereafter, simply agent v) adopts in period t, he earns $v - \omega c$ and leaves the game; if not, his payoff for the period is normalized to 0. Throughout the analysis, we maintain the following assumption to focus our attention on relevant cases.

Assumption 1 *$c \geq 1 > mc$.*

The first inequality ensures that learning is essential for all agents,[10] while the second inequality rules out the trivial case where no agents adopt. The agents maximize the discounted sum of expected payoffs with a discount factor $\delta \in (0, 1)$. For now, we assume that the principal is benevolent and maximizes the discounted sum of the payoffs of all agents, which we call the total surplus for clarity, with the same discount factor δ; in Section 6.1, we extend the analysis

[9] We focus on the bad-news case because it is more relevant in many instances of policy experimentation. For instance, in the context of vaccine uptake, it is more likely that the government would receive feedback when the vaccine is of bad quality.

[10] If $c < 1$, it is optimal for agents $v \in [c, 1]$ to adopt regardless of the state. The optimal decision for those agents is hence trivial in that they always adopt immediately in period 1 in any mechanism.

to incorporate a more general objective function that can capture other factors such as payoff externalities and the welfare of third parties (e.g., firm profit).

4 Analysis

In the two-period example, there is only one possible history, which is $\mu_1 = m$, and any disclosure mechanism is entirely summarized by a single distribution of period-2 beliefs $P_2(\cdot \mid m)$. By contrast, when there are more than two periods, the problem becomes substantially more complicated because the set of possible histories could expand exponentially over time. In the general setting, we need to consider how the current distribution of beliefs potentially affects their future evolution in the continuation game, thereby giving rise to a much wider array of feedback mechanisms.

4.1 Preliminaries

Before we proceed to the general analysis, we first provide a more detailed account of three key constraints in this general setup: IC constraint, Bayes plausibility, and consistency. Let v_t^* be the threshold type who is indifferent between adopting and waiting in period t and let $v_0^* = 1$. Once a mechanism P is given, we can pin down a non-increasing sequence of thresholds up to period t for any given history μ^t. More precisely, fix a mechanism P and let H_t^s be the set of all subsequences $\mu^{t,t+s} := (\mu_{t+1}, \ldots, \mu_{t+s})$ such that $v_{t+s}^* < v_t^* = v_{t+s-1}^*$ under P. Then, for a given history μ^t and the corresponding sequence of thresholds up to period $t-1$, the threshold type in period t is determined by

$$v_t^* - \mu_t c = \sum_{s=1}^{T-t} \sum_{\mu^{t,t+s} \in H_t^s} \delta^s P_{t+1}(\mu_{t+1}|\mu^t) \cdots P_{t+s}(\mu_{t+s}|\mu^{t+s-1})(v_t^* - \mu_{t+s}c), \tag{2}$$

if a solution $v_t^* \in [0, v_{t-1}^*)$ exists. If no solution exists in $[0, v_{t-1}^*)$, no agents are induced to adopt in period t, and we let $v_t^* = v_{t-1}^*$. In what follows, we call (2) the IC constraint for clarity.

The principal's problem is restricted by two conditions in our setting. First, as usual, the Bayes plausibility requires that the expectation of posteriors must coincide with the prior conditional on μ^{t-1}, i.e.,

$$\sum_{\mu_t \in M_t(\mu^{t-1})} \mu_t P_t(\mu_t \mid \mu^{t-1}) = \mu_{t-1}. \tag{3}$$

Second, any feasible mechanism must also be consistent with the principal's private information. Specifically, at any point in time, the principal is in either one of the two information states: *pessimistic* (has observed bad news) or *optimistic* (has observed no news). In the bad-news case, the principal's belief is always $\mu_t^p = 1$ in the pessimistic state, and

$$\mu_t^p = \underline{\mu}(v_{t-1}^*) := \frac{me^{-\lambda(1-F(v_{t-1}^*))}}{me^{-\lambda(1-F(v_{t-1}^*))} + 1 - m}$$

in the optimistic state. Since the principal mixes these two information states to generate a belief, any feasible belief must be bounded between $\underline{\mu}(v_{t-1}^*)$ and 1. For later use, we write this consistency restriction more generally as

$$M_t(\mu^{t-1}) \subset [\underline{\mu}(v_{t-1}^*), \overline{\mu}(v_{t-1}^*)], \tag{4}$$

where $\overline{\mu}(v_{t-1}^*) = 1$ for all $v_{t-1}^* \in [0, 1)$ in the bad-news case.[11]

4.2 Main result

The principal's problem is to determine the distribution of the sequence of beliefs to maximize the expected total surplus subject to the three constraints noted above. Define $W_t(\mu_t, x, y)$ as the principal's continuation payoff (or simply the continuation surplus) in period t when the belief is μ_t and agents $v \in [x, y)$ adopt in period t. Given some mechanism P, the continuation surplus is given by

$$\begin{aligned} W_t(\mu_t, v_t^*, v_{t-1}^*) = &\int_{v_t^*}^{v_{t-1}^*} (v - \mu_t c) \mathrm{d}F(v) \\ &+ \sum_{s=1}^{T-t} \sum_{\mu^{t,t+s} \in H_t^s} \delta^s P_{t+1}(\mu_{t+1}|\mu^t) \cdots P_{t+s}(\mu_{t+s}|\mu^{t+s-1}) W_{t+s}(\mu_{t+s}, v_{t+s}^*, v_t^*). \end{aligned}$$

The principal's problem is then defined as

$$\max_P \; W_1(m, v_1^*, 1),$$

subject to the Bayes plausibility (3) and consistency (4), where $(v_t^*)_{t=1}^T$ is a non-increasing sequence determined by the IC constraint (2) whenever $v_t^* < v_{t-1}^*$. Recall that since the principal has no additional information over and above the prior in period 1, the Bayes plausibility

[11] We use this slightly more general notation to make it applicable to the good-news case we consider later.

requires $P_1(m) = 1$. Our focus is therefore on period 2 and beyond.

A few more notations would be helpful in formally stating the main result. First, for $t \geq 2$, we divide the support $M_t(\mu^{t-1})$ into two disjoint sets $M_t^A(\mu^{t-1})$ and $M_t^N(\mu^{t-1})$, where

$$M_t^A(\mu^{t-1}) := \{\mu_t \in M_t(\mu^{t-1}) : v_t^* < v_{t-1}^*\},$$
$$M_t^N(\mu^{t-1}) := \{\mu_t \in M_t(\mu^{t-1}) : v_t^* = v_{t-1}^*\},$$

with $M_1(\mu^0) = M_1^A(\mu^0) = \{m\}$. We say that the game is in the *adoption phase* in period t if $\mu_t \in M_t^A(\mu^{t-1})$ and in the *no-adoption phase* in period t if $\mu_t \in M_t^N(\mu^{t-1})$. Also, for the subsequent analysis, let

$$q(v_t^*, v_{t-1}^*) := \underline{\mu}(v_{t-1}^*)e^{-\lambda(F(v_{t-1}^*)-F(v_t^*))} + 1 - \underline{\mu}(v_{t-1}^*)$$

denote the probability of receiving no news in period t conditional on being optimistic. Given this, for any $\beta \in [0, 1]$, let $\tilde{\mu}(v_{t-1}^*, v_{t-2}^*; \beta)$ denote the belief that satisfies the Bayes plausibility in period t when the mechanism assigns $\mu_t = \underline{\mu}(v_{t-1}^*)$ with probability $\beta q(v_{t-1}^*, v_{t-2}^*)$ and $\mu_t = \tilde{\mu}(v_{t-1}^*, v_{t-2}^*; \beta)$ with the remaining probability conditional on $\mu_{t-1} = \underline{\mu}(v_{t-2}^*)$. Formally, it is given by the solution to the following equation:

$$\beta q(v_{t-1}^*, v_{t-2}^*)\underline{\mu}(v_{t-1}^*) + (1 - \beta q(v_{t-1}^*, v_{t-2}^*))\tilde{\mu}(v_{t-1}^*, v_{t-2}^*; \beta) = \underline{\mu}(v_{t-2}^*).$$

Definition 1 *A mechanism P is a binary-message termination mechanism (or, simply, a binary-message mechanism) if for each $t = 2, \ldots, T$,*

(a) for any $\mu_{t-1} \in M_{t-1}^A(\mu^{t-2})$, there is $\beta_t \in [0, 1]$ such that

$$P_t(\mu_t \mid \mu^{t-1}) = \begin{cases} \beta_t q(v_{t-1}^*, v_{t-2}^*) & \text{if } \mu_t = \underline{\mu}(v_{t-1}^*), \\ 1 - \beta_t q(v_{t-1}^*, v_{t-2}^*) & \text{if } \mu_t = \tilde{\mu}(v_{t-1}^*, v_{t-2}^*; \beta_t), \\ 0 & \text{otherwise,} \end{cases}$$

(b) for any $\mu_{t-1} \in M_{t-1}^N(\mu^{t-2})$, the mechanism assigns all the probability on μ_{t-1}, i.e.,

$$P_t(\mu_t \mid \mu^{t-1}) = \begin{cases} 1 & \text{if } \mu_t = \mu_{t-1}, \\ 0 & \text{otherwise,} \end{cases}$$

where $(v_t^*)_{t=1}^T$ is the induced sequence of thresholds when the game stays in the adoption phase *for all t.*[12]

The following facts are useful for interpreting binary-message mechanisms. Throughout the paper, all the proofs are relegated to Appendix A.

Lemma 1 *In a binary-message termination mechanism,*

(i) $\underline{\mu}(v^*_{t-1}) \in M^A_t(\mu^{t-1})$ *and* $\tilde{\mu}(v^*_{t-1}, v^*_{t-2}; \beta_t) \in M^N_t(\mu^{t-1})$ *if* $\mu_{t-1} = \underline{\mu}(v^*_{t-2}) \in M^A_t(\mu^{t-2})$*,*

(ii) $\mu_t \in M^N_t(\mu^{t-1})$ *if* $\mu_t = \mu_{t-1} \in M^N_{t-1}(\mu^{t-2})$*,*

where $(v^*_t)^T_{t=1}$ *is the induced sequence of thresholds when the game stays in the adoption phase for all t.*

Along with Lemma 1, Definition 1 provides two essential properties, which we call *termination* and *caution* for clarity. First, part (b) of the definition states the termination property where if the game is in the no-adoption phase for a period, it will never revert back to the adoption phase; as such, $\mu_t \in M^N_t(\mu^{t-1})$ indicates permanent termination of the adoption process. Second, part (a) requires the principal to provide the most accurate information $\mu_t = \underline{\mu}(v^*_{t-1})$ with some probability in period t if the game is in the adoption phase in period $t-1$. This alternatively implies that the principal must be optimistic whenever the game is in the adoption phase.[13] Given this, when the game is in the adoption phase in period $t-1$ (implying that the principal is optimistic in period $t-1$), she remains optimistic (by observing no news) with probability $q(v^*_{t-1}, v^*_{t-2})$, in which case she randomizes between continuation (with probability β_t) and termination (with probability $1-\beta_t$). With the remaining probability, the principal turns pessimistic (by observing bad news), in which case she surely terminates the adoption process. We refer to this as the caution property, because such a mechanism is biased towards termination in that it terminates surely in the pessimistic state but may also terminate with some probability in the optimistic state (where the principal has received no bad news).

Any binary-message mechanism is characterized solely by a sequence $(\beta_t)^T_{t=1}$, where each β_t admits an economically meaningful interpretation: $\beta_t = 1$ corresponds to full disclosure while $\beta_t = 0$ corresponds to no disclosure. The following statement is the main characterization result of this paper.

Theorem 1 *There exists an optimal mechanism that is a binary-message termination mechanism.*

[12]Under Definition 1, there is a unique sequence of thresholds as long as the game stays in the adoption phase.

[13]To realize the lowest possible belief, the principal must send a fully separating message in the optimistic state. The conclusion then follows since $\mu_t \in M^A_t(\mu^{t-1})$ if and only if $\mu_t = \underline{\mu}(v^*_{t-1})$ by Lemma 1 .

4.3 Sketch of the proof

The proof of Theorem 1 is lengthy, and we relegate its technical details to Appendix A. Below, we provide a sketch of the proof to summarize the key steps towards establishing our characterization result and to build intuition for it.

The proof proceeds in two steps. The first step is to show that there is no gain from adopting a more complicated "stop-and-go" mechanism that switches back and forth between the adoption phase and the no-adoption phase. This can be seen most clearly by inspecting a three-period example. Consider a mechanism P that induces the lowest possible belief, which we denote by $\underline{\mu}$, with probability $\beta_2 q(v_1^*, 1)$ in period 2 and the belief $\tilde{\mu}(v_1^*, 1; \beta_2)$ leading to no adoption with the remaining probability. Suppose further that in the history with $v_1^* = v_2^*$ (no adoption in period 2), P induces the lowest possible belief $\mu_3 = \underline{\mu}$ with probability $\gamma_3 q(v_1^*, 1)$ in period 3. The total surplus attained by this mechanism is then given by

$$W_1(m, v_1^*, 1) = \int_{v_1^*}^{1} (v - mc)\mathrm{d}F(v) + \beta_2 q(v_1^*, 1)\delta W_2(\underline{\mu}, v_2^*, v_1^*)$$
$$+ (1 - \beta_2 q(v_1^*, 1))\gamma_3 q(v_1^*, 1)\delta^2 W_3(\underline{\mu}, v_3^*, v_1^*),$$

where the IC constraint in period 1 is

$$v_1^* - mc = \beta_2 q(v_1^*, 1)\delta(v_1^* - \underline{\mu}c) + (1 - \beta_2 q(v_1^*, 1))\gamma_3 q(v_1^*, 1)\delta^2(v_1^* - \underline{\mu}c).$$

On one hand, the IC constraint is unaffected as long as $\beta_2 q(v_1^*, 1) + (1 - \beta_2 q(v_1^*, 1))\gamma_3 q(v_1^*, 1)\delta$ is fixed. On the other hand, we have $W_2(\underline{\mu}, v_2^*, v_1^*) \geq W_3(\underline{\mu}, v_3^*, v_1^*)$ because in period 2, the principal could always replicate what she would do in period 3 (by providing no information in period 3). This means that we can improve the original mechanism by reducing γ_3 and increasing β_2 in a way to keep $\beta_2 q(v_1^*, 1) + (1 - \beta_2 q(v_1^*, 1))\gamma_3 q(v_1^*, 1)\delta$ constant. In the end, this implies $\gamma_3 = 0$, i.e., it is optimal to front-load all the information. This argument directly implies the termination property that substantially reduces the class of mechanisms we need to consider.

Given this, the next step is to show that we can always improve the total surplus by applying a mean-preserving spread to any interior belief. Intuitively, such an operation is beneficial for the agents because their value functions are convex; it is in fact relatively straightforward to establish this when there are only two periods. However, when there are more than two periods, the situation becomes more complicated because the game continues after the period in which we apply a mean-preserving spread. To illustrate this point, suppose we start from some interior belief $\mu_t \in (\underline{\mu}(v^*_{t-1}), \frac{v^*_{t-1}}{c})$ at which a positive measure of agents adopt in period $t < T$. To improve the payoffs of those agents by providing more information, their behavior must be contingent on the new information. This alternatively implies that we must split the belief μ_t into two outer beliefs such that they choose to adopt at the lower belief but not to adopt at the higher belief. While this operation clearly benefits those agents who are induced to adopt in period t, it necessarily lowers the continuation probability (more frequent termination), which could adversely affect the remaining agents who may adopt in some future periods.

Despite this complication, we can still show that the benefit of more precise information always dominates the cost of more frequent termination. Specifically, we establish this result by considering a mean-preserving spread that splits an interior belief $\mu_t \in (\underline{\mu}(v^*_{t-1}), \frac{v^*_{t-1}}{c})$ for $t \geq 2$ into a slightly lower belief $\mu_t - \varepsilon$ and the highest possible one $\overline{\mu}(v^*_{t-1})$ (which equals 1 for any $v^*_{t-1} < 1$) and showing that it is Pareto-improving.[14] This observation directly implies that the principal must provide the most accurate information in the adoption phase by terminating surely in the pessimistic state while randomizing in the optimistic state. The reason why this caution property holds is the same as in the two-period example: more precise information at the low end of the belief distribution helps agents with lower valuations make more informed decisions but is irrelevant for those with higher valuations. The principal can exploit this fact to alleviate the free-rider problem and achieve the optimal outcome.

In any period t, the principal is either pessimistic with $\mu^p_t = 1$ or optimistic with $\mu^p_t = \underline{\mu}(v^*_{t-1})$. The caution property suggests that whenever the game is in the adoption phase, it must induce the lowest possible belief, i.e., $M^A_t(\mu^{t-1}) = \{\underline{\mu}(v^*_{t-1})\}$, which implies that the principal surely terminates in the pessimistic state. When the game switches to the no-adoption phase, the mechanism can specify any beliefs resulting in no adoption, but the termination property suggests that we can merge all these beliefs into a single one without loss of generality, i.e., $M^N_t(\mu^{t-1}) = \{\tilde{\mu}(v^*_{t-1}, v^*_{t-2}; \beta_t)\}$. The resulting mechanism uses only two messages, one indicating the continuation of the adoption process and the other indicating its permanent termination, which amounts to a simple binary-message termination mechanism.

[14]In the proof, we exploit the fact that the highest possible belief is 1 (for any $v^*_t < 1$) in the bad-news case, at which point there is no potential loss of surplus (because the state is bad for sure). As it turns out, though, this property also holds in the good-news case where the highest possible belief is strictly less than 1. We expand on this point in Section 6.2.

4.4 When information design matters

Given Theorem 1, the principal's problem can now be substantially simplified. Note that the principal's objective function can be written as

$$W_1(m, v_1^*, 1) = \sum_{t=1}^{T} \delta^{t-1} B_t \left(m e^{-\lambda(1-F(v_{t-1}^*))} \int_{v_t^*}^{v_{t-1}^*} (v-c)\mathrm{d}F(v) + (1-m) \int_{v_t^*}^{v_{t-1}^*} v\mathrm{d}F(v) \right),$$

where $B_t := \Pi_{s=1}^{t} \beta_s$. Let $r(v_t^*, v_{t-1}^*)$ denote the joint probability of state being bad and receiving no news in period t conditional on being optimistic. Formally,

$$r(v_t^*, v_{t-1}^*) := \underline{\mu}(v_{t-1}^*) e^{-\lambda(F(v_{t-1}^*)-F(v_t^*))}.$$

Then, the IC constraint in each period t is simplified to

$$v_t^* - \underline{\mu}(v_{t-1}^*)c = \delta\beta_{t+1}(q(v_t^*, v_{t-1}^*)v_t^* - r(v_t^*, v_{t-1}^*)c). \tag{5}$$

From these, the principal's problem is redefined as

$$\max_{(\beta_t)_{t=1}^T} W_1(m, v_1^*, 1),$$

subject to (5).

Observe that the principal has no private information and has no choice but to set $P_1(m) = 1$, which implies $\beta_1 = 1$. For $t \geq 2$, a change in β_t affects v_{t-1}^* and also $v_t^*, \ldots, v_T^*$ through its effect on v_{t-1}^*. Taking derivative with respect to β_t yields

$$\begin{aligned} \delta^{1-t}\frac{\partial W_1}{\partial \beta_t} =& \sum_{s=t}^{T} \delta^{s-t} B_s \lambda r(v_{s-1}^*, 1) f(v_{s-1}^*) \frac{\partial v_{s-1}^*}{\partial \beta_t} \int_{v_s^*}^{v_{s-1}^*} (v-c)\mathrm{d}F(v) \\ &+ \sum_{s=t}^{T} \delta^{s-t} \frac{B_{s+1}}{\beta_t} q(v_{s-1}^*, 1) \int_{v_s^*}^{v_{s-1}^*} (v - \underline{\mu}(v_{s-1}^*)c)\mathrm{d}F(v). \end{aligned} \tag{6}$$

Notice that the second term in this expression is always positive. Thus, if there in an interior solution for the optimal β_t, then the first term must be negative. This illustrates the key trade-off faced by the principal: lower transparency, i.e. a lower β_t, makes the agents more willing to experiment now (as captured by the first term) but leads to more distortion and more premature termination (as captured by the second term). It is easy to see that the first term of (6) disappears as $\beta_t \to 0$ while the second term is unaffected. Therefore, it is never optimal to set $\beta_t = 0$ (no disclosure) for any t.

By contrast, there may be cases where it is optimal to fully disclose all the information by setting $\beta_t = 1$, despite the fact that the amount of adoption is generally insufficient due to the free-rider problem. This is because the principal must respect the IC constraint and, in some cases, cannot do any better than fully disclosing all the information; formally, this occurs when we have a corner solution for some period (i.e., $\beta_t = 1$ for some $t \geq 2$), in which case information design is of no use for that period. To illustrate this possibility, consider the problem of choosing β_T. Given that $v_T^* = \underline{\mu}(v_{T-1}^*)c$, the derivative with respect to β_T is obtained as

$$\delta^{1-T}\frac{\partial W_1}{\partial \beta_T} = \beta_T \lambda r(v_{T-1}^*, 1) f(v_{T-1}^*) \frac{\partial v_{T-1}^*}{\partial \beta_T} \int_{\underline{\mu}(v_{T-1}^*)c}^{v_{T-1}^*} (v - c)\mathrm{d}F(v) \\ + q(v_{T-1}^*, 1) \int_{\underline{\mu}(v_{T-1}^*)c}^{v_{T-1}^*} (v - \underline{\mu}(v_{T-1}^*)c)\mathrm{d}F(v) \tag{7}$$

where, by (5),

$$\frac{\partial v_{T-1}^*}{\partial \beta_T} = \frac{\delta(q(v_{T-1}^*, v_{T-2}^*)v_{T-1}^* - r(v_{T-1}^*, v_{T-2}^*)c)}{1 - \delta\beta_T[\lambda r(v_{T-1}^*, v_{T-2}^*)(v_{T-1}^* - c) + q(v_{T-1}^*, v_{T-2}^*)]}.$$

Note that $\frac{\partial v_{T-1}^*}{\partial \beta_T} \to 0$ as $\delta \to 0$, and hence the first term of the left-hand side of (7) disappears. Since the second term is always strictly positive, we end up with a corner solution when δ is sufficiently small. Intuitively, when the agents are almost myopic, there is no information free-riding problem, and therefore, there is no gain from lowering transparency.

As we have shown, no disclosure is never optimal but full disclosure can sometimes be optimal. Against this backdrop, an economic question of practical importance is when full disclosure is not optimal. As one might expect from our discussions so far, full disclosure is not optimal when the agents are more forward-looking because they wait more patiently for news, thereby aggravating the free-rider problem. To see this, suppose the principal sets $\beta_{t+1} = 1$ for some t and let $\delta \to 1$. The period-t IC constraint then converges to

$$v_t^* - \underline{\mu}(v_{t-1}^*)c = q(v_t^*, v_{t-1}^*)v_t^* - r(v_t^*, v_{t-1}^*)c. \tag{8}$$

Provided that $c > v_t^*$, the only solution that can satisfy this is $v_t^* = v_{t-1}^*$,[15] i.e., no agents

[15] To see, note that $r(v_t^*, v_{t-1}^*) = q(v_t^*, v_{t-1}^*) + \mu(v_{t-1}^*) - 1$. Substituting for $r(v_t^*, v_{t-1}^*)$ and rearranging, the only solution that can satisfy this equation when $c > v_t^*$ must have $q(v_t^*, v_{t-1}^*) = 1$. This is only possible when $v_t^* = v_{t-1}^*$.

adopt in period t, resulting in an extreme form of procrastination that produces no useful information along the way. We can then show that setting $\beta_{t+1} = 1$ is not optimal for the same reason employed in establishing the termination property. This observation suggests that there is always room for information design to improve welfare when the agents are sufficiently forward-looking.

Theorem 2 *In the optimal mechanism, for* $t = 2, \ldots, T$,

(i) no disclosure is never optimal, i.e, $\beta_t > 0$,

(ii) partial disclosure is optimal when the agents are sufficiently patient, i.e., $\beta_t < 1$, *if* δ *is sufficiently close to 1.*

4.5 Asymptotic outcome and the limit of social learning

We have thus far observed that the optimal disclosure mechanism exercises caution and occasionally terminates even when the principal has observed no news and is relatively optimistic as a result. One way to illustrate the trade-off generated by the caution property is to look at the asymptotic outcome as T tends to infinity. To this end, we start with the full-disclosure policy. It is easy to observe that social learning never ceases under full disclosure as long as no bad news is observed.[16] When the state is bad, the game continues indefinitely until bad news arrives, and the true state is eventually identified. When the state is good, the principal's belief when she is optimistic converges to 0, albeit very slowly, and so does the threshold type. This means that the true state can be identified almost surely under full disclosure, suggesting that there is no inherent upper limit of social learning in the bad-news case.

The fact that the true state can be identified asymptotically under full disclosure clarifies a distortion introduced by the optimal mechanism. For the sake of argument, suppose δ is close to 1, so that $\beta_t < 1$ for some t. By the same argument as in the previous paragraph and also the termination property, social learning never ceases in the optimal mechanism as long as the adoption phase continues. Therefore, when the state is bad, the adoption process terminates almost surely as under full disclosure, though at a faster rate. However, when the state is good, the adoption process may terminate prematurely, so that the true state may not be fully identified even asymptotically; this is an obvious distortion associated with the caution property, giving rise to the essential trade-off between the cost of premature termination

[16] A strictly positive measure of agents adopt in period t if $v^*_{t-1} - \underline{\mu}(v^*_{t-1})c > \delta(q(v^*_t, v^*_{t-1})v^*_{t-1} - r(v^*_t, v^*_{t-1})c)$. As $v^*_{t-1} \to v^*_t$, this reduces to $v^*_{t-1} > \underline{\mu}(v^*_{t-1})c$ if $\delta \in [0, 1)$, which holds for any v^*_{t-1} that satisfies the IC constraint.

and the gain of faster social learning. As we can see from Theorem 2, the cost arising from this distortion is always outweighed by faster social learning when the agents are sufficiently forward-looking.

5 Policy implications

Our analysis suggests that it is *ex ante* optimal to terminate the adoption process with some probability, even when the principal has observed no news. There are, however, some conceptual issues in implementing this outcome in practice. First, it may be deemed inappropriate, or even unethical, for the government to knowingly manipulate information; this is especially so in health-related issues such as vaccine uptake. Second, the implementation of this outcome requires credible randomization, but it is certainly inappropriate for the government to flip a coin to determine when to terminate.

In the context of collective experimentation, it is more reasonable to interpret our model environment as the one in which the principal designs and commits to the way she collects and analyzes data to produce evidence and transforms it into a binary recommendation at the outset of the game. More precisely, in our model, we assume that the principal observes conclusive bad news with some probability in each period, but in reality, such conclusive news is often prohibitively costly or simply infeasible to attain; in most cases, the principal must evaluate various pieces of evidence with varying degrees of precision to form a belief. The information-design problem of our analysis can be regarded as a reduced form of this complicated deliberation process. If it is costly to acquire more accurate information, it may be inevitable to admit some statistical errors in the decision-making process. The relevant questions to ask are then what type of statistical error—type I (false positive) or type II (false negative)—should be condoned and to what extent.

Formally, consider a decision-making process in which the principal designs the information structure that determines how stochastic signals are generated and a decision rule that maps the observed signal into a binary recommendation. Given that decision rules that involve randomization are hard to implement in practice, we here only consider pure (deterministic) decision rules. Below, we build on this framework to explore ways in which to achieve the (second-best) optimal outcome and highlight two practical implications of our analysis.[17]

[17] In the following argument, we abstract away from the cost of making the information structure more accurate to illustrate the key insights. If it is costly to raise the precision of the information structure, the principal must trade off the benefit of information design against the cost, but even then, the same insights should follow.

False positives versus false negatives. To illustrate what type of statistical error should be more tolerated, consider the two-period example with an augmented information structure. Suppose the signal distribution is now given by

$$\mathbb{P}[s_1 = b \mid \omega, v_1^*] = \varepsilon_I e^{-\lambda\omega(1-F(v_1^*))} + (1-\varepsilon_{II})(1-e^{-\lambda\omega(1-F(v_1^*))}).$$

The efficiency of the information structure is measured by $(\varepsilon_I, \varepsilon_{II})$, where ε_I is the probability of a type I error and ε_{II} is interpreted loosely as representing the probability of a type II error. The information structure in the baseline model corresponds to the one with $\varepsilon_I = \varepsilon_{II} = 0$. Under this specification, the probability of observing no signal, denoted by $\hat{q}(v_1^*)$, is

$$\begin{aligned}\hat{q}(v_1^*) &= 1 - m[\varepsilon_I e^{-\lambda(1-F(v_1^*))} + (1-\varepsilon_{II})(1-e^{-\lambda(1-F(v_1^*))})] - (1-m)\varepsilon_I \\ &= (1-\varepsilon_I)q(v_1^*, 1) + \varepsilon_{II} m(1-e^{-\lambda(1-F(v_1^*))}).\end{aligned}$$

Similarly, define $\hat{r}(v_1^*)$ such that

$$\begin{aligned}\hat{r}(v_1^*) &= m[1 - \varepsilon_I e^{-\lambda(1-F(v_1^*))} - (1-\varepsilon_{II})(1-e^{-\lambda(1-F(v_1^*))})] \\ &= (1-\varepsilon_I)r(v_1^*, 1) + \varepsilon_{II} m(1-e^{-\lambda(1-F(v_1^*))}).\end{aligned}$$

Since there are only two signals, the only non-trivial (pure) decision rule is to continue after $s_1 = \varnothing$ and terminate after $s_1 = b$. Given this decision rule, the principal designs the information structure by choosing $(\varepsilon_I, \varepsilon_{II})$ at the outset of the game. Observe that the IC constraint becomes

$$v_1^* - mc = \delta(\hat{q}(v_1^*)v_1^* - \hat{r}(v_1^*)c).$$

for a given $(\varepsilon_I, \varepsilon_{II})$. We argue that the principal can implement any outcome that is achievable by a binary-message mechanism. To this end, let $v^*(\beta)$ be the threshold (in the two-period model) when $\beta_2 = \beta$, i.e., the threshold that satisfies

$$v^*(\beta) - mc = \delta\beta(q(v^*(\beta), 1)v^*(\beta) - r(v^*(\beta), 1)c),$$

The principal's goal is to implement $v_1^* = v^*(\beta)$. To achieve this outcome, we must have $\hat{q}(v^*(\beta)) = \beta q(v^*(\beta), 1)$ and $\frac{\hat{r}(v^*(\beta))}{\hat{q}(v^*(\beta))} = \underline{\mu}(v^*(\beta)) = \frac{r(v^*(\beta),1)}{q(v^*(\beta),1)}$, which together imply $\hat{r}(v^*(\beta)) = \beta\hat{r}(v^*(\beta))$. Note that these conditions can be satisfied if and only if $\varepsilon_I = 1-\beta$ and $\varepsilon_{II} = 0$. This argument suggests that false positives (termination in the good state) are more acceptable than false negatives (continuation in the bad state) in that the optimal information structure must occasionally admit false positives but no false negatives.

Virtue of overly cautious mechanisms. We extend the same idea further by introducing a second signal that is not fully informative. Suppose $s_1 \in \{b, w, \varnothing\}$, and consider the following signal distribution:

$$\mathbb{P}[s_1 \mid \omega, v^*] = \begin{cases} (1-\tilde{\varepsilon})(1 - e^{-\lambda\omega(1-F(v^*))}) & \text{if } s_1 = b, \\ \varepsilon_I e^{-\lambda\omega(1-F(v^*))} + \tilde{\varepsilon}(1 - e^{-\lambda\omega(1-F(v^*))}) & \text{if } s_1 = w, \end{cases}$$

where we suppose for illustration that $\tilde{\varepsilon}$ is fixed due to some exogenous factors. In this specification, the strong signal $s_1 = b$ is still conclusive, corresponding, e.g., to a catastrophic event that unambiguously reveals that the state is bad and leads to immediate termination. By contrast, the weak signal $s_1 = w$ represents the ambiguous state that requires the principal's judgment.[18] Let $\mu_w(v^*)$ be the principal's belief associated with the weak signal, which is given by

$$\mu_w(v^*) = \frac{m[\varepsilon_I e^{-\lambda(1-F(v^*))} + \tilde{\varepsilon}(1 - e^{-\lambda(1-F(v^*))})]}{m[\varepsilon_I e^{-\lambda(1-F(v^*))} + \tilde{\varepsilon}(1 - e^{-\lambda(1-F(v^*))})] + (1-m)\varepsilon_I}.$$

In this setting, the principal can achieve the optimal outcome by setting ε_I at $1-\beta$ (while $\tilde{\varepsilon} \in (0,1)$ can be arbitrary) and terminating after $s_1 \in \{b, w\}$.

Note that $\mu_w(v^*)$ may take any value between $\underline{\mu}(v^*)$ and 1 and is strictly decreasing in $\tilde{\varepsilon}$. As $\tilde{\varepsilon}$ gets smaller, therefore, the weak signal becomes almost equivalent to no news. Under the optimal decision rule, however, the principal is still instructed to terminate the adoption process upon observing $s_1 = w$. The decision rule considered above is thus overly cautious (when $\tilde{\varepsilon}$ is relatively small) in the sense that the principal "overreacts" to insignificant events that are hardly informative. Alternatively, this argument points to a possibility that when it is not feasible to commit to a decision rule *ex ante*, delegating decision-making authority to a cautious leader can work as a compromised solution; this can be done by appointing someone who is known to be cautious or by institutional design to punish failures severely and somewhat excessively. In either case, our analysis suggests a virtue of mechanisms that may appear overly cautious, even though we assume that all the concerned parties are risk-neutral.

[18] As before ε_I corresponds to type I error when the decision rule is to terminate after any signal.

6 Extensions

In the main body of the analysis, we deliberately consider the simplest setting to illustrate the key insights in a transparent manner. As it turns out, though, our main characterization result, Theorem 1, is robust to various alterations to the underlying setup and hence holds more generally. Below, we discuss two possible extensions—one on the principal's objectives and the other on the information-generating process—to ensure that our framework can be applied widely to a range of situations of practical importance.

6.1 More general objective function

So far, we have considered a benevolent principal whose goal is to maximize the sum of the payoffs of all agents. In practice, however, the principal's objective can be more diverse and include benefits that are not fully internalized by the agents. For instance, in the case of vaccine uptake, one major social goal is to achieve herd immunity, the benefit of which can be shared equally among all citizens irrespective of their vaccination decisions. The vaccination process also involves third parties, such as pharmaceutical companies, and the government may have some interest in their well-being as well.

To capture this situation, we now suppose that the principal's objective function includes an additional term that depends on past thresholds. Specifically, suppose that the principal's payoff in period t is now more generally given by

$$\int_{v_t^*}^{v_{t-1}^*} (v - \mu_t c)\mathrm{d}F(v) + \alpha H(v_{t-1}^*, v_t^*),$$

where $H : [0,1]^2 \to \mathbb{R}_+$. In this specification, $H(v_{t-1}^*, v_t^*)$ measures the external benefit that is contingent on the current and previous thresholds, and $\alpha \geq 0$ is the welfare weight given to the external benefit with our baseline model corresponding to a special case with $\alpha = 0$. Note that $1 - v_t^*$ represents the stock of agents who have adopted up to period t while $v_{t-1}^* - v_t^*$ represents the flow of agents who adopt in period t. This extended specification of the objective function can thus accommodate a range of goals that depend on the stock and the current flow in a flexible manner: for instance, the level of herd immunity should be positively correlated with the stock of citizens who have been vaccinated, whereas a pharmaceutical firm's profit in a given period should be positively correlated with the flow of citizens who get vaccinated in that period. In what follows, we make the following assumption on the form of the external benefit.

Assumption 2 *The following conditions are satisfied:*

(a) $H_2(v^*_{t-1}, v^*_t) \leq 0$ *for any* (v^*_{t-1}, v^*_t)*;*

(b) $H_2(v^*_{t-1}, v^*_t) + H_1(v^*_t, v^*_{t+1}) \leq 0$ *for any* $(v^*_{t-1}, v^*_t, v^*_{t+1})$*.*

In any period t, an increase in the current threshold v^*_t implies a decrease in the stock and the flow. It is therefore reasonable to assume $H_2 \leq 0$. By contrast, since an increase in the previous threshold v^*_{t-1} implies an increase in the flow, it is *a priori* difficult to determine the sign of H_1. Note that the previous threshold itself was a current threshold in the previous period. Taking this into account, in part (b) of Assumption 2, we assume that the overall impact of an increase in v^*_t is weakly negative. An implication of this assumption is that for any two sequences $(v^*_s, \ldots, v^*_{s'})$ and $(\hat{v}^*_s, \ldots, \hat{v}^*_{s'})$ such that $v^*_s = \hat{v}^*_s$, $v^*_{s'} = \hat{v}^*_{s'}$, and $v^*_t \geq \hat{v}^*_t$ for all $t = s+1 \ldots s'-1$,

$$\sum_{t=s}^{s'-1} \delta^{t-s} H(\hat{v}^*_t, \hat{v}^*_{t+1}) \geq \sum_{t=s}^{s'-1} \delta^{t-s} H(v^*_t, v^*_{t+1}),$$

i.e., if we have two sequences that start from and end at the same thresholds, the one that is consistently lower (i.e. the one that has faster adoption rate) yields a weakly higher surplus.[19] For instance, Assumption 2 is satisfied if the external benefit is given by

$$H(v^*_{t-1}, v^*_t) = \eta(v^*_{t-1} - v^*_t) + h(v^*_t),$$

where $\eta > 0$ is a weight given to the flow and $h : [0,1] \to \mathbb{R}_+$ is some decreasing function that captures the impact of the stock. We can show that our main characterization result holds in this extended setup.

Theorem 3 *Theorem 1 holds for any* $\alpha > 0$ *under Assumption 2.*

6.2 Good-news case

In the baseline model, we focus on the case where the principal may observe bad news (a "breakdown") when the state is bad. Here, we consider the opposite case where the principal

[19]To see this, if there is an infinitesimal increase in v^*_t (while fixing all other thresholds constant), the change in the surplus is $H_2(v^*_{t-1}, v^*_t) + \delta H_1(v^*_t, v^*_{t+1})$. Since the first term is positive by assumption, the change must be nonnegative for any $\delta \in [0, 1]$.

may observe good news (a "breakthrough") when the state is good. Specifically, let $s_t \in \{\varnothing, g\}$, and assume

$$\mathbb{P}[s_t = g \mid \omega, n_t] = 1 - e^{-\lambda(1-\omega)n_t}$$

where n_t is the measure of agents who adopt in period t. The principal's belief jumps down to 0 once she observes good news so that $\underline{\mu}(v^*_{t-1}) = 0$ for any $v^*_{t-1} \in [0, 1]$. In the absence of good news, the principal's belief is given by

$$\mu^p_t = \overline{\mu}(v^*_{t-1}) := \frac{m}{m + (1-m)e^{-\lambda(1-F(v^*_t))}},$$

which is the highest possible belief that can be induced in period t. The consistency restriction is then defined in the same way, except that μ_t must now be bounded between 0 and $\overline{\mu}(v^*_{t-1})$.

Note that in the good-news case, the optimistic state is when the principal has observed good news while the pessimistic state is when she has observed no news. Given this, the caution property is still the same, where a binary-message mechanism terminates surely in the pessimistic state, whereas it randomizes and continues with some probability in the optimistic state. With abuse of notation, redefine

$$q(v^*_t, v^*_{t-1}) := (1 - \underline{\mu}(v^*_{t-1}))(1 - e^{-\lambda(1-F(v^*_t))})$$

as the probability of observing good news in period t, where

$$\underline{\mu}(v^*_{t-1}) = \begin{cases} m & \text{if } v^*_{t-1} = 1, \\ 0 & \text{if } v^*_{t-1} < 1. \end{cases}$$

With this modified definition, in a binary-message mechanism, the probability that the game stays in the adoption phase is still $\beta_t q(v^*_{t-1}, v^*_{t-2})$ as in the bad-news case.[20]

Before we state the result, it is worth noting that the good-news case is not the mirror image of the bad-news case. In the proof of Theorem 1, we consider a mean-preserving spread that splits an interior belief into a slightly lower belief and the highest possible one. As noted in Section 4.3, while this mean-preserving spread provides more accurate information and benefits the agents who adopt now, it could be detrimental to the remaining agents who might adopt in the future because it necessarily lowers the continuation probability. In the bad-news case, however, the highest possible belief is always 1, in which case the state is bad for sure.

[20] As we will see below, the game always ends in two periods, so that β_t for $t \geq 3$ is actually irrelevant.

Therefore, the decrease in the continuation probability, induced by a mean-preserving spread, entails no loss of surplus, and we can take a more direct route to show that such an operation benefits all agents. This same argument cannot be applied to the good-news case because the highest possible belief is always bounded away from 1. In this situation, terminating in the pessimistic state (after no news) entails a loss of potential surplus and yields the aforementioned trade-off, which forces us to explore a different route to establish the characterization result. As it turns out, though, the gain from alleviating the free-rider incentive always dominates the loss of surplus due to the decrease in the continuation probability, such that the same principles—termination and caution—still hold in this alternative scenario.

Theorem 4 *Theorem 1 holds in the good-news case.*

The theorem states that the caution property continues to hold in the good-news case. As a consequence, the period-2 belief is either $\mu_2 = 0$ (the optimistic state) or $\mu_2 = \tilde{\mu}(v_1^*, 1; \beta_2) > \frac{v_1^*}{c}$ (the pessimistic state). If $\mu_2 = 0$, the state is good for sure, and all remaining agents adopt in period 2. If not, no agent will adopt after period 2 by the termination property. In either case, therefore, social learning ceases in two periods. This feature of our model draws a clear contrast with Che and Hörner (2018) who also look at the good-news case in their baseline specification. In their model, the optimal mechanism always recommends adoption after good news and also with some probability after no news. The strategy of mixing good news with no news can extend the learning process even after no news and is valuable in their setting. In our framework where agents are forward-looking, however, the same strategy would not work because that would intensify the incentive to wait for news. Although terminating the adoption process after no news implies a loss of surplus because the principal's belief is still bounded away from 1, it is nonetheless part of the optimal mechanism to deter information free-riding, and the principal cannot fare any better than this.

The fact that social learning cannot extend beyond the second period also suggests that there is a natural upper bound of social learning in the good-news case. In the best case, the true state can be discovered with a breakthrough, but this occurs only with probability $(1-m)(1-e^{-\lambda(1-F(v_1^*))})$, which is strictly bounded away from 1. In case no news arrives, the principal's belief goes down slightly, but she must still terminate the adoption process altogether, and there is no feasible way to incentivize exploration beyond this point, no matter how large T becomes. This asymptotic property of the good-news case is clearly different from that of the bad-news case where there is no such limit, and increasingly more information can be acquired as T gets larger. The crucial difference is that news that leads to termination (i.e., bad news) takes a form of conclusive bad news in the bad-news case, so that the true state is precisely revealed when the adoption process terminates under full disclosure. In the good-news case, by contrast, the adoption process terminates after inconclusive news (i.e., no news), making it inherently infeasible to fully identify the true state even asymptotically.

Finally, it is also worth noting that although there is an upper limit of social learning, the free-rider problem tends to be less severe in the good-news case because of the way information is generated in this case. In the bad-news case, the arrival of news indicates that the state is bad for sure. Since no one has an incentive to adopt when the state is bad, this information is valuable for all agents regardless of their valuations. This is not necessarily the case in the good-news case where the principal's belief is bounded away from 1 in the pessimistic state (no news). In this case, additional information may not be informative enough for agents with sufficiently high valuations for whom there is no free-rider incentive. Although we do not analyze this aspect in depth as it is outside the scope of this paper, it is of some interest, at least theoretically, to explore the difference between the bad-news and good-news cases from this perspective.

7 Conclusion

Social learning is often hampered by citizens' desire to wait and see others' experiences. In this paper, we explore optimal ways to resolve this issue via information design and characterize the optimal feedback mechanism to deter information free-riding. The optimal mechanism is cautious in the sense that it certainly terminates when the principal is pessimistic and terminates with some probability even when she is optimistic. The key driving force of this mechanism is the observation that the value of information varies across agents with different valuations: more precise information at the low end of the belief distribution is beneficial for agents with lower valuations but irrelevant for those with higher valuations. As such, a mechanism that is consistently more precise at the low end but obscure at the high end can induce agents with higher valuations to adopt earlier, thereby alleviating the free-rider problem, while still enabling those with lower valuations to make more informed decisions. We show that the optimal mechanism takes a simple form, which effectively reduces the principal's problem to an optimal stopping problem. We also show that the key principles identified in this paper are robust to various alterations in the underlying setup and can be applied to a range of social situations.

Appendix A: Proofs

Proof of Lemma 1. (i) We consider the incentive of the threshold agent v_{t-1}^*, given that $\mu_{t-1} = \underline{\mu}(v_{t-2}^*) \in M_{t-1}^A(\mu^{t-2})$. First, suppose on the contrary that $\mu_t = \underline{\mu}(v_{t-1}^*) \in M_t^N(\mu^{t-1})$. Then, by part (b) of the definition, the belief stays at $\underline{\mu}(v_{t-1}^*)$ for all future periods, meaning that there are no gains from waiting. Therefore, if the threshold agent deviates and chooses to wait in period $t-1$, he must choose to adopt in period t given that $\frac{v_{t-1}^*}{c} \geq \underline{\mu}(v_{t-2}^*) > \mu_T = \underline{\mu}(v_{t-1}^*)$. However, this is a contradiction because $\mu_t = \underline{\mu}(v_{t-1}^*) \in M_t^N(\mu^{t-1})$. Given this, we next show $\tilde{\mu}(v_{t-1}^*, v_{t-2}^*; \beta_t) \in M_t^N(\mu^{t-1})$. Suppose otherwise. Then, the threshold agent must adopt in period t with probability 1 if he deviates. This is a contradiction because the threshold agent would strictly prefer to adopt in period $t-1$.

(ii) We again consider the incentive of the threshold agent v_{t-1}^*. If $\mu_{t-1} \in M_t^N(\mu^{t-2})$, we have $v_{t-1}^* = v_{t-2}^*$ by definition. Also, by part (b), the belief stays constant for all future periods. Therefore, if it is not optimal for the threshold agent to adopt in period $t-1$, it is not optimal for him to adopt in period t. ■

Proof of Theorem 1. The proof consists of two properties, which we label as *termination* and *caution*. These two properties together imply the optimality of binary-message termination mechanisms.

Termination. We show that if there is no adoption in period t ($v_t^* = v_{t-1}^*$), then there will be no adoption in any future period as well ($v_s^* = v_t^*$ with probability 1 for all $s > t$). Consider some mechanism P that pauses the adoption process for one period and then resumes after that.[21] Fix some history μ^{t-1} and let v_{t-1}^* denote the threshold following μ^{t-1}. Also, let $v_t^*(\mu_t)$ and $v_{t+1}^*(\mu_t, \mu_{t+1})$ be the thresholds following (μ^{t-1}, μ_t) and $(\mu^{t-1}, \mu_t, \mu_{t+1})$, respectively. We then have some $\mu_t' \in M_t(\mu^{t-1})$ and $\mu_{t+1}' \in M_{t+1}(\mu^{t-1}, \mu_t')$ such that $v_{t+1}^*(\mu_t', \mu_{t+1}') < v_t^*(\mu_t') = v_{t-1}^*$. We now propose an alternative mechanism $\hat{P}$ that modifies P following $\mu_t = \mu_t'$ (while keeping everything else the same). Specifically, suppose the modified mechanism splits μ_t' into all $\mu_t \in M_{t+1}(\mu^{t-1}, \mu_t')$ and μ_t' such that $\hat{P}_t(\mu_t \mid \mu^{t-1}) = \delta P_t(\mu_t' \mid \mu^{t-1}) P_{t+1}(\mu_t \mid \mu^{t-1}, \mu_t')$ for all $\mu_t \in M_{t+1}(\mu^{t-1}, \mu_t')$ and $\hat{P}_t(\mu_t' \mid \mu^{t-1}) = (1-\delta) P_t(\mu_t' \mid \mu^{t-1})$. Moreover, (i) following any $\mu_t \in M_{t+1}(\mu^{t-1}, \mu_t')$ in period t, $\hat{P}$ implements the same allocation as P and induces no adoption in period T; (ii) following μ_t' in period t, $\hat{P}$ induces no adoption for all subsequent periods. Note that this modified mechanism keeps the continuation payoffs of all remaining agents, and hence the continuation surplus, unchanged (while the allocation following $\mu_t = \mu_t'$

[21] The same argument applies for a mechanism that pauses for multiple periods.

is clearly suboptimal). It is therefore without loss of generality to front-load all the information by setting $P_t(\mu \mid \mu^{t-1}) = 0$.

This argument suggests that there is no need to consider temporary suspension of the adoption process, as we can always split this into continuation and permanent termination in a profitable way. Alternatively, if $v_t^* = v_{t-1}^*$ for some t, that must mean permanent termination ($v_s^* = v_t^*$ with probability 1 for all $s > t$).

Caution. We show that we can increase the total surplus by stretching out any interior belief $\mu_t \in (\underline{\mu}(v_{t-1}^*), \frac{v_{t-1}^*}{c})$ into two outer points. Alternatively, this suggests that it is not optimal to assign a positive probability to any interior belief. We establish this result via induction.

We begin with the period-T surplus which can be written as

$$W_T(\mu_T, \mu_T c, v_{T-1}^*) = \int_{\mu_T c}^{v_{T-1}^*} (v - \mu_T c)\mathrm{d}F(v).$$

Taking derivative of W_T with respect to μ_T then yields $-c(F(v^*) - F(\mu_T c))$, which is increasing in μ_T, i.e., the period-T surplus is convex in μ_T. This ensures the caution property for period T. In the following, we show that if this property holds for all periods $s > t$, it also holds for period t.

Consider some arbitrary history μ^{t-1} and an interior belief $\mu_t = \mu_t' \in (\underline{\mu}(v_{t-1}^*), \frac{v_{t-1}^*}{c})$, where v_{t-1}^* is determined by μ^{t-1}. Let $v_t^* = v^*$ denote the threshold following (μ^{t-1}, μ_t'). Under the induction hypothesis, the IC constraint is given by

$$v^* - \mu_t' c = \delta P_{t+1}(\underline{\mu}(v^*) \mid \mu^{t-1}, \mu_t')(v^* - \underline{\mu}(v^*)c).$$

The Bayes plausibility, along with the consistency restriction, implies that there is an upper bound $\overline{P}_{t+1}(\underline{\mu}(v_t^*) \mid \mu^t)$ of the probability that can be assigned to the lowest possible belief. To obtain the upper bound, let ν be the probability that the principal is uninformed for a given μ_t, which must satisfy

$$\nu\underline{\mu}(v_{t-1}^*) + 1 - \nu = \mu_t.$$

Solving this then yields $\nu = \frac{1-\mu_t}{1-\underline{\mu}(v_{t-1}^*)}$. Since the principal observes no news in period t with probability $q(v_t^*, v_{t-1}^*)$, the upper bound is obtained as

$$\begin{aligned}\overline{P}_{t+1}(\underline{\mu}(v_{t-1}^*) \mid \mu^{t-1}, \mu_t') &= \frac{1-\mu_t'}{1-\underline{\mu}(v_{t-1}^*)} q(v_t^*, v_{t-1}^*) \\ &= (1-\mu_t') \frac{\underline{\mu}(v_{t-1}^*) e^{-\lambda(F(v_{t-1}^*)-F(v_t^*))} + 1 - \underline{\mu}(v_{t-1}^*)}{1-\underline{\mu}(v_{t-1}^*)} \\ &= (1-\mu_t') \frac{m e^{-\lambda(1-F(v_t^*))} + 1 - m}{1-m}, \end{aligned} \tag{9}$$

which is determined independently of v_{t-1}^*.

We now consider a modified mechanism $\hat{P}$ that applies a mean-preserving spread to μ_t' as in the two-period example. In the general case, however, we need to adopt a different strategy because the way we split a belief may affect the continuation probability due to the upper bound constraint. Specifically, we let $\hat{P}$ split μ_t' into $\mu_t' - \varepsilon$ and $\mu_t'' > \frac{v^*}{c}$ with probabilities $\frac{\mu_t'' - \mu_t'}{\mu_t'' - \mu_t' + \varepsilon}$ and $\frac{\varepsilon}{\mu_t'' - \mu_t' + \varepsilon}$, respectively, where $\varepsilon > 0$ is some arbitrarily small number, and moreover mix μ_t'' with the highest belief on the support of $P_t(\cdot \mid \mu^{t-1})$. Denote the higher belief by $\hat{\mu}$. Since the resulting mechanism $\hat{P}$ is arbitrarily close to P, all the past thresholds up to period $t-2$ (if $t \geq 3$) are unaffected by this modification. For $t > 1$, this modification increases the continuation payoff of agent v_{t-1}^*, raising the threshold slightly. Let $v_{t-1}^* = \hat{v}_{t-1}^*$ be the threshold after the modification and $v_t^* = \hat{v}^*$ the threshold following $\mu_t = \mu_t' - \varepsilon$.

To make the situation directly comparable, it is convenient to combine two contingencies ($\mu_t = \mu_t' - \varepsilon$ and $\mu_t = \hat{\mu}$) to represent the IC constraint. Since agent $\hat{v}^*$ adopts after $\mu_t = \mu_t' - \varepsilon$ but not after $\mu_t = \hat{\mu}$, given the induction hypothesis, the IC constraint combining the two contingencies can be written as

$$\frac{\mu_t'' - \mu_t'}{\mu_t'' - \mu_t' + \varepsilon} [\hat{v}^* - (\mu_t' - \varepsilon) c] = \delta \hat{R}_{t+1} (\hat{v}^* - \underline{\mu}(\hat{v}_t^*) c). \tag{10}$$

where

$$\hat{R}_{t+1} := \frac{\mu_t'' - \mu_t'}{\mu_t'' - \mu_t' + \varepsilon} \hat{P}_{t+1}(\underline{\mu}(\hat{v}_t^*) \mid \mu^{t-1}, \mu_t' - \varepsilon).$$

From (9), the upper bound for the modified mechanism is given by

$$\hat{P}_{t+1}(\underline{\mu}(\hat{v}^*) \mid \mu^{t-1}, \mu_t' - \varepsilon) \leq (\mu_t'' - \mu_t' + \varepsilon) \frac{m e^{-\lambda(1-F(\hat{v}^*))} + 1 - m}{1-m},$$

from which the upper bound for $\hat{R}_{t+1}$ is obtained as

$$\hat{R}_{t+1} \leq \frac{\mu_t'' - \mu_t'}{\mu_t'' - \mu_t' + \varepsilon}(\mu_t'' - \mu_t' + \varepsilon)\frac{me^{-\lambda(1-F(\hat{v}^*))} + 1 - m}{1-m} = (\mu_t'' - \mu_t')\frac{me^{-\lambda(1-F(\hat{v}^*))} + 1 - m}{1-m}.$$

Note that the upper bound depends on and is increasing in μ_t''. To ensure maximum flexibility, we set $\mu_t'' = 1$, so that the upper bound is given by

$$\hat{R}_{t+1} \leq (1 - \mu_t')\frac{me^{-\lambda(1-F(v^*))} + 1 - m}{1-m} = \overline{P}_{t+1}(\underline{\mu}(v^*) \mid \mu^{t-1}, \mu_t'),$$

when $\hat{v}^* = v^*$. Note also that since

$$\frac{1-\mu_t'}{1-\mu_t'+\varepsilon}[v^* - (\mu_t' - \varepsilon)c] > v^* - \mu_t' c,$$

$\hat{R}_{t+1}$ must be set higher than $P_{t+1}(\underline{\mu}(v^*) \mid \mu^{t-1}, \mu_t')$ to implement the same threshold as in P.

There are two cases to consider, depending on whether the upper bound is binding. We start with the case where $P_{t+1}(\underline{\mu}(v^*) \mid \mu^{t-1}, \mu_t') < \overline{P}_{t+1}(\underline{\mu}(v^*) \mid \mu^{t-1}, \mu_t')$, so that the upper bound is not binding. In this case, we can find $\hat{R}_{t+1} \in (P_{t+1}(\underline{\mu}(v^*) \mid \mu^{t-1}, \mu_t'), \overline{P}_{t+1}(\underline{\mu}(v^*) \mid \mu^{t-1}, \mu_t'))$ that can implement the same threshold as P. With the same threshold, the modified mechanism $\hat{P}$ can implement the same allocation as P for all future periods following $\mu_{t+1} = \underline{\mu}(v^*)$. The only difference is that $\mu_{t+1} = \underline{\mu}(v^*)$ is now realized with a higher probability so that agents $v \in [0, v^*)$ are better off. In addition, agents $v \in [v^*, \hat{v}^*)$ also benefit from the modification.

Next, consider the case where $P_{t+1}(\underline{\mu}(v^*) \mid \mu^{t-1}, \mu_t') = \overline{P}_{t+1}(\underline{\mu}(v^*) \mid \mu^{t-1}, \mu_t')$, so that the upper bound is binding. In this case, since $\hat{R}_{t+1}$ cannot be raised any further, we have a lower threshold $\hat{v}^* < v^*$ such that

$$\frac{1-\mu_t'}{1-\mu_t'+\varepsilon}[\hat{v}^* - (\mu_t' - \varepsilon)c] = \delta \hat{R}_{t+1}(\hat{v}^* - \underline{\mu}(\hat{v}^*)).$$

With a lower threshold, we can induce a belief lower than $\underline{\mu}(v^*)$. This means that we can implement the same allocation following $\mu_t = \mu_t' - \varepsilon$ by inducing $\mu_{t+1} = \underline{\mu}(v^*)$ (instead of the lowest possible belief). Since this modification, though suboptimal, gives the same payoff to all agents $v < \hat{v}^*$ following $\mu_{t+1} = \underline{\mu}(v^*)$, it is weakly beneficial for those agents. Moreover, since $\mu_{t+1} = \underline{\mu}(v^*)$ is still realized with probability $\overline{P}_{t+1}(\underline{\mu}(v^*) \mid \mu^{t-1}, \mu_t')$, we have

$$\frac{1-\mu_t'}{1-\mu_t'+\varepsilon}[v - (\mu_t' - \varepsilon)c] > \delta \overline{P}_{t+1}(\underline{\mu}(v^*) \mid \mu^{t-1}, \mu_t')(v - \underline{\mu}(v^*)c),$$

for agents $v \in [\hat{v}^*, v^*)$, meaning that they are also better off. Finally, it is easy to see that the payoffs of agents $v \in [v^*, \hat{v}_{t-1})$ increase while the payoffs of agents $v \in [\hat{v}^*_{t-1}, 1]$ remain unchanged. This argument shows that the proposed modification is Pareto improving and hence increases the total surplus. ■

Proof of Theorem 2. Since we have already established part (i) in the main text, here we focus on part (ii). We first show $\beta_T < 1$ as $\delta \to 1$. To this end, it is important to observe that if $\beta_T = 1$, v^*_{T-1} converges to v^*_{T-2} as $\delta \to 1$, i.e., no agents adopt in period $T-1$, and agents $v \in [\mu_{T-1}c, v^*_{T-2})$ adopt in period T. However, the same allocation can be implemented a period earlier by setting $\beta_T = 0$, in which case agents $v \in [\mu_{T-1}c, v^*_{T-2})$ adopt in period $T-1$ (while no agents adopt in period T). As $\delta \to 1$, these two allocations are equivalent and yield the same continuation surplus. The existence of an interior solution ($\beta_T < 1$) is then implied by part (i) which states that it is never optimal to set $\beta_T = 0$.

This argument suggests that v^*_{T-1} is bounded away from v^*_{T-2} as $\delta \to 1$. Let $v^*_{T-1} = v^*_a$ and $v^*_T = v^*_b$ be the optimal thresholds in the limit where $v^*_{T-2} > v^*_a > v^*_b$. Now consider period $T-1$ and suppose $\beta_{T-1} = 1$ in the optimal mechanism as $\delta \to 1$. Then, in the limit, v^*_{T-2} converges to v^*_{T-3} (where $v^*_0 = 1$ if $T = 3$), and hence the period-$T-2$ surplus converges to 0. However, the principal can instead implement the same allocation a period earlier by inducing $v^*_{T-2} = v^*_a$ and $v^*_{T-1} = v^*_T = v^*_b$. Observe that the implemented allocation is suboptimal because we must have $v^*_T < v^*_{T-1}$ in the optimal mechanism. We thus conclude $\beta_{T-1} < 1$ in the optimal mechanism if δ is sufficiently close to 1. We can then apply this argument repeatedly to establish the proposition. ■

Proof of Theorem 3. Recall that we establish the termination property by showing that if the game is in the no-adoption phase in one period and moves back to the adoption phase in the next, we can always implement the same allocation a period earlier by a modified mechanism. Since the modified sequence is entirely lower than the original sequence, Assumption 2 ensures that the principal's payoff is still weakly higher in the modified mechanism, and the termination property continues to hold. Note also that the caution property holds for a fixed sequence of thresholds if the upper bound defined in the proof of Theorem 1 is not binding. If it is binding, we consider a modified mechanism (a mean-preserving spread) that implements a threshold slightly lower than the original one in period t and the same allocation after period t. Again, Assumption 2 ensures that the principal's payoff is weakly higher in the modified mechanism, and the caution property continues to hold. This proves Theorem 1 for any $\alpha > 0$. ■

Proof of Theorem 4. Observe that the termination property is independent of the underlying information-generating process and holds as it is. Below, we thus focus on establishing the caution property, i.e., $M_t^A(\mu^{t-1})$ contains only the lowest possible belief $\underline{\mu}(v_{t-1}^*) = 0$ when the game is in the adoption phase. To this end, we first note that $M_t^A(\mu^{t-1}) = [\underline{\mu}(v_{t-1}^*), \frac{v_{t-1}^*}{c})$ and $M_t^N(\mu^{t-1}) = [\frac{v_{t-1}^*}{c}, \overline{\mu}(v_{t-1}^*)]$ in any optimal mechanism satisfying the termination property. Alternatively, we need to show that $v_t^* < v_{t-1}^*$ if and only if $\mu_t \in [\underline{\mu}(v_{t-1}^*), \frac{v_{t-1}^*}{c})$. The necessity is obvious because it is strictly better to wait for all agents $v \in [0, v_{t-1}^*)$ if $\frac{v_{t-1}^*}{c} \leq \mu_t$. To show the sufficiency, suppose on the contrary that $v_t^* = v_{t-1}^*$ even if $\frac{v_{t-1}^*}{c} > \mu_t$. Then, by the termination property, the belief stays constant in all future periods. In period T, agents whose valuations are arbitrarily close to v_{t-1}^* have an incentive to adopt because $\frac{v}{c} > \mu_T = \mu_t$ for v arbitrarily close to v_{t-1}^*. As this is a contradiction, we must have $M_t^A(\mu^{t-1}) = [\underline{\mu}(v_{t-1}^*), \frac{v_{t-1}^*}{c})$ and $M_t^N(\mu^{t-1}) = [\frac{v_{t-1}^*}{c}, \overline{\mu}(v_{t-1}^*)]$ in an optimal mechanism.

Given this, we next establish the following result.

Lemma 2 *In the optimal mechanism, either* $\{0\} \in M_t^A(\mu^{t-1})$, $\{\overline{\mu}(v_{t-1}^*)\} \in M_t^N(\mu^{t-1})$ *or both.*

Proof. We show that if $M_t^A(\mu^{t-1}) \neq \{0\}$ and $M_t^N(\mu^{t-1}) \neq \{\overline{\mu}(v_{t-1}^*)\}$ in some mechanism P, we can construct a modified mechanism $\hat{P}$ that can improve upon P for a given history μ^{t-1}. First, observe that we can merge all the beliefs in $[\frac{v_{t-1}^*}{c}, \overline{\mu}(v_{t-1}^*)]$ into a single one without affecting the thresholds. Therefore, it is without loss of generality to focus on mechanisms that have only one element, say μ_H, in $M_t^N(\mu^{t-1})$. Pick some $\mu_L \in M_t^A(\mu^{t-1}) \setminus \{\underline{\mu}(v_{t-1}^*)\}$ and apply a mean-preserving spread to μ_H, so that μ_H is split into μ_L and $\mu_H + \varepsilon$ for some arbitrarily small $\varepsilon > 0$. This operation raises the payoff of agent v_{t-1}^* and thus (slightly) raises the threshold in period $t-1$ to some $\hat{v}_{t-1}^* > v_{t-1}^*$. Note that with a higher threshold $\hat{v}_{t-1}^*$, $\overline{\mu}(\hat{v}_{t-1}^*)$ is smaller than $\underline{\mu}(v_{t-1}^*)$ but is still higher than $\mu_H + \varepsilon$. In period t, some agents adopt if $\mu_t \in (0, \frac{\hat{v}_{t-1}^*}{c})$ while no agents adopt if $\mu_t \in [\frac{\hat{v}_{t-1}^*}{c}, \overline{\mu}(\hat{v}_{t-1}^*))$. Note that $\hat{v}_{t-1}^* > v_{t-1}^*$, $\hat{P}_t(\mu_L \mid \mu^{t-1}) > P_t(\mu_L \mid \mu^{t-1})$, and $\hat{P}_t(\mu_H + \varepsilon \mid \mu^{t-1}) < P_t(\mu_H \mid \mu^{t-1})$. Therefore, the probability that μ_t lies in $(0, \frac{\hat{v}_{t-1}^*}{c})$ becomes weakly higher, leading to an increase in the continuation probability. This means that the payoffs of agents adopting in period t and after increase while the payoffs of those adopting before period t remain the same, suggesting that the modification is Pareto-improving. ■

Given the lemma, we now show $M_t^A(\mu^{t-1}) = \{0\}$, which allows us to establish the theorem. To this end, suppose there is some interior belief $\mu_t > 0$ in $M_t^A(\mu^{t-1})$ for some μ^{t-1} and $t \geq 2$. We then claim that for any history that follows μ^{t-1}, there must be some period $s > t$ such that $M_s^A(\mu^{s-1}) = \{0\}$ (where s is the earliest period that has this feature). Suppose otherwise.

We then have some history μ^{T-1} such that $M_T^A(\mu^{T-1})$ includes some interior belief $\mu_T > 0$ or is empty. If $M_t^A(\mu^{T-1})$ includes an interior belief, since this is the last period, we can split this into $\underline{\mu}(v_{T-1}^*)$ and $\frac{v_{T-1}^*}{c}$ and improve the payoffs of all types. In period T, therefore, we either have $M_T^A(\mu^{T-1}) = \{\underline{\mu}(v_{T-1}^*)\}$ or $M_T^A(\mu^{T-1}) = \varnothing$. If there is a history μ^{T-1} such that $M_T^A(\mu^{T-1}) = \varnothing$, then period $T-1$ is effectively the last period, and we can apply the same argument to show that $M_{T-1}^A(\mu^{T-2}) = \{\underline{\mu}(v_{T-2}^*)\}$ or $M_{T-1}^A(\mu^{T-2}) = \varnothing$. We can apply this argument repeatedly to provide the claim.

Given this, in period $s-1$, we have

$$v_{s-1}^* - \mu_{s-1}c = \delta P_s(0 \mid \mu^{s-1}) v_{s-1}^*,$$

from which we obtain

$$\delta P_s(0 \mid \mu^{s-1}) = \frac{v_{s-1}^* - \mu_{s-1}c}{v_{s-1}^*}.$$

Note that the continuation payoff of agent $v \in (0, v_{s-1}^*)$ is

$$\delta P_s(0 \mid \mu^{s-1}) v = \frac{v_{s-1}^* - \mu_{s-1}c}{v_{s-1}^*} v.$$

Now consider a mean-preserving spread that splits μ_{s-1} into $\frac{v_{s-2}^*}{c}$ with probability $\frac{\mu_{s-1}c}{v_{s-2}^*}$ and 0 with the remaining probability. Note that this mean-preserving spread does not affect the payoff of agent v_{s-2}^* but raises the payoffs of agents $v \in [v_{s-1}^*, v_{s-2}^*]$. Also, with this modification, agents $v \in (0, v_{s-1}^*)$ adopt in period $s-1$ when the belief is 0 and obtain $\frac{v_{s-2}^* - \mu_{s-1}c}{v_{s-2}^*} v$, which is larger than the payoff in the original mechanism because $v_{s-2}^* > v_{s-1}^*$. Therefore, $M_t^A(\mu^{t-1}) = \{0\}$ for all t as long as the game is in the adoption phase, which proves that Theorem 1 holds in the good-news case. ■

Appendix B: Free-rider incentive in collective experimentation

In this appendix, we show that the amount of adoption under full disclosure is generally insufficient compared to the first best for any (continuous) type distribution and any discount factor $\delta \in (0, 1]$. Consider a two-period model in which each agent's valuation $v \in [0, 1]$ is distributed according to some distribution function F with full support. The payoff to agent v of adopting is $v - \omega c$, where $c \geq 1$ is some constant.

Let v_F^* be the threshold under full disclosure that solves

$$v_F^* - mc = \delta(q(v_F^*)v_F^* - r(v_F^*)c). \tag{11}$$

Here, we use $q(v^*) := m(1 - p(1 - v^*)) + 1 - m$ and $r(v^*) := m(1 - p(1 - v^*))$ to save notation where $p : [0, 1] \to [0, 1]$ is a strictly increasing function. The total surplus for a given threshold v^* in the two-period case can be written as

$$W_1(m, v^*, 1) = \int_{v^*}^{1} (v - mc)\mathrm{d}F(v) + \delta \int_{\underline{\mu}(v^*)c}^{v^*} (q(v^*)v - r(v^*)c)\mathrm{d}F(v),$$

where $\underline{\mu}(v^*) = \frac{r(v^*)}{q(v^*)}$ is the principal's belief when she is uninformed. Now consider a social planner who can unilaterally impose v^* to maximize the total surplus. The first-order condition is then given by

$$\frac{\partial}{\partial v^*} W_1(m, v^*, 1) = -[v^* - mc - \delta(q(v^*)v^* - r(v^*)c)]f(v^*) + \delta m p'(1 - v^*) \int_{\underline{\mu}(v^*)c}^{v^*} (v - c)\mathrm{d}F(v). \tag{12}$$

Observe that the first term is simply the IC constraint that must equal 0 at v_F^*. Evaluating (12) at $v^* = v_F^*$ thus yields

$$\frac{\partial}{\partial v^*} W_1(m, v_F^*, 1) = \delta m p'(1 - v_F^*) \int_{\underline{\mu}(v_F^*)c}^{v_F^*} (v - c)\mathrm{d}F(v).$$

The threshold under full disclosure is too high if $\frac{\partial W_1}{\partial v^*} < 0$. A sufficient condition for this is $v_F^* \leq c$, which always holds because it is optimal for any agent to adopt in period 1 if $v \geq c$.

References

Au, Pak Hung, "Dynamic Information Disclosure," *Rand Journal of Economics*, 2015, *46* (4), 791–823.

Baccara, Mariagiovanna, Gilat Levy, and Ronny Razin, "Research Waves," *mimeo*, 2024.

Ball, Ian, "Dynamic Information Provision: Rewarding the Past and Guiding the Future," *Econometrica*, 2023, *91* (4), 1363–1391.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 1992, *100* (5), 992–1026.

Bolton, Patrick and Christopher Harris, "Strategic Experimentation," *Econometrica*, 1999, *67* (2), 349–374.

Bonatti, Alessandro and Johannes Hörner, "Collaborating," *American Economic Review*, 2011, *101*, 632–663.

Che, Yeon-Koo and Johannes Hörner, "Recommender Systems as Mechanisms for Social Learning," *The Quarterly Journal of Economics*, 2018, *133* (2), 871–925.

_ **and Konrad Mierendorff**, "Optiaml Dynamic Allocation of Attention," *American Economic Review*, 2019, *109* (8), 2993–3029.

Correia da Silva, Joao and Takuro Yamashita, "Information Design in Repeated Interaction," *mimeo*, 2024.

Ely, Jeffrey C., "Beeps," *American Economic Review*, 2017, *107* (1), 31–53.

_ **and Martin Szydlowski**, "Moving the Goalposts," *Journal of Political Economy*, 2020, *128* (2), 393–782.

Frick, Mira and Yuhta Ishii, "Innovation Adoption by Forward-Looking Social Learners," *Theoretical Economics*, 2024, *forthcoming.*

Halac, Marina, Navin Kartik, and Qingmin Liu, "Optimal Contracts for Experimentation," *Review of Economic Studies*, 2016, *83* (3), 1040–1091.

– , – , **and** – , "Contests for Experimentation," *Journal of Political Economy*, 2017, *125* (5), 1523–1569.

Hamel, Liz, Lunna Lopes, and Mollyann Brodie, "KFF COVID-19 Vaccine Monitor: What Do We Know About Those Who Want to "Wait and See" Before Getting a COVID-19 Vaccine?," *KFF Poll*, 2021.

Kamenica, Emir and Matthew Gentzkow, "Bayesian Persuasion," *American Economic Review*, 2011, *101* (6), 2590–2615.

Keller, Godfrey and Sven Rady, "Breakdowns," *Theoretical Economics*, 2015, *10* (1), 175–202.

– , – , **and Martin Cripps**, "Strategic Experimentation with Exponential Bandits," *Econometrica*, 2005, *73* (1), 39–68.

Knoepfle, Jan and Julia Salmi, "Dynamic Evidence Disclosure: Delaying the Good to Accelerate the Bad," *arXiv: 2406.11728*, 2024.

Kremer, Ilan, Yishay Mansour, and Motty Perry, "Implementing the "Wisdom of the Crowd"," *Journal of Political Economy*, 2014, *122* (5), 988–1012.

Manso, Gustavo, "Motivating Innovation," *Journal of Finance*, 2011, *66* (5), 1823–1860.

Nikandrova, Arina and Romans Pancs, "Dynamic Project Selection," *Theoretical Economics*, 2018, *13* (1), 115–143.

Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov, "Persuading the Principal to Wait," *Journal of Political Economy*, 2020, *128* (7), 2473–2837.