Optimal Experimentation Design with Secret Repetition

Zheng Gong and Qianjun Lyu*

September 8, 2018

Abstract

We study a persuasion game with limited commitment in which a biased sender designs and conducts costly experiments to acquire information which he can conceal or reveal. The sender commits to the experiment design, but he can secretly repeat experiments and selectively report the outcomes. In the benchmark model, the optimal experiment turns out to be a one-round experiment and the sender truthfully discloses the experiment outcome. The cost of an experiment is a measure of credibility. Higher credibility leads to less informative experiment which lowers the receiver's payoff. With general payoff function of the sender, the above results remain with mild restrictions. We geometrically characterize the optimal experiment using the same concavification with Kamenica and Gentzkow (2011) but within a refined belief space.

JEL classification: D83, D82

Keywords: persuasion, information acquisition, selective information disclosure, concavification, credibility frontier.

*Rotman School of Management, Univerisity of Toronto, zheng.gong16@rotman.utoronto.ca; Faculty of Business and Economics, University of Hong Kong, lyu2014@hku.hk. Qianjun would like to thank Wing Suen for his invaluable comments and suggestions. We would also like to thank Heski Bar-Isaac, Ettore Damiano, Xianwen Shi, Yupu Zhang and Ruizhi Zhu for helpful discussions.

1 Introduction

Studies show that empirical researchers may hide the negative tests, provide multiple tests without proper adjustment, or engage in data mining and "p-hacking" to produce "significant" results (R.Harvey (2017), Brodeur et al. (2016)). Firms can switch to another rating agency in the hope of improving their rating (Sangiorgi et al. (2009), Skreta and Veldkamp (2009), Bongaerts et al. (2012), Bolton et al. (2012)). Similarly, financial analysts who are incapable to pass CFA Exam on the first attempt usually retake the examination in order to convince potential employers of their quality. In these situations, intending to change a receiver's action, a biased sender implements costly experiments and strategically repeats it in order to provide results of their interests. Repetition is unobserved by the receiver.

This paper formalizes a persuasion game with limited commitment where a biased sender commits to his endogenous experiment design but he can secretly repeat the experiment and selectively present the experiment outcomes. Each time he conducts the experiment, he has to pay a fixed cost. Following Kamenica and Gentzkow (2011), we focus on the case that the sender and the receiver share common prior beliefs. We assume that the sender's preference only depends on the receiver's action, regardless of the true state of the world. Then, with a rational receiver who realizes the sender's intention of selective information disclosure, is the sender still capable to affect her decision? To what extent will the receiver discount sender's evidence? What is the sender's optimal experiment design and how informative is it?

Consider a simple example of a think tank (he) trying to persuade the policymaker (she) to accept his policy suggestion. There are two states of the world: *Good* and *Bad*. The think tank and policymaker share the same prior belief $\mu_0(Good) = 0.3$. If the policymaker accepts the suggestion, she gets utility 4 when the state is *Good* and she loses utility 6 when the state is *Bad*. If she remains status quo, her utility is 0. Therefore she accepts as long as the belief of state being *Good* exceeds 0.6. The think tank gains utility 7 if his suggestion is approved and 0 otherwise, regardless of the state.

The think tank designs an experiment to acquire information and proposes this design to the policymaker so that he cannot switch to another experiment. Different from full commitment, the think tank can repeat the experiment and selectively disclose the experiment outcomes. However, at each round he implements the experiment, he needs to pay a fixed cost $c = 1$. The policymaker will observe the experiment outcomes disclosed by the think tank but she cannot distinguish at which round the outcomes are produced.

If the think tank has to communicate every experiment outcome to the policymaker (full commitment), then Kamenica and Gentzkow (2011) suggest that the uniquely optimal experiment is binary signal:

$$\begin{aligned} Pr(g|Good) &= 1 \qquad & Pr(b|Good) &= 0 \\ Pr(g|Bad) &= \frac{2}{7} \qquad & Pr(b|Bad) &= \frac{5}{7} \end{aligned} \tag{1}$$

Here, the policymaker is just indifferent between accepting or not when she observes a good signal. The think tank's ex-ante payoff is $3.5 - 1 = 2.5$.

However, this payoff is not attainable in our game. Suppose the think tank performs the same experiment with the above signal structure and the policymaker accepts the suggestion when one good signal is disclosed. If the first experiment generates a bad signal, the think tank's posterior belief reduces to 0. But the expected gain of next round is $\frac{2}{7} \times 7 = 2$ which is higher than the cost of an additional round of experiment. Thus, even if the think tank realizes the true state is *Bad*, he will keep repeating until he obtains one good signal. The policymaker can anticipate this and therefore does not update her belief if only one good signal is disclosed. Thus, she will not accept the think tank's suggestion.

We suggests that the optimal signal of this example is more informative than the above one:

$$\begin{aligned} Pr(g|Good) &= 1 \qquad & Pr(b|Good) &= 0 \\ Pr(g|Bad) &= \frac{1}{7} \qquad & Pr(b|Bad) &= \frac{6}{7} \end{aligned} \tag{2}$$

With this information structure, it is less frequently to generate in *Bad* state. Therefore, the think tank is more reluctant to repeat. Conditional on getting a bad signal at the first round, the expected gain of playing another round is $\frac{1}{7} \times 7 = 1$ which is equal to the cost. In other words, the think tank is indifferent between stopping and performing another experiment in the hope of obtaining a good signal. In this case, we assume he stops. With one bad signal in hand, the think tank can only report truthfully since he can not falsify the information. On the other hand, if he is lucky to obtain one good signal in the first round, he will report it for sure. Thus, the policymaker believes the disclosed outcome and she strictly prefers to accept the suggestion if being provided a good signal. The ex-ante payoff of the think tank is $0.3 \times 7 + 0.7 \times \frac{1}{7} \times 7 - 1 = 1.8$ which is smaller than the ex-ante payoff with full commitment. This is commonly true. Because without full commitment, the think tank's optimization problem is constrained by his incentive constraints. This lowers his ex-ante payoff.

One lesson of this simple example is that implementing a more informative experiment (i.e. an experiment which generates good signal less frequently in *Bad* state) is a commitment device to restore his credibility: a more informative experiment restricts the sender from repetition since the expected gain from additional round cannot compensate the cost. In particular, the sender only performs the optimal experiment once even he is allowed to repeat. The second thing is that the receiver is better off compared to the full commitment case: in KG, the sender designs an experiment to make the receiver indifferent between the two actions. However, here, the optimal signal is more informative, therefore, the receiver is able to make more precise decisions.

There are three features of the optimal binary signal structure in this example: (1) the probability of good signal conditional on *Good* state is 1; (2) one good signal is informative enough to convince the receiver; (3) the sender only performs the experiment once and truthfully discloses the experiment outcomes. These three features generally hold in our benchmark model where both the state space and the receiver's action space are binary (shown in section 2.2). For the first feature, if the state is *Good*, the sender and receiver share common interests because they both want the suggestion to be accepted. Thus, conditional on the state being *Good*, the sender wants the experiment to generate the positive signal for sure since it is the fastest way to collect all positive signals required by the receiver. For the second feature, what if the sender proposes a less informative experiment design such that the receiver would require more than one good signal to accept his suggestion? At the first glance, it may be a better alternative because now the sender has to undertake the experiment more than once and this increase his expected cost which restores his credibility. However, we prove that the gain from higher credibility is smaller than the cost increment. Therefore, the optimal experiment follows the feature that one good signal is informative enough to change the receiver's action. Given the first two features, the sender' posterior belief will reduce to 0 if he observes a bad signal in the first round. Thus, one incentive constraint for him is that his expected gain of an additional round of experiment cannot exceed the cost, otherwise, he will keep repeating the experiment until he obtains one good signal. The sender's optimality requires this constraint to bind so that he is just indifferent between stopping with a bad signal and repeating in the hope of one good signal. All these suggest the last feature: the sender only performs the optimal experiment once no matter which signal realization is produced.

In section 2.3, we show how the optimal experiment changes with the experiment cost. The experiment cost is a measure of credibility. Higher cost means higher credibility. Thus, the sender can increase the chance of generating good signal in *Bad* state without violating his incentive constraint. The optimal experiment is less informative since the posterior belief of *Good* state given one good signal is lower. If the cost is high enough, this posterior belief could be low enough and equal to the receiver's threshold of approval. In other words, the receiver is indifferent between accepting and rejecting. The optimal experiment design is the same as the case when the sender has full commitment power.

In Section 3, we generalize the sender's payoff function. To solve the problem, we combine the strategy-based approach and the belief-based approach (adopted by the Bayesian Persuasion literature). With some restrictions on the sender's payoff function, we find that the sender's optimal experiment is still a one-shot experiment, which can be geometrically characterized using the same concavification with Kamenica and Gentzkow (2011) but within a refined belief space. The refined belief space is characterized by a credibility frontier such that if the sender's payoff is lower than this frontier, he can credibly implement the experiment only once.

Related Literature. Our paper is related to the literature of persuasion. Kamenica and Gentzkow (2011) (KG from now on) uses a belief-based approach to pin down the highest payoff that the sender is able to achieve if he has full commitment power (i.e the sender commits to the information structure and the realization of signals is truthfully communicated to the receiver). Chakraborty and Harbaugh (2010) discuss persuasion by cheap talk and they assume the sender's preference is state-independent. They suggest that if there are multidimensional states, the sender can benefit from strategic communication if his preference is quasiconvex. Lipnowski and Ravid (2018), (hereinafter referred to as LR) adopt the approach in Bayesian persuasion, they claim a general geometric characterization of the value of cheap talk, pinned down by the quasiconcave envelope of the sender's value function. To achieve this result, they need the receiver to mix between different actions. Actually, KG and LR identify the range that the highest ex-ante payoff the sender can achieve, where upper bound of this range is generated by full credibility while the lower bound comes from no credibility in the sense of cheap talk.

We also speak to the literature that studies selective information disclosure and limited commitment. With exogenous information generating process, Di Tillio et al. (2017) investigate the persuasion outcome if a researcher can use private information to manipulate the experiment, for instance, he can assign subjects to treatment based on their baseline outcomes. They show that in some situations, the researcher can even benefit from his manipulation. In our paper, the information generating process is exogenous and our sender manipulates the experiment by selective information disclosure. Lipnowski et al. (2018) discuss persuasion vis weak institutions. Specifically, the sender privately learns the research outcomes and whether he can influence the report to the message of his choice. They conclude that the receiver can benefit from a reduction of the sender's credibility and small decreases in credibility often lead to large payoff losses for the sender. They measure credibility by the exogenous proportion of the sender who cannot influence the report. In our paper, the cost of each experiment measures the level of credibility. But the sender can covertly raise his credibility by run experiment more than once. Henry and Ottaviani (2018) analyze a dynamic costly information acquisition problem of an informer aiming to persuade an evaluator to approve an activity. One implication is that the ex-ante payoff of the informer converges to the value of optimal signal in KG when both discount factor and cost of research converge to zero. In our paper, even the cost is zero, the optimal signal structure departs from KG.

2 Benchmark Model and Equilibrium Analysis

2.1 The Model

There are one sender (he) and one receiver (she), named Sender and Receiver separately. There is an unknown state of the world, $\omega \in \{B, G\}$, with prior belief $\mu_0 \in (0,1)$ that $\omega = G$, shared by both Sender and Receiver. Receiver is a Bayesian decision maker and her action space is binary, $A = \{0, 1\}$. 1 is for approval and 0 is for remaining status quo. Receiver's preference is state-dependent: $u : A \times \Omega \to \mathbb{R}$. Receiver gets positive utility if she accepts Sender's suggestion in state G or she rejects Sender's suggestion in state B. Otherwise, she gets negative utility. Same as KG, we assume that Receiver takes the action that maximizes Sender's expected utility if she is indifferent between some actions at a given belief. Thus, Receiver's decision rule can be reduced to a cut off policy:

$$\begin{cases} a^*(\mu^R) = 1 & if\ \mu^R \geq \mu^* \\ a^*(\mu^R) = 0 & if\ \mu^R < \mu^* \end{cases} \tag{3}$$

To make the solution non-trivial, we assume $\mu^* > \mu_0$ so that Receiver remains status quo under prior belief. Sender's preference only depends on Receiver's action, regardless of the true state of the world. His utility equals to a positive number Q if $a = 1$, otherwise his utility is 0. Since Receiver's decision rule only depends on her belief, we can write down Sender's value function as a function of Receiver's belief, $v(\mu^R)$.

$$\begin{cases} v(\mu^R) = Q & if\ \mu^R \geq \mu^* \\ v(\mu^R) = 0 & if\ \mu^R < \mu^* \end{cases} \tag{4}$$

In order to persuade Receiver to accept his suggestion, Sender can take experiments and selectively disclose results from these experiments. An experiment π consists of a finite realization space $\mathcal{S}$ and a set of distributions $\pi(\cdot|\omega)_{\omega\in\Omega}$ over $\mathcal{S}$. Π is the family of all possible experiments. Each round Sender conducts the experiment, he pays a cost $c < Q$. At the first stage, Sender chooses whether to enter this persuasion game. If he enters, he chooses an experiment and Receiver observes Sender's choice of experiment. Departs from KG, Sender can secretly repeat the experiment and selectively disclose the experiment outcomes. The disclosed outcomes are viewed as hard evidence that cannot be falsified. We assume the repeated experiments are i.i.d. conditional on the state. At round n Sender implements the experiment, he obtains the realization $h_n \in \mathcal{S}$. $H_n = h_1 \cup h_2 \cup ...h_n$ is the set of all realizations over historical n-round experiments he has conducted. We denote $n = N$ as the last round Sender conducts experiment. Given the assumption that information is hard evidence, the message Sender discloses to Receiver must be in the set of historical realizations, $m \subset H_N$. Receiver can only observe the realizations disclosed by Sender and she cannot distinguish at which round(s) those realization(s) is produced. Thus, the message space M is actually a production set of $\mathcal{S}$ for all possible rounds of experiment. Therefore, $H_N \subset M$.

The game is an extensive game where at the first stage Sender designs an experiment and commits to this experiment design, then the game moves to the second stage[1]: Simultaneously, Receiver chooses her accepting standard and Sender chooses his stopping rule and disclosure policy.

The solution concept we use is (Sender Preferred) Subgame Perfect Bayesian Equilibrium: given Sender's choice of experiment π and the disclosed experiment outcomes m, Receiver forms a posterior belief μ_m and then takes an action $a^*(\mu_m) = arg\ \max_{a \in A} u(a, \mu_m)$. Then taking Receiver's behavior and the experiment π as given, Sender first decides whether to stop or repeat after each H_n. We assume that whenever Sender is indifferent between stopping and continuing, he stops. If he chooses to stop, he discloses the message $m \subset H_N$ that maximizes his utility in the subgame. Furthermore, the posterior belief is obtained from μ_0, given Sender's stopping rule and disclosure policy in the subgame. Anticipating Sender and Receiver's behavior in the subgame, Sender chooses an experiment design which maximized his ex-ante utility at the first stage.

In the second state (subgame): Receiver's accepting strategy $\alpha : \Pi \times M \to A$. Sender's strategy consists of two maps: (1) stopping rule $\xi : \Pi \times H_n \to \{0, 1\}$, where 0 is for stop and 1 is for continue, (2) disclosure policy: $\delta : H_N \to H_N$. The belief system $\beta : \Pi \times M \to \Delta\Omega$. Back to the first stage, Sender chooses his optimal experiment π^* given all (α, ξ, δ). Thus, a Sender preferred equilibrium[2] consists of

[1] The second stage actually consists of many substages because Sender may conduct the experiment many times. But it is not necessary to describe all possible paths. Therefore we simply use the subgame of the second stage to describe the game after Sender commits to the experiment design.

[2] There may exists multiple equilibria in the second stage's subgame given a pair of (p, q). Sender chooses the one gives him the highest ex-ante payoff.

three maps and the experiment design: $(\alpha^*, \xi^*, \delta^*, \pi^*)$.

1. $\alpha(\pi, m) = \arg\max_{a \in A} u(a, \beta(\pi, m))$.

2. $\xi(\pi, H_n) = \arg\max_{s \in \{0,1\}} \{v_n(\alpha(\pi, \delta(H_n))|s = 0), Ev_{n+1}(\alpha(\pi, \delta(H_{n+1}))|s = 1) - c\}$.

3. $\delta(H_N) = \arg\max_{m \subset H_N} v(\alpha(m))$.

4. β is obtained given μ_0, Sender's stopping rule and disclosure policy, using Bayes' rule whenever possible.

5. $\pi^* = \arg\max_{\pi \in \Pi} W(\pi|(\alpha, \xi, \delta))$.

$W(\pi|(\alpha, \xi, \delta))$ is the ex-ante payoff of Sender given his choice of experiment π and the strategies in the subgame (α, ξ, δ). Note that W not only contains Sender's expected gain from Receiver's approval but also contains Sender's expected cost from conducting the experiment. $(\alpha^*, \xi^*, \delta^*)$ is the Sender preferred PBE in the subgame conditional on Sender's choice of experiment π^*. Denote $W^* = W(\pi^*|(\alpha^*, \xi^*, \delta^*))$.

We focus on experiment with binary structures, since both state space and Receiver's action space are binary. We denote the experiment as a pair $\pi = (p, q)$, where $(p, q) = (\pi(g|G), \pi(b|B))$. Without loss of generality, we assume $p \geq 1 - q$, so g is a good signal in the sense that it is evidence of state G. The time-line of this game is shown in Figure 1.

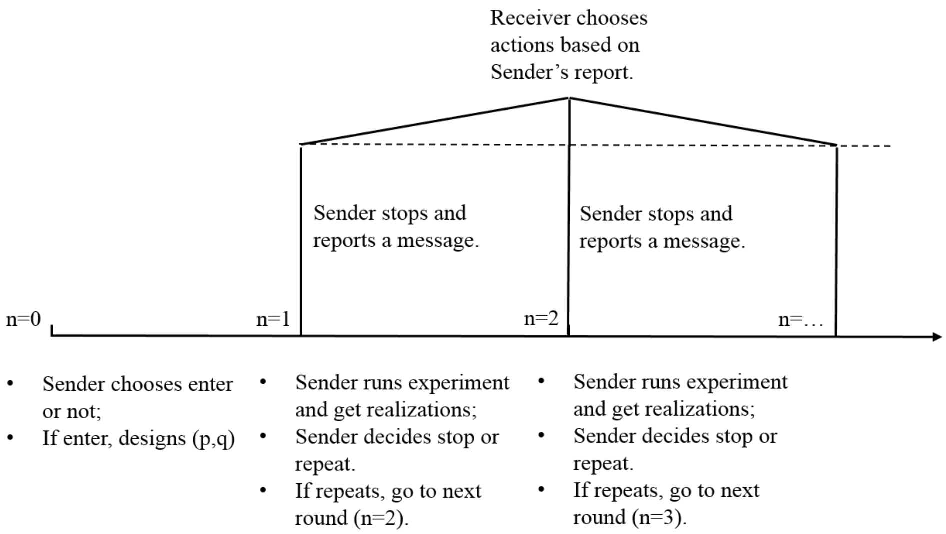

Figure 1: Time-line

2.2 A Preliminary Result: Optimal experiment when $\mathbf{c = 0}$

We start from a special case such that the cost of running experiment is 0. This case sheds some light on the key intuition of this persuasion game.

Proposition 1. *If $c = 0$, there are multiple optimal experiments: $\pi^* = (p, 1)$ for all $p \in (0, 1]$. Sender's highest ex-ante payoff is $W^* = \mu_0 Q$.*

Proof. First, we verify that all experiments with $q < 1$ is not optimal for Sender by contradiction. Suppose Sender designs an experiment π with $q < 1$ and Receiver takes $a = 1$ with some message m. Note that m must contain at least one good signal. Given Receiver's strategy, Sender's stopping rule is obvious: he keeps taking experiments until m is a subset of the history of his outcome, regardless of the state. In particular, even if he is sure that the state is B, he still has an incentive to repeat because state B can generate good signal and there is no cost to repeat. Thus, this message m is not informative given Sender's stopping rule, therefore, Receiver will deviate to $a = 0$ when she observes m. This contradicts with the assumption on Receiver's strategy is in equilibrium.

Secondly, we verify that $\pi^* = (p, 1)$, for all $p > 0$, are optimal experiments. $q = 1$ implies that only state G can generate good signal. Thus, as long as Receiver observes a good signal, she knows the true state is G and she takes action $a = 1$. Then, if Sender's belief of state G is positive, he will keep taking experiments until he gets one good signal since there is no cost. With $\pi^* = (p, 1)$, no matter how many bad signal Sender obtains, his posterior belief will not be zero. Thus, from an ex-ante point of view, even if p is extremely small, Sender will obtain one good signal for sure conditional on state G. After Sender obtains one good signal, he discloses it to Receiver and Receiver's posterior belief jumps to 1. Back to the first stage, Sender's ex-ante payoff $W(\pi^*) = \mu_0 Q$ for all possible π^*. □

Though there are infinitely many experiments which are optimal to Sender, they are all the same to Receiver, in terms of informativeness: If the state is G, Sender will provide one good signal anyway given his stopping rule. If the state is B, he will never be able to obtain one good signal. Thus, all the optimal experiments are equivalent to a perfect signal structure (i.e. $p = 1$ and $q = 1$) to Receiver. Therefore, when cost is zero, Receiver can make decisions with the true state revealed.

In general, Sender and Receiver share aligned interests when the true state is G, but they have opposite interests under state B. Thus, as long as Sender commits to an experiment in which the good signal can only be produced in state G, he aligns his interest with Receiver in the second stage of the game, regardless of the true state.

Thus, in equilibrium, Sender is indeed playing a full revelation experiment when he has the least credibility (cost is zero). Besides, it is interesting to note that Sender can always commit to a perfect signal $\pi = (1, 1)$ and he can always credibly convince Receiver the true state is G given one realization of good signal, even if the cost is positive. Thus, if we don't consider the cost, the gain of the optimal experiment is always (weakly) higher than $\mu_0 Q$, which is smaller than the payoff in KG, but higher than that in LR.

2.3 Characterization of Optimal Experiment Design

Now we start to analyze the whole game. Since we restrict the signal structure to binary structure, we can reduce the complicated equilibrium notions in the model setting to a simpler one. Recall that given Sender's experiment design at the first stage, the solution concept for the second stage's subgame is Perfect Bayesian Equilibrium: given Sender's choice of $\pi = (p, q)$ at the first stage, Receiver decides her accepting strategy which can be reduced to a cutoff strategy k: if the message contains at least $k \in \mathbb{N}^+$ good signals, she takes $a = 1$, otherwise she takes $a = 0$. Sender's best response to Receiver's cutoff k consists of two things: (1) his stopping rule $s(p, q, k)$, which is a collection of maps from possible histories to his stopping decision. (2) his reporting rule: when Sender stops after the round n of experiments, if H_n contains k good signals, he reports k good signals. Otherwise reports nothing[3]. Furthermore, to make the second stage's subgame a PBE, the belief system should be consistent: Receiver's posterior belief is updated over μ_0 and given Sender's best response to k, and indeed her posterior belief should exceed μ^* if being provided k good signals. Then by backward induction, Sender chooses his optimal design π^* in the first stage. Therefore, for simplicity, we denote the tuple (p, q, k) as a notion of an equilibrium in the subgame given (p, q), and this tuple is sufficient to represent all players' strategies in the second stage. Sender's ex-ante payoff at the first stage is $W(p, q|k, s(p, q, k))$ for any feasible experiment (p, q). And we denote $(p^*, q^*) = \arg\max_{(p,q)} W(p, q|k, s(p, q, k))$ as the set of Sender's optimal experiment(s). Thus, the Sender Preferred Subgame Perfect Bayesian Equilibrium can be denoted as (p^*, q^*, k^*).

Definition: Supposing Receiver takes $a = 1$ if the message contains at least k good signals, then an experiment (p, q) is ***credible*** for k if Receiver updates her posterior belief $\mu_m > \mu_0$ after observing a message m containing at least k good signals.

Define $\hat{q}_k$ as the threshold such that Sender is indifferent between stopping and re-

[3] Actually, Sender can report any subset of H_n to Receiver. But they are all the same to Sender given Receiver's strategy. Thus, we polish the message space and simply let Sender reports nothing.

peating to get k good signals, when sender knows that the true state is B:

$$\frac{k}{1-\hat{q}_k}c = Q \Rightarrow \hat{q}_k = 1 - kc/Q \tag{5}$$

Where $\frac{k}{1-\hat{q}_k}$ is the expected round Sender would take to get k good signals conditional on the state is B. So the left hand side is the expected cost while the right hand side is the gain if Receiver takes $a = 1$ with k good signals.

Lemma 1. *Any tuple (p, q, k) with $q < \hat{q}_k$ is not a PBE of the subgame in the second stage.*

Proof. For any fixed k, if $q < 1 - \frac{kc}{Q}$, $\frac{k}{1-\hat{q}_k}c < Q$. Therefore even if Sender knows that true state is B, he will keep repeating till he obtains k good signals. Thus, Sender will disclose k good signals with probability one, regardless of true state. Therefore, the experiments (p, q) with $q < \hat{q}_k$ are not *credible* for k, in particular uninformative to Receiver. Taking action if being provided k good signals is not Receiver's best response. □

For simplicity of exposition, from now on, we relaxed Receiver's incentive constraints: we assume that for a given k, as long as (p, q) is *credible* for this k, then Receiver take $a = 1$ if the message contains k good signals. It is worth to note that even though we relaxed Receiver's incentive constraints, but we restrict the discussion in the following within the *credible* experiments, in particular within the (p, q, k) space excluding those described in Lemma 1. Lemma 2 below further narrows down the set that we search for the optimal experiment for Sender.

Lemma 2. *Any experiment design (p, q) is not optimal for Sender if $p \neq 1$.*

Proof. By proving inequality (6), we prove Lemma 2.

$$W(1, q|k, s(1, q, k)) \geq W(1, q|k, s(p, q, k)) > W(p, q|k, s(p, q, k)) \quad \forall k > 0,\ q \geq \hat{q}_k \tag{6}$$

The first inequality is because that $s(1, q, k)$ is the optimal stopping rule given $(1, q, k)$. Thus any other stopping rule given by a different p leads to a weakly lower payoff. The second inequality is slightly tricky. Given the same stopping rule $s(p, q, k)$, if the underlying state is B, then the expected payoff under $(1, q, k)$ and (p, q, k) are the same, since only q and k matters. If the underlying state is G, then under $(1, q, k)$, there is only one path that leads to stopping: a sequence of k good signals, which is also the shortest path to make Receiver take action. This gives a sure return of $W(1, q, k) = Q - kc$, which is the highest payoff Sender can achieve given Receiver's strategy. If $p < 1$, there would be several different paths that lead to stopping. However, each path gives a lower payoff (comparing to $Q - kc$) as long as Sender is not lucky enough to have continuous k good signals in sequence. Thus $p = 1$ is optimal for Sender. A formal proof is shown in the appendix. □

The intuition for Lemma 2 is that Sender and Receiver share aligned interest when true state is G. Therefore Sender can increase the frequency of obtaining good signal in state G as much as he wants, i.e. p=1. Then conditional on state G, Receiver can make the right decision for sure. The next step is to find the optimal q given $p = 1$. Before doing that, we need to specify Sender's stopping rule under $p = 1$.

Lemma 3. *(Sender's stopping rule given Receiver's strategy k)*
Supposing Receiver takes $a = 1$ if the message contains k good signals and $p = 1$:
(1) With $q \in [0, \hat{q}_k)$, Sender stops only if he obtains k good signals;
(2) With $q \in [\hat{q}_{k-n+1}, \hat{q}_{k-n})$, $n = 1, 2...k$, if Sender does not observe n good signals at the first n round, he stops; otherwise he continues until he obtains k good signals.

Proof. It is important to notice that when $p = 1$, the first time Sender obtains a bad signal, the posterior belief jumps to 0. Thus by Lemma 1, when $q \in [0, \hat{q}_k)$, even after Sender observes a bad signal at the first experiment, he will repeat the experiments until he collects enough good signals. When $q \in [\hat{q}_{k-n+1}, \hat{q}_{k-n})$, $\frac{k-n}{1-q}c \leq Q < \frac{k-n+1}{1-q}c$. On the one hand, if Sender does not obtain n good signals at the beginning n rounds, the best case is having $n - 1$ good signals. However, even in this case, Sender's posterior is 0 and the expected cost of obtains the additional $k - n + 1$ will exceed the benefit Q. Thus he will choose to optimally stop. On the other hand, if he is lucky enough to obtain n good signals firstly, the expected cost of getting additional $k - n$ good signals can be compensated by the benefit Q. Thus he will stop only if he successfully collects those required signals and the additional $k - n$ good signals are actually uninformative. □

Lemma 4. *For any k, any experiments $(1, q)$ is not optimal for Sender if $q > \hat{q}_k$.*

Proof. We prove Lemma 4 by inequality (7).

$$W(1, \hat{q}_k, k|s(1, \hat{q}_k, k)) \geq W(1, q, k|s(1, q, k)), \ \forall q \in [\hat{q}_k, 1], k \geq 1 \tag{7}$$

In state G, any experiment with $p = 1$ will lead to the same payoff for Sender, since he will obtain k good signals continuously. Therefore, Sender's payoff only defers among experiments with different q when the true state is B. In state B, by Lemma 3, if Sender takes the first experiment and gets one good signal, he will continue till he obtains k good signals, otherwise, he stops. Therefore, the chance of gaining Q is $1 - \hat{q}_k$. If he carries out experiments other than this, i.e. $q \in [\hat{q}_{k-n+1}, \hat{q}_{k-n})$, his ex-ante payoff will be lower: conditional on his stopping rule, to finally convince Receiver taking $a = 1$, he has to obtains n good signal at the first n round, which happens with lower chance. Furthermore, if he is not lucky enough to obtain n good signals continuously, he has to pay the sunk cost for (weakly) more than one round. But when $q = \hat{q}_k$, he only need to pay a sunk cost of the first round if he fails to have one good signal. In the appendix, we prove this formally. □

With Lemma 4, we narrow down sender's potential set of optimal experiments to $(1, \hat{q}_k)$ given each $k \in \mathbb{N}^+$. In Lemma 5, we compare Sender's ex-ante payoff given his optimal experiment under each k and we find that $(1, \hat{q}_1)$ yields the highest ex-ante payoff for Sender.

Lemma 5. *Under the assumption that Receiver's incentive constraint being relaxed, the optimal experiment design for Sender is* $(1, \hat{q}_1)$.

$$W(1, \hat{q}_1, 1|s(1, \hat{q}_1, 1)) > W(1, \hat{q}_k, k|s(1, \hat{q}_k, k)), \ \forall k \in \mathbb{N}^+ \tag{8}$$

Conditional on $p = 1$, higher cost makes Sender more reluctant to repeat, which restores Sender's commitment power. In particular, higher cost gives Sender a larger set of credible q for each k. Hence Sender may have an incentive to strategically choose an experiment which is credible for higher k. But he also needs to afford a higher expected cost in this case. The Sender's ex-ante payoff under $(1, \hat{q}_k, k)$ is shown below:

$$\begin{aligned} W(1, \hat{q}_k, k|s(1, \hat{q}_k, k)) &= \mu_0(Q - kc) + (1-\mu_0)[(1-\hat{q}_k)(Q - \frac{k-1}{1-\hat{q}_k}c - c) + \hat{q}_k(-c)] \\ &= \underbrace{[\mu_0 + (1-\mu_0)kc/Q]Q}_{gain} - \underbrace{kc}_{cost} \\ &= \mu_0(Q - kc) \end{aligned} \tag{9}$$

The net gain is decreasing in k, and strictly positive when $k = 1$. The second line of equation (9) characterizes the gain and cost of each equilibrium. It is easy to see that an additional requirement of good signal (increase k by one) can increase Sender's gain by enlarging the set of credible q, but the increment of gain is only $(1 - \mu_0) < 1$. In the other word, the gain from restoring credibility cannot compensate for the loss of sustaining this credibility. Hence when $k = 1$, $(1, \hat{q}_1)$ leads to Sender's highest ex-ante payoff.

From now on, we restore Receiver's incentive constraint: Receiver will not take action if her posterior belief μ_m given Sender's stopping rule is below μ^*. The optimal design we find in Lemma 5 may fail this constraint. Proposition 2 below characterizes Sender's optimal experiment design in equilibrium:

Proposition 2. ***(Optimal experiment)***
If $\mu^* \leq \frac{\mu_0}{c/Q}$, *there exists an optimal experiment:*
(1) If $\mu^* \in (\mu_0, \frac{\mu_0}{\mu_0 + (1-\mu_0)c/Q}]$, *the optimal experiment design is* $(1, 1 - \frac{c}{Q})$;
(2) If $\mu^* \in (\frac{\mu_0}{\mu_0 + (1-\mu_0)c/Q}, \frac{\mu_0}{c/Q}]$, *the optimal experiment is* $(1, \frac{\mu^* - \mu_0}{\mu^*(1-\mu_0)})$;
(3) The game ends in one round.
If $\mu^* > \frac{\mu_0}{c/Q}$, *Sender won't enter the market.*

Proof. The first statement is directly from Lemma 4. This case is shown in Figure 2(a). $\frac{\mu_0}{\mu_0+(1-\mu_0)c/Q}$ is the posterior belief of Receiver when she receives one good signal given $(1, \hat{q}_1)$. If $\frac{\mu_0}{\mu_0+(1-\mu_0)c/Q} \geq \mu^*$, then Receiver will indeed take $a = 1$. Therefore $(p^* = 1, q^* = \hat{q}_1, k^* = 1)$ is the equilibrium in the whole game. Sender's ex-ante payoff is:

$$W(1, \hat{q}_1|(1, \hat{q}_1, 1)) = \underbrace{\mu_0 Q + (1-\mu_0)c}_{gain} - \underbrace{c}_{cost} = \underbrace{\mu_0(Q-c)}_{net\ gain} \tag{10}$$

The gain from Sender's preferred experiment is smaller than the KG, since the experiment pinned down by KG can not be credibly carried out.

As for the second statement, shown in Figure 2(b), $(1, \hat{q}_1, 1)$ now is not an equilibrium path since Receiver's posterior belief given one good signal is smaller than her threshold of acceptance, μ^*. Sender's most preferred path $(1, \hat{q}_1, 1)$ is not informative enough to persuade Receiver to take action, as shown in Figure 2(b). In this case, Sender will increase q and make Receiver just indifferent between accepting or not. The experiment is $(1, \frac{\mu^*-\mu_0}{\mu^*(1-\mu_0)})$, Note that this is exactly the same experiment as in KG, which leads to the upper bound of Sender's ex-ante payoff. It is important to note that $\frac{\mu^*-\mu_0}{\mu^*(1-\mu_0)} > \hat{q}_1$. Hence, by Sender's stopping rule given $k = 1$, if in the first experiment Sender obtains a good signal, he reports this good signal, while if he observes a bad signal in the first round, he stops and the game ends. In other words, Sender can use one experiment to achieve the highest gain attainable with full commitment power, this experiment design is therefore optimal. Sender's ex-ante payoff will be:

$$\underbrace{\frac{\mu_0}{\mu^*}Q - c}_{net\ gain} = \underbrace{\frac{\mu_0}{\mu^*}Q}_{gain} - \underbrace{c}_{cost} \tag{11}$$

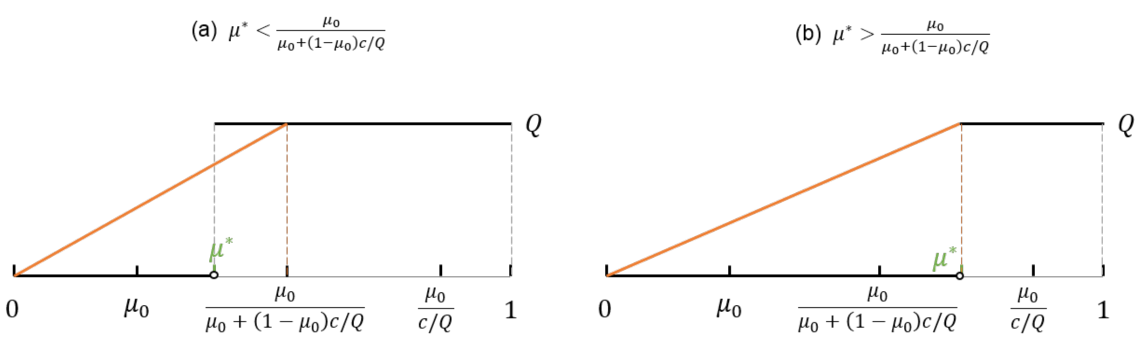

Figure 2: optimal experiment design

Lastly, Receiver's standard of acceptance cannot be extremely high. If $\mu^* > \frac{\mu_0}{c/Q}$, then the highest net gain from experiment $(1, \frac{\mu^*-\mu_0}{\mu^*(1-\mu_0)})$ is negative, Sender will not choose to enter the market. □

The optimal experiment here has a very simple structure. If the state is G, $p = 1$ implies that the experiment provides a good signal for sure. If the state is B, either $q = \hat{q}_1$ which implies that Sender is just indifferent between stopping and repeating in the hope of getting one good signal or $q = \frac{\mu^* - \mu_0}{\mu^*(1-\mu_0)}$ which implies that Receiver is just indifferent between her two actions. Under both cases, Sender will stop as long as the first experiment fails to generate a good signal and Sender learns that the true state is B. Thus, the experiment Sender implements in equilibrium (from ex-ante point of view) is indeed the experiment Sender proposes to Receiver at the very beginning. In other words, even if we assume Sender has the ability to repeat, he actually designs a more informative experiment as a commitment device to restore his credibility. This is because whenever Sender prefers to manipulate, Receiver will anticipate this incentive, which will re-adjust Sender's gain.

2.4 Comparative Statistics

Proposition 3. ***(comparative statistics)***
(1) Informativeness of experiment is weakly decreasing in c and increases in Q.
(2) Ex-ante payoff: both Sender's and Receiver's ex-ante payoff are decreasing in c and increasing in Q.

Proof. We use the same notion with Blackwell (1953) order on distributions of beliefs. Denote $\tau \in \Delta(\Delta\Omega)$ as distribution of posterior beliefs in equilibrium. Then if τ is a mean-preserving spread of τ', we say τ is more informative than τ', denoted as $\tau \succ \tau'$.

Figure 3 (a) characterizes the optimal experiment design over belief space when $\frac{c}{Q} < \frac{\mu_0(1-\mu^*)}{\mu^*(1-\mu_0)}$ (just rearrange $\mu^* \leq \frac{\mu_0}{\mu_0 + (1-\mu_0)c/Q}$). If $c > c'$, Sender is more reluctant to repeat so that he can credibly conduct a pair of posterior beliefs that is closer to KG. Figure 3 (b) shows the CDFs over beliefs. With a higher c, Sender is able to successfully persuade Receiver more frequently, which implies a less informative experiment: $\tau(c'/Q) \succ \tau(c/Q)$. Thus, the informativeness of the optimal experiment is decreasing in c. With the same logic, if Q is higher, Sender is less credible in the sense that he is more likely to manipulate. Hence, he should provide a more informative experiment design in order to convince Receiver that he will not manipulate the result.

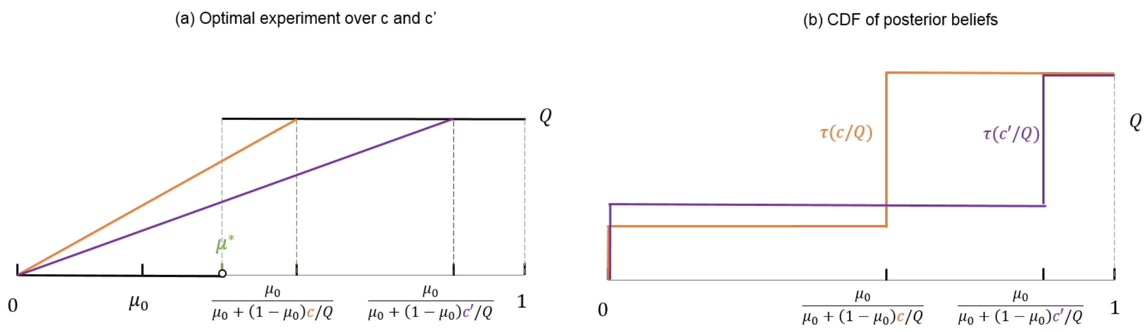

Figure 3: Informativeness

When $\frac{c}{Q} \geq \frac{\mu_0(1-\mu^*)}{\mu^*(1-\mu_0)}$, Sender's cost c (benefit Q) is large (small) enough so that he can credibly achieve the experiment design in KG, in which the signal structure is only related to Receiver's cutoff of taking action. Hence the informativeness won't change due to the cost increment. Given Receiver prefers more informative signals, she will be better off if c (Q) is smaller (larger).

By equation (10) and (11), no matter what the optimal experiment is, the net gain (ex-ante payoff) of Sender is always decreasing in c and increasing in Q. $\square$

3 Extension: General Payoff Function

The optimal experiment with binary step payoff function has very a simple structure. In this section, we relax this strong assumption and aim to discuss the optimal experiment under a general payoff function of Sender, denoted as $v(\mu)$, where μ is the posterior belief of Receiver. For now, we focus on experiment with binary signal structure (p, q) and the case $k = 1$, in which Receiver will take action when she is given one good signal in equilibrium.

Proposition 4. *If $v(\mu)$ ($v(0)$ is normalized to be 0) satisfies A1 and A2, then Sender will either enter the market and design a one-shot experiment with $p = 1$ or leave with $v_0 := v(\mu_0)$.*

A1. $v(\mu) \leq \overline{v}(\mu)$ when $\mu \leq \mu_0$, where $\overline{v}(\mu) := \frac{v_0}{\mu_0}\mu$

A2. $v(\mu)$ is increasing and bounded from above when $\mu > \mu_0$.

The first assumption basically says that $v(\mu)$ to the left side of prior is bounded from above by a linear payoff function. If this is not true, for example, $v(\mu)$ is extremely concave such that $v'(0) \to \infty$. Thus, compared to the optimal signal structure in the last section, Sender can always gain by reducing p a little bit and increasing q at the same time so that he can credibly implement this experiment only once. Even though he suffers a small loss of chance of success (obtain a good signal), he gains a large first order benefit from the little increment of posterior belief given a bad signal. Hence $p = 1$ won't be optimal.

For the second assumption, if in equilibrium Sender probably performs the experiment more than once, his posterior belief is different from Receiver's. Receiver knows that Sender might generate the good signal at different rounds, and Receiver takes expectation over all the possibilities. Then Receiver's posterior belief μ_m can be lower than the case that Sender commits to conduct the experiment only once. If $v(\mu)$ is non-monotonic when $\mu > \mu_0$, Sender might have incentive to design a multi-round experiment so to lower the posterior belief of Receiver when she sees a good signal. A numerical example is $c = \mu_0 = 0.3$. $v(\mu) = 0$ when $\mu \in [0, 0.3]$; $v(\mu) = 2$ when $\mu \in (0.3, \frac{10}{17})$ and $v(\mu) = 1$ when $\mu \in [\frac{10}{17}, 1]$. The optimal experiment design for one round is $(p = 1, q = 0.7)$. However, $(p = 0.9, q = 0.88)$ gives a higher payoff to the Sender and the experiment might be run twice. This is because at $(p = 0.9, q = 0.88)$, upon receiving a good signal, the posterior belief of Receiver lies in the region of $v(\mu) = 2$, which is of Sender's interest. A formal exposition of this example is in the Appendix.

To prove this proposition, we construct a two-step proof. In the first step, we use strategy-based approach to analyze the optimal experiment in an auxiliary game; In the second step, we use belief-based approach to analyze the experiment design in the original game with the results we get in the auxiliary game. Note that for all experiment (p, q) Sender proposes to Receiver, there exists an equivalent pair of posterior beliefs (μ_F, μ_S) that pins down the equilibrium experiment Sender is playing. Later on, we will switch between these two notations as needed[4]. If the experiment in equilibrium is the same as the one Sender proposes to Receiver, we will just denote it as (μ_b, μ_g).

3.1 First Step: Solution for the Auxiliary Game

Definition of Auxiliary game: Sender's payoff in the auxiliary game satisfies A1 and $v(\mu) = V$ when $\mu > \mu_0$.

Given Receiver's strategy such that she takes action with one good signal, Sender's stopping rule is obvious: for any (p, q) he proposes, there exists a maximum rounds[5]

[4] To avoid abusing notations, whenever we use beliefs to illustrate the experiment design, we will not omit (μ_F, μ_S).

[5] If there does not exist a N, which means that Sender will never stop even he gets infinite number of bad signals. Thus the evidence he provides is not informative and Receiver will not update her

N such that if he obtains a sequence of N bad signals, he will stop and report no good signal, since his posterior belief will be too low for him to repeat and intend for a good signal next round. Denote the posterior belief of Sender after a series of bad signal as $\mu_1^s(b), \mu_2^s(b), \ldots, \mu_N^s(b)$. The formula for posterior belief after n bad signals, $\mu_n^s(b)$ is:

$$\mu_n^s(b) = \frac{\mu_0(1-p)^n}{\mu_0(1-p)^n + (1-\mu_0)q^n} \tag{12}$$

Hence we can view this game in the following process: Sender proposes a tuple (p, q, N) to the receiver, where N is the maximum round of trials he will carry out. Sender can commit to (p, q) but not N. Thus, (p, q) must satisfy all Sender's incentive compatible constraints such that N is indeed the maximum round of experiments that he will play if he keeps getting bad signals.

Fix N, from an ex-ante point of view, there are only two events for Sender: successfully provide a good signal (event S) within N rounds, or fail to provide a good signal (event F). The ex-ante probability of the first event is denoted as $1 - B_N$, while the ex-ante probability of the second event is B_N:

$$B_N = \mu_0(1-p)^N + (1-\mu_0)q^N \tag{13}$$

If Sender fails to provide a good signal, he obtains N bad signals in state G with probability $(1-p)^N$ while he the chance of failure in state B is q^N. Besides, in state G, the chance of successfully providing one good signal is $1-(1-p)^N$ and the chance is $1 - q^N$ in state B. Therefore, the posterior beliefs are the following:

$$\mu_F = \frac{\mu_0(1-p)^N}{\mu_0(1-p)^N + (1-\mu_0)q^N} \tag{14}$$

$$\mu_S = \frac{\mu_0(1-(1-p)^N)}{\mu_0(1-(1-p)^N) + (1-\mu_0)(1-q^N)} \tag{15}$$

μ_S and μ_F are the posterior beliefs of Receiver. They are also the ex-ante posterior beliefs of Sender.

Denote $E_N(N, p, q)$ as the expected round that Sender will play given the maximum round and the experiment he proposes.

$$\begin{aligned} E_N(N,p,q) &= \mu_0[\sum_{n=1}^{N} n(1-p)^{n-1}p + N(1-p)^N] + (1-\mu_0)[\sum_{n=1}^{N} nq^{n-1}(1-q) + Nq^N] \\ &= \mu_0\frac{1-(1-p)^N}{p} + (1-\mu_0)\frac{1-q^N}{1-q} \end{aligned} \tag{16}$$

This is actually a weighted average of the manipulation between the game in equilibrium and the game Sender proposes to Receiver: with the game Sender proposes to Receiver, the probability to have a good signal in state G is p. However, for each N,

beliefs. This cannot be an optimal signal structure for Sender.

Sender is actually playing a game where the probability of obtaining a good signal is $1-(1-p)^N$ in state G. Thus, the ratio of $\frac{1-(1-p)^N}{p}$ is the expected round Sender will play in state G, which also measures the manipulation Sender implements (this ratio is increasing in N). In the meanwhile, the ratio $\frac{1-q^N}{1-q}$ measures the same thing in state B. The formal proof of this part is in the appendix. Hence, Sender's optimization problem is as following:

$$\max_{\{p,q,N\}} W(p,q,N) = (1-B_N)V + B_N v(\mu_F) - E_N(N,p,q)c \tag{17}$$

with IC constraints:

$$\left[\mu^s_{N-1}(b)p + (1-\mu^s_{N-1}(b))(1-q)\right](V - v(\mu_F)) \geq c \tag{18}$$

$$\left[\mu^s_{N}(b)p + (1-\mu^s_{N}(b))(1-q)\right](V - v(\mu_F)) \leq c \tag{19}$$

The first constraint (18) implies that Sender will repeat if he obtains $N-1$ bad signals, while the second constraint (19) implies that he will indeed stop if he obtains continuously N bad signals. These two constraints together imply N is the maximum round that Sender will play with the experiment design he proposes. It is easy to see that whenever (19) binds, (18) always hold, which is because the posterior belief with one less bad signal is higher and this leads to a higher expected chance of getting good signal in the next round. Thus, we can discuss the relaxed original optimization problem only with (19), and then check if the optimal solution of this relaxed problem locates in the region of (18).

Lemma 6. *If $V > \frac{v_0}{\mu_0} + c$, then the optimal experiment design of the auxiliary game is $(p=1, q=1-\frac{c}{Q})$ and the game ends in one round. If $V < \frac{v_0}{\mu_0} + c$, then Sender does not enter the market. If $V = \frac{v_0}{\mu_0} + c$, entering the market and not entering deliver the same payoff and both are optimal to the sender.*

Instead of discussing the relaxed original optimization problem, we construct a new optimization problem, denoted as Problem $\hat{P}$.

The new objective function:

$$\hat{W}(p,q,N) = (1-B_N)V + B_N \overline{v}(\mu_F) - E_N(N,p,q)c \geq W(p,q,N) \tag{20}$$

and the new IC, denote as IC':

$$(1-q)(V - \overline{v}(\mu_F)) \leq c \tag{21}$$

Note that when the experiment design is $(1, 1-\frac{c}{V})$, we have $\hat{W}(1, 1-\frac{c}{V}, N) = W(1, 1-\frac{c}{V}, N) = \mu_0(V-c)$ and both IC (19) and IC' (21) are binding. Besides, we substitute the upper boundary $\overline{v}(\mu)$ into the original objective function (17). Thus,

$\hat{W}(p,q,N)$ is higher than $W(p,q,N)$. If this higher new $\hat{W}(p,q,N) \leq W(1, 1-\frac{c}{V}, 1) = \mu_0(V-c)$ (given constraint (19)), then the optimal solution is $(1, 1-\frac{c}{V})$. As for the new constraint (21), we first substitute the upper boundary $\overline{v}(\mu)$, which makes the LHS of (19) lower. Since $p \geq 1-q$, we do one more scale down as $\mu^s_{N-1}(b)p + (1-\mu^s_{N-1}(b))(1-q) \geq 1-q$. Now that the LHS of new IC' is lower than the original one. More specifically, for any given p, IC' binding leads to a lower q. In other word, the region of the new IC' contains the original one. Therefore, if this higher objective function (20) is still smaller than $W(1, 1-\frac{c}{V}, 1)$ everywhere within the set of IC', the original constrained optimization problem cannot have a solution better than $(1, 1-\frac{c}{V})$. We formally show this is true when $V \geq \frac{v_0}{\mu_0} + c$ in the appendix, which is the constraint that implementing $(1, 1-\frac{c}{V})$ is better than not entering the market. If $V < \frac{v_0}{\mu_0} + c$, it is better for Sender to leave with v_0.

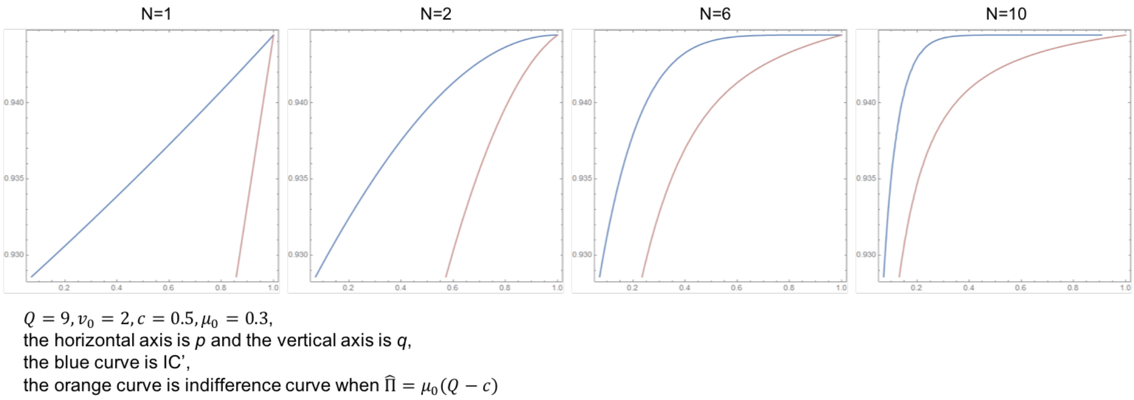

Figure 4: Simulation when $V \geq \frac{v_0}{\mu_0} + c$

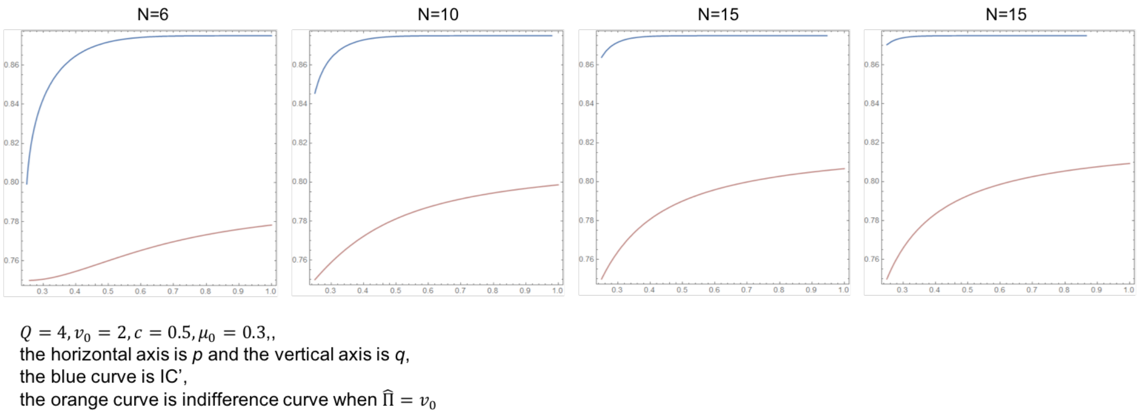

Figure 5: Simulation when $V < \frac{v_0}{\mu_0} + c$

Figure 4 and 5 illustrate the simulation results in both cases. The blue curve is the binding IC' and the red curve is Sender's indifference curve. The region above the blue curve is the case such that IC' is slack, while the region below the indifference curve gives a higher utility level. When $V \geq \frac{v_0}{\mu_0} + c$, the indifference curve is at utility level of $\mu_0(V - c)$ since it crosses the point $(p = 1, q = 1 - \frac{c}{V})$. As the indifference curve does not intersect with the new incentive constraint IC' for every possible N. Therefore, $(p = 1, q = 1 - \frac{c}{V})$ is the optimal design. In Figure 5, the ex-ante payoff Sender can achieve with credible signal structure (region above the blue curve) is even worse than not entering the market. Thus, Sender will choose not to enter the market. The formal proof is in the appendix.

3.2 Optimal Experiment Design

In this section, we transfer to belief-based approach to construct the optimal experiment design of any general payoff function with A1 and A2, using the results of the auxiliary game. To do this, we first need to define **credibility frontier** $V^{cf}(\mu_g)$, which is the maximum payoff at each $\mu_g > \mu_0$ such that the experiment of $(\mu_b = 0, \mu_g)$ is credible for $N = 1$ (on $V^{cf}(\mu_g)$, with a bad signal in the first round, Sender is indifferent between stopping and continuing for a good signal).

$$V^{cf}(\mu_g) = \frac{(1 - \mu_0)\mu_g c}{\mu_0(1 - \mu_g)} \tag{22}$$

$V^{cf}(\mu_g)$ is globally convex and increasing, and $V^{cf}(\mu_g)|_{\mu_g \to 1} = \infty$. Intuitively, when μ_g is higher, the chance of getting a good signal in state B is lower. Thus even with a higher V, Sender can still be indifferent between stopping and continuing. In the meanwhile, if $\mu_g = 1$, the experiment is full revelation. No matter how high the gain from good signal is, Sender will not continue the experiments when getting a bad signal since he will never get a good signal the next time. Thus, $v(\mu)$ must cross V^{cf} at least once, since $v(\mu)$ is bounded from above implies $v(1) < V^{cf}(1)$. Denote all the crossing points as $C_1, C_2, ..., C_N$, and the posterior beliefs associated to those points as $\mu_{C_1}, \mu_{C_2}, ..., \mu_{C_N}$, where C_N is the highest crossing point.

Define a set of $\mathcal{C}$, as following:

$$\mathcal{C} := \{\mu_g | v(\mu_g) \geq \frac{v_0}{\mu_0} + c \ and \ \mu_g \geq \mu_{C_N}\} \tag{23}$$

This set $\mathcal{C}$ is always non-empty since there exists a $C_N < 1$ for sure.

Proposition 5. *For any payoff function with A1 and A2, (1) If $v(1) < \frac{v_0}{\mu_0} + c$, Sender will leave the market with v_0. (2) If $v(1) \geq \frac{v_0}{\mu_0} + c$, the highest payoff Sender can achieve through experiment is $\hat{V}(\mu_0) - c$ while the optimal experiment is pinned down by the concave envelope $\hat{V}(\mu)$ (defined in KG) over the belief space $\mathcal{B} = 0 \cup \mathcal{C}$. If $\hat{V}(\mu_0) - c > v_0$, then Sender enters the market and choose the optimal design. Otherwise Sender leaves the market with v_0.*

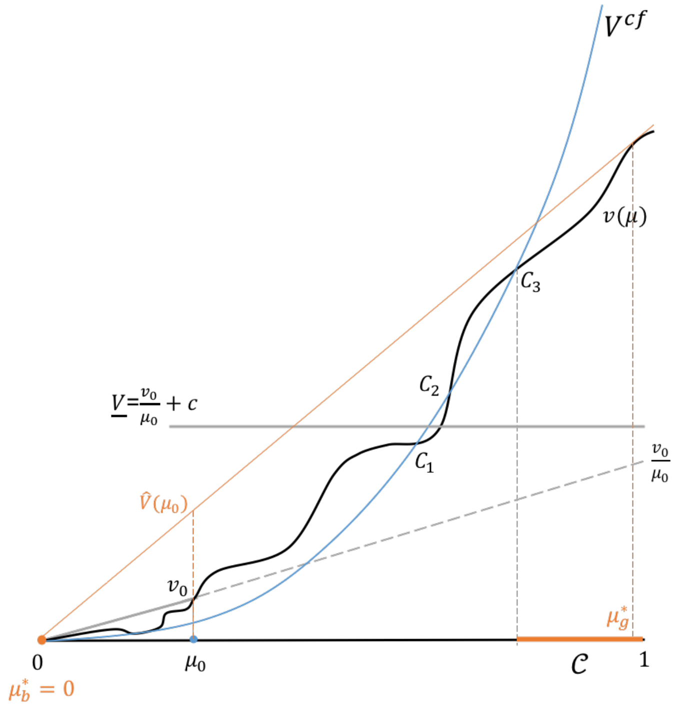

Figure 6: optimal experiment design with general payoff function

Figure 6 gives an example of a general payoff function of Sender satisfying A1 and A2, shown as the black curve. The blue curve is the credibility frontier V^{cf} which divides the original payoff function into several parts. In this example, $v(1)$ is higher than $\frac{v_0}{\mu_0}+c$ and the refined belief space $\mathcal{C}$ is highlighted as an orange set. By Proposition 5, the optimal pair of posterior beliefs are located within the belief space $\mathcal{C}$. Drawing the concave envelop over this refined belief space, we pin down the optimal pair of posterior beliefs (μ_b^*, μ_g^*). Furthermore, we can derive the optimal experiment (p^*, q^*) directly from (μ_b^*, μ_g^*) since in the signal structure Sender proposes to Receiver is exactly the same with the structure he will play in equilibrium. In the following part, we formally prove Proposition 5.

Proof. Step 1. Any experiment design with μ_g such that $v(\mu_g) < \frac{v_0}{\mu_0}+c$ is not worth trying. For $v(\mu_g) < \frac{v_0}{\mu_0}+c$, suppose Sender will enter the market and design an experiment with (μ_b, μ_g, N), then the gain from good signal is $v(\mu_g)$. From Lemma 6, we know that $W(\mu_b, \mu_g, N) \equiv W(\mu_b, \mu_g, N|\textit{auxliary game with } Q = v(\mu_g)) < v_0$. Thus, we refine the belief space of μ_g first time such that $v(\mu_g) \geq \frac{v_0}{\mu_0}+c$. And if the highest payoff $v(1) < \frac{v_0}{\mu_0}+c$, Sender will not enter the market.

Step 2. For $v(\mu_g) \geq \frac{v_0}{\mu_0}+c$, any design with μ_g such that $v(\mu_g) > V^{cf}(\mu_g)$ is not optimal.

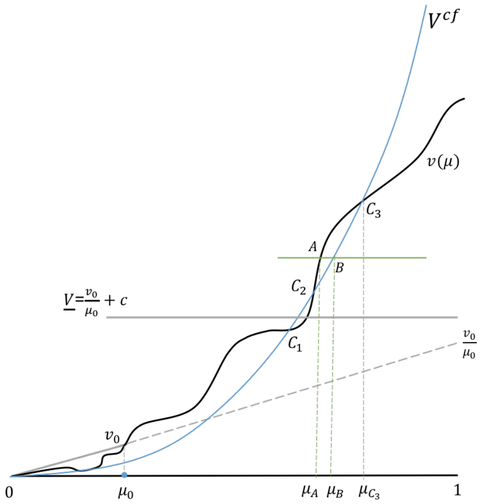

Figure 7: graph illustration of step 2

Suppose that optimal design is some $(\mu_b, \mu_g = \mu_A)$ such that $v(\mu_A) > V^{cf}(\mu_A)$, then we can draw a horizontal line crossing A and intersecting $V^{cf}(\mu)$ at B, shown in Figure 7. It is important to note that $W(\mu_b, \mu_A, N) < W(0, \mu_B, 1) < W(0, \mu_{C_3}, 1)$. The first inequality comes from Lemma 6. The second inequality comes from the convexity of V^{cf} and monotonicity of $v(\mu)$. Thus, we refine the set of posterior beliefs of μ_g again such that $v(\mu_g) \leq V^{cf}(\mu_g)$.

Step 3. The optimal design is pinned down by the concave envelope over the belief set, $\mathcal{B} = 0 \cup \mathcal{C}$.

$$
\begin{aligned}
W(0, \mu_g^*, 1) &\geq W(0, \mu_{C_N}, 1) \\
&> W(\mu_b, \mu_g, 1 | \mu_b \in [0, \mu_0), \mu_g < \mu_{C_N} \; and \; v(\mu_g) < V^{cf}(\mu_g)) \\
&> W(\mu_F, \mu_S, N > 1 | \mu_F \in [0, \mu_0), \mu_S < \mu_{C_N} \; and \; v(\mu_S) < V^{cf}(\mu_S))
\end{aligned} \tag{24}
$$

The first inequality comes from the fact of concavification. The second inequality comes from A1, A2 and V^{cf} being convex. The last inequality comes from the fact that for each fixed pair of equilibrium posterior beliefs, one-shot experiment is the cheapest. □

Now we also close the proof for Proposition 4, since $\mu_b^* = 0$ is equivalent to $p = 1$. It is interesting that the optimal experiment is pinned down by the same way as KG but within a refined subset of the belief space. If the refined belief space is large enough to contain the optimal pair of beliefs in KG, Sender can achieve the highest gain (even if he does not have a full commitment power) and pay a one shot cost c. Otherwise, he cannot achieve the experiment design with full commitment.

3.3 Comparative Statistics

Proposition 6. *If c is higher, the optimal experiment of any general payoff function with A1 and A2 is weakly less informative.*

It is easy to show that V^{cf} is increasing in c. Thus, the set of $\mathcal{B}$ is larger. Intuitively, higher cost gives a higher credibility, which leads to a larger refined belief set. Figure 8 shows the optimal pair of posterior beliefs with c and c' separately, where $c > c'$. As $\mathcal{B}(c') \subset \mathcal{B}(c)$, $\mu_g^*(c)$ must be weakly smaller than $\mu_g^*(c')$. If this is not true, Sender should have chosen experiment with $\mu_g^*(c)$ when the cost is c' and $\mu_g^*(c')$ is not optimal. The distribution of posterior beliefs $\tau(c')$ is a mean preserving spread of $\tau(c)$. Therefore, $(\mu_b = 0, \mu_g^*(c'))$ is more informative.

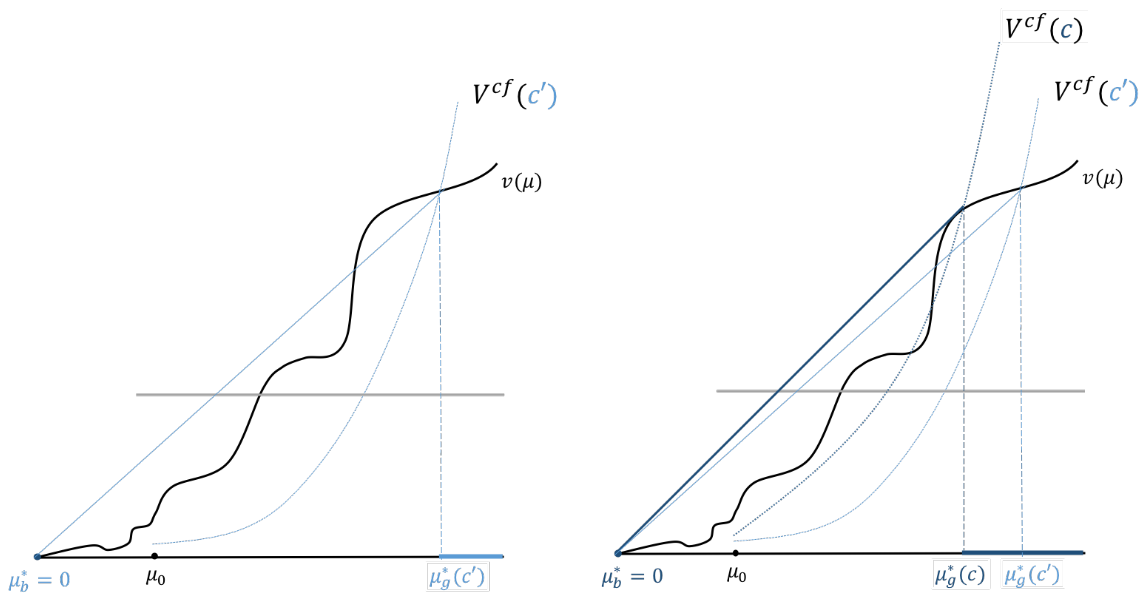

Figure 8: Optimal experiment with different costs

However, the impact of cost on Sender's ex-ante payoff is ambiguous since the increment of gain due to the change of experiment is not comparable with the increment of cost under the assumption of general payoff function.

4 Conclusion

The binary state space gives us tractability of this path dependent information acquisition problem. We find some interesting results that Sender restricts himself to end the game in one round when he is allowed to acquire information dynamically, and also he restricts his manipulation power when he is allowed to use it. We think it would be interesting to dig more on the credibility frontier which may give us a more general geometric illustration of optimal information disclosure policy.

Reference

Patrick Bolton, Xavier Freixas, and Joel Shapiro. The credit ratings game. The Journal of Finance, 67(1):85–111, 2012.

Dion Bongaerts, KJ Martijn Cremers, and William N Goetzmann. Tiebreaker: Certification and multiple credit ratings. The Journal of Finance, 67(1):113–152, 2012.

Abel Brodeur, Mathias Lé, Marc Sangnier, and Yanos Zylberberg. Star wars: The empirics strike back. American Economic Journal: Applied Economics, 8(1):1–32, 2016.

Archishman Chakraborty and Rick Harbaugh. Persuasion by cheap talk. American Economic Review, 100(5):2361–82, 2010.

Alfredo Di Tillio, Marco Ottaviani, and Peter Norman Sørensen. Persuasion bias in science: can economics help? The Economic Journal, 127(605):F266–F304, 2017.

Emeric Henry and Marco Ottaviani. Research and approval process: The organization of persuasion. working paper, 2018.

Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. American Economic Review, 101(6):2590–2615, 2011.

Elliot Lipnowski and Doron Ravid. Cheap talk with transparent motives. working paper, 2018.

Elliot Lipnowski, Doron Ravid, and Denis Shishkin. Persuasion via weak institutions. working paper, 2018.

Campbell R.Harvey. Presidential address: The scientific outlook in financial economics. The Journal of Finance, 72(4):1399–1440, 2017.

Francesco Sangiorgi, Jonathan Sokobin, and Chester Spatt. Credit-rating shopping, selection and the equilibrium structure of ratings. Technical report, Working Paper, Stockholm School of Economics and Carnegie Mellon University, 2009.

Vasiliki Skreta and Laura Veldkamp. Ratings shopping and asset complexity: A theory of ratings inflation. Journal of Monetary Economics, 56(5):678–695, 2009.

Appendix

Proof of Inequality (7)

Proof. Firstly, the expected payoff function under $(1, \hat{q}_k, k)$:

$$\begin{aligned} W(1, \hat{q}_k, k|s(1, \hat{q}_k, k)) &= \mu_0(Q - kc) + (1-\mu_0)[(1-\hat{q}_k)(Q - \frac{k-1}{1-\hat{q}_k}c - c) + \hat{q}_k(-c)] \\ &= \mu_0 Q - kc + (1-\mu_0)(1-\hat{q}_k)Q \\ &= \mu_0(Q - kc) \end{aligned} \tag{25}$$

Next, for any given k and $2 \leq n \leq k$, and $\forall q \in [\hat{q}_{k-n+1}, \hat{q}_{k-n})$, the expected payoff function for the sender with $(1, q, k)$ is

$$\begin{aligned} W(1, q, k|s(1, q, k)) &= \mu_0(Q - kc) + (1-\mu_0)\left[-c\sum_{s=0}^{n-1}(1-q)^s + (1-q)^n(Q - \frac{k-n}{1-q}c)\right] \\ &= \mu_0(Q - kc) + (1-\mu_0)\left[-c\sum_{s=0}^{n-2}(1-q)^s + (1-q)^n(Q - \frac{k-n+1}{1-q}c)\right] \\ &\leq \mu_0(Q - kc) + (1-\mu_0)\left[-c\sum_{s=0}^{n-2}(1-q)^s + (1-q)^n(Q - \frac{k-n+1}{1-\hat{q}_{k-n+1}}c)\right] \\ &= \mu_0(Q - kc) + (1-\mu_0)\left[-c\sum_{s=0}^{n-2}(1-q)^s\right] \\ &< \mu_0(Q - kc) \end{aligned} \tag{26}$$

Also, for any k, $n = 1$ and $q \in [\hat{q}_k, \hat{q}_{k-1})$

$$\begin{aligned} W(1, q, k|s(1, q, k)) &= \mu_0(Q - kc) + (1-\mu_0)[(1-q)(Q - \frac{k-1}{1-q}c - c) + q(-c)] \\ &= \mu_0(Q - kc) + (1-\mu_0)[(1-q)(Q - c) - qc - (k-1)c] \\ &= \mu_0(Q - kc) + (1-\mu_0)[(1-q)Q - kc] \\ &\leq \mu_0(Q - kc) + (1-\mu_0)[(1-\hat{q}_k)Q - kc] \\ &= \mu_0(Q - kc) \end{aligned} \tag{27}$$

Hence (7) holds. □

Example when A2 of Proposition 4 is Violated

A numerical example is as following $c = \mu_0 = 0.3$. $v(\mu) = 0$ when $\mu \in [0, 0.3]$; $v(\mu) = 2$ when $\mu \in (0.3, \frac{10}{17})$ and $v(\mu) = 1$ when $\mu \in [\frac{10}{17}, 1]$.

When Sender wants to design an experiment such that the experiment run at most one time, by Proposition 2, the optimal experiment is $(p = 1, q = 0.7)$, where it leads to $(\mu_b = 0, \mu_g = \frac{10}{17})$ and the ex-ante payoff for Sender is 0.21.

We argue that $(p = 0.9, q = 0.88)$ gives strictly higher ex-ante payoff for Sender, and hence when A2 is violated, Proposition 4 can be wrong. Firstly, we show that under $(p = 0.9, q = 0.88)$, there exists an equilibrium such that Sender will do the experiment for at most 2 times, and Receiver's posterior belief conditional on receiving good signal lies in $v(\mu) = 2$ region. We have to prove four constraints: Receiver's posterior belief lies in $v(\mu) = 2$ area; Sender has incentive to do a second experiment when fail once; Sender has no incentive to do a third experiment when fail twice; ex-ante payoff is higher than that of $(p = 1, q = 0.7)$.

Constraint 1. By (15), the posterior belief of receiver under $(p = 0.9, q = 0.88, N = 2)$ is $\mu_S = 0.541 \in (0.3, \frac{10}{17})$.

Constraint 2. After one bad signal, the posterior belief of Sender is determined by (12): $\mu_1^s(b) = \frac{\mu_0(1-p)}{\mu_0(1-p)+(1-\mu_0)q}$. This equals to 0.0464. Then run the same experiment again, the probability to get a good signal is 0.1562. This is greater than $0.3/2 = 0.15$. Thus, Sender wants to run a second experiment to get a higher expected payoff.

Constraint 3. After two fails, the posterior of Sender drops to 0.0055. And expected probability to get a good signal is $0.1243 < 0.15$. Hence Sender will stop and not carrying out round 3 after two fails.

Constraint 4. Equation (17) determines the objective function. At $(p = 0.9, q = 0.88, N = 2)$, the expected payoff is $0.443 > 0.21$.

Proof of Lemma 6

The new optimization problem $\hat{P}$:

$$\begin{aligned} \hat{W}(p,q,N) =& [\mu_0(1-p)^N + (1-\mu_0)q^N] \frac{v_0(1-p)^N}{\mu_0(1-p)^N + (1-\mu_0)q^N} \\ & + (1 - \mu_0(1-p)^N - (1-\mu_0)q^N)V \\ & - \left[\mu_0 \frac{1-(1-p)^N}{p} + (1-\mu_0)\frac{1-q^N}{1-q}\right] c \end{aligned} \tag{28}$$

The incentive constraint IC':

$$(1-q)\left(V - \frac{v_0(1-p)^N}{\mu_0(1-p)^N + (1-\mu_0)q^N}\right) \le c \tag{29}$$

Now we prove Lemma 6 by contradiction.

Case 1: $V \ge \frac{v_0}{\mu_0} + c$

Suppose there exists a (p, q, N) satisfies IC' such that $\hat{W}(p, q, N) > \mu_0(V - c)$.

Begin with:

$$\hat{W}(p,q,N) - \mu_0(V-c) > 0 \tag{30}$$

Plug $-\frac{c}{1-q} \leq \frac{v_0(1-p)^N}{\mu_0(1-p)^N+(1-\mu_0)q^N} - V$ into (30), the LHS of (30) is larger, then we have:

$$-\mu_0(1-p)^N V + \frac{v_0(1-p)^N(1-\mu_0+\mu_0(1-p)^N)}{\mu_0(1-p)^N+(1-\mu_0)q^N} + \mu_0 c - \mu_0 c\frac{1-(1-p)^N}{p} > 0 \tag{31}$$

Plug $V \geq \frac{v_0}{\mu_0} + c$ into (31), we have:

$$f(p,q,N) \equiv \frac{v_0(1-\mu_0)(1-p)^N(1-q^N)}{\mu_0(1-p)^N+(1-\mu_0)q^N} - \mu_0 c(1-p)\frac{1-(1-p)^N}{p} > 0 \tag{32}$$

Notice that we can transform $f(p,q,N)$ to:

$$f(p,q,N) = [1-(1-p)^N]\left[v_0(1-\mu_0)\frac{1-q^N}{1-(1-p)^N}\frac{(1-p)^N}{\mu_0(1-p)^N+(1-\mu_0)q^N} - \mu_0 c\frac{1-p}{p}\right] \tag{33}$$

Notice that the first bracket is always positive and the second bracket is decreasing in N. Then $f(p,q,1) > 0$ is a necessary condition for $f(p,q,N) > 0$. Look back to IC' (29): if (p,q,N) satisfies IC', $(p,q,1)$ must be satisfying IC'. Thus, we only need to discuss the case of $N=1$. $f(p,q,1) > 0$ implies:

$$\frac{v_0(1-\mu_0)(1-q)}{\mu_0(1-p)+(1-\mu_0)q} > \mu_0 c \tag{34}$$

And at $(p,q,1)$, IC' holds implies:

$$(1-q)(V - \frac{v_0(1-p)}{\mu_0(1-p)+(1-\mu_0)q}) \leq c \tag{35}$$

Since $V \geq \frac{v_0}{\mu_0} + c$, (35) becomes:

$$\frac{v_0(1-\mu_0)(1-q)}{\mu_0(1-p)+(1-\mu_0)q} \leq \mu_0 c \tag{36}$$

Inequality (34) and (36) cannot hold together implies that within the credible space of $(p,q|N)$, there does not exist an experiment in equilibrium such that the ex-ante payoff of Sender exceeds $W^*(1, 1-\frac{c}{V}, 1) = \mu_0(V-c)$. Besides, the signal structure with $p=1, q=1-\frac{c}{V}$ is located within the original incentive constraints (18) and (19). Thus this structure guarantees Sender the highest ex-ante payoff.

Case 2: $V \leq \frac{v_0}{\mu_0} + c$

Similar to the method in last subsection, Suppose there exists a (p,q,N) satisfying IC' such that $\hat{W}(p,q,N) - v_0 \geq 0$.

Subtract $(1-\mu_0)\frac{1-q^N}{1-q}$ times IC' (29) from $\hat{W}(p,q,N) \geq v_0$, and we have:

$$0 \leq \frac{v_0(1-\mu_0)(1-q^N)(1-p)^N}{\mu_0(1-p)^N + (1-\mu_0)q^N} + \left(\mu_0 V - v_0 - \frac{\mu_0 c}{p}\right)(1-(1-p)^N) \equiv g(p,q,N) \tag{37}$$

$g(p,q,N)$ can be transformed into the following:

$$g(p,q,N) = (1-(1-p)^N)\left[\frac{v_0(1-\mu_0)(1-q^N)(1-p)^N}{(\mu_0(1-p)^N + (1-\mu_0)q^N)(1-(1-p)^N)} + (\mu_0 V - v_0 - \frac{\mu_0 c}{p})\right] \tag{38}$$

With similar analysis as above, it is easy to see that $g(p,q,N) > 0$ is more slack when N decreases. Thus we only need to analyze the $N=1$ case. $g(p,q,1) > 0$ implies:

$$\frac{v_0(1-\mu_0)(1-q)(1-p)}{\mu_0(1-p) + (1-\mu_0)q} + \mu_0 pV - pv_0 - \mu_0 c > 0 \tag{39}$$

Next, IC' can be transformed into the following:

$$(1-q)\frac{v_0}{\mu_0} - \frac{v_0(1-q)(1-p)}{\mu_0(1-p) + (1-\mu_0)q} + (1-q)V - (1-q)\frac{v_0}{\mu_0} - c \leq 0 \tag{40}$$

Which simplifies to:

$$\frac{v_0(1-\mu_0)(1-q)q}{\mu_0(1-p) + (1-\mu_0)q} + \mu_0(1-q)V - (1-q)v_0 - \mu_0 c \leq 0 \tag{41}$$

First, if $V \leq \frac{v_0}{\mu_0}$, (39) and (41) cannot hold together. Second, when $V \in (\frac{v_0}{\mu_0}, \frac{v_0}{\mu_0} + c]$, we prove by contradiction. Suppose there exists a (p,q) such to make both condition holds. Then, for any $V' = V + \frac{k}{1-q}$, $c' = c + k$, where $k \geq 0$, LHS of (41) remains unchanged and LHS of (39) increases. Then (39) and (41) holds simultaneously as well under (V', c') and the same (p,q). Let $\Delta = \frac{v_0}{\mu_0} + c - V$. It is easy to verify that there exists $k = \frac{(1-q)\Delta}{q}$ such to make $V' = \frac{v_0}{\mu_0} + c'$. However, in the case of $V \geq \frac{v_0}{\mu_0} + c$, we already showed that the payoff is bounded by $\mu_0(V-c)$. Therefore, under (V', c'), the payoff should be bounded by $\mu_0(V' - c') = v_0$ for any signal structure. Hence (39) and (41) cannot hold altogether under (V', c'). In this way we construct a contradiction.

To sum up, we prove that when $V \leq \frac{v_0}{\mu_0} + c$ and $v(\mu)$ is upper bounded by $\frac{v_0 \mu}{\mu_0}$ to the left of μ_0, the optimal strategy for the sender is always not designing any experiment and just take v_0 away.