

Bureaucratic Norms and Dynamic Bayesian Persuasion

Yihang Zhou*

March, 2024

Abstract

A developer seeks to persuade a welfare-maximizing bureaucrat to award a contract to her. The official has a short tenure, and his successor's decisions are subject to a bureaucratic norm: an official's decision must be based on evidence that is either recorded by his predecessor or presented to him. Thus, Bayesian inference is *restricted* when the first official fails to record evidence. The first official can exploit this and induce the developer to conduct more informative experiments. I focus on parameter values where the static values of persuasion are zero to the bureaucracy and strictly positive for the developer. I show that there are two possibilities in the dynamic game. Either the developer conducts a more informative experiment and the official decides immediately, so that social welfare is greater. Or there is delay, where the cost of delay to the bureaucracy offsets the benefits of a more informed decision. In either case, the developer is worse off compared to static persuasion. With unrestricted inference, there exists an intuitive PBE that replicates the static outcome, so that the bureaucratic norm may increase social welfare and never reduces it.

Keywords: Dynamic Bayesian persuasion; Restricted Bayesian inference

*University of Texas at Austin, yhzhou@utexas.edu

I thank V. Bhaskar for helpful comments and his valuable input. I would also like to thank Simon Board, William Fuchs, Vasiliki Skreta, Caroline Thomas, Thomas Wiseman, and the audiences at Midwest Economics Theory Conference and Texas Theory Camp for their helpful comments.

Declarations of interest: none

1 Introduction

Government agencies and organizations are usually subjected to greater scrutiny and more external supervision than private corporations. For example, the director of the FBI participates in over 18 meetings yearly with oversight bodies, which are noted for their extended duration and complexity. State Department's reporting task is under close and often hostile scrutiny, which leads to a culture of caution (Wilson, 1989). Thus, governmental organizations have to answer more questions about their decisions and also face more constraints than private firms. These constraints pertain especially to the use of information. As Wilson (1989) notes in the context of public procurement, public officials may not, unlike private firm employees, incorporate their personal knowledge and past experience in determining their purchasing decisions. Instead, they are required to follow strict procedural guidelines, such as competitive bidding. Otherwise, they may be accused of collusion, favoritism, and unethical agreements. Assessing the impact of these constraints and scrutiny on governmental efficiency is crucial for understanding the unique challenges faced by public sector management.

This paper examines one implication of the scrutiny of the decisions of public servants. I assume that a government official must be able to justify his decision to a third party (e.g., his superiors), by citing the evidence that is presented, either directly to him or as noted by his predecessor. In particular, the official cannot cite the *absence of a record* as a justification for his decision. I model this situation as a dynamic model with restricted Bayesian inference and show that, in this case, the restriction on inference may be socially beneficial.

The most well-known example of a restriction on inference is in the context of *Miranda right*. The US Supreme Court decided, in *Miranda vs. Arizona*, that the Fifth Amendment guard against self-incrimination implies that the jury cannot draw a negative inference from the suspect's refusal to answer police questions. Seidmann (2005) analyzes the effects of *Miranda*, using a model of restricted Bayesian inference: at the information set where the suspect has exercised his Miranda rights, the jury's decision is based only on the prior and upon witness statements. At other information sets, the inference is unrestricted.¹

As a leading example, I present a model with two short-tenure bureaucrats (he) who are being lobbied by an interest group, such as a local developer (she). Both bureaucrats have identical objectives and seek to maximize (discounted) social welfare. They must decide whether to award the contract to the developer or an outsider, and prior beliefs favor the outsider.² The developer seeks to persuade the first bureaucrat by conducting a Bayesian experiment. After observing the experiment, bureaucrat 1 must decide whether to award the contract to the outsider, to the local developer, or to postpone the decision so that it is made by bureaucrat 2. If he defers the decision, he may choose not to record the results of the experiment that he has observed, in which case bureaucrat 2 must decide based only on the prior and upon the experiment that he personally observes.

My main finding is that in equilibrium, the bureaucratic norm that restricts Bayesian inference may increase social welfare, the objective of the organization. For most of the paper, I focus on a class of information design problems where the value of persuasion to the

¹His main finding is that *Miranda* right can be socially beneficial: it reduces the wrongful conviction of innocent suspects. It also decreases the conviction rate and keeps the confession rate unchanged. See also Seidmann and Stein (2000).

²An alternative application: the bureaucrats must decide the location of a local public good, such as a park, and the lobby group favors location A over location B.

bureaucrat is zero in the static game, but strictly positive for the developer. In the absence of a norm, there exists an equilibrium of the dynamic game where the outcome of the static equilibrium is replicated, so that there is no delay and the bureaucrat has a zero value of persuasion. However, the bureaucratic norm may force the developer to conduct more informative experiments; if the developer conducts an experiment where the first bureaucrat only marginally prefers to award the contract to the developer after some signal realization, then the bureaucrat will fail to record this and defer the decision to his successor, and the developer will be forced to provide more information to convince him. To avoid this, the developer may provide more information to bureaucrat 1, so he is, after a positive signal, enthusiastic enough about the developer that she avoids costly delay. The reader may ask, why does the developer not defer all information provision to the second period? In this case, bureaucrat 1 correctly anticipates that the social value of second-period information will be zero, and thus he chooses the outside contractor immediately since delay is costly.

The main results can be extended to a more general case with more than two actions. I assume that bureaucrats divide the project into several parts and decide how many parts to be assigned to the developer. The developer wants a larger share of the project. Bureaucrats want to assign the project entirely to the developer in one state and entirely to the outsider in the other, and have a quadratic loss payoff function. I also assume that the developer's payoff is concave in the shares assigned to her. These assumptions ensure that the bureaucrat has zero value of persuasion in the static game. We have a more complicated situation with multiple actions. In the binary-action case, fixing the initial prior, when the first bureaucrat delays and does not record the experiment, the optimal experiment in period 2 always induces the same actions. However, in the multiple-action case, the actions induced in period 2 will

change with the bureaucrat 1's posterior in period 1. Despite the complexity, we still have a similar result as the leading example: when players are patient enough, in equilibrium, two scenarios arise. The first bureaucrat has a positive value of persuasion and chooses only immediate actions. Alternatively, he has a zero value of persuasion and delays when receiving a positive signal.

The organization of this paper is as follows. Section 2 discusses the implications of the bureaucratic norm via a leading example with binary actions. Section 3 analyzes the general case with multiple actions. Section 4 is the benchmark where the bureaucrat is not restricted by the norm. Section 5 concludes.

1.1 Related literature

Kamenica and Gentzkow (2011) initiated the study of Bayesian persuasion. Dynamic models of persuasion include Ely (2017), Honryo (2018), Orlov, Skrzypacz and Zryumov (2020), Smolin (2020), Bizzotto, Rüdiger and Vigier (2021).

The bureaucratic norm implies that the second-period interaction between the developer and the bureaucrat is one where they effectively have different priors, and my paper builds on the analysis of Alonso and Cámara (2016), who study heterogeneous beliefs in Bayesian persuasion. The bureaucratic norm also says that the updating is not Bayesian when the information is hidden. Such non-Bayesian updating is also analyzed by Levy, de Barreda and Razin (2018), de Clippel and Zhang (2022), and Galperti (2019).

An alternative interpretation of the model is that bureaucrat 2 is naive in the sense of regarding the absence of communication as the absence of information. This interpretation

is related to the self-deception problem studied by Bénabou and Tirole (2002). However, in their model, the incentive to manipulate the information comes from time-inconsistent preferences, while I assume identical preferences for two bureaucrats.

2 Binary actions

2.1 Setup

As a leading example, I present a two-period dynamic information design model with binary actions. There are three players: one long-lived local developer and 2 short-lived bureaucrats (B_1 and B_2). The local developer seeks to persuade the bureaucrat to contract with her instead of an outsider via Bayesian experiments.

Bureaucrats have the decision set $\mathcal{A} = \{1, 0\}$: $a = 1$ stands for contracting with the developer, and $a = 0$ stands for contracting with the outsider.

Let $\omega \in \{1, 0\} := \Omega$ denote the unobserved payoff-relevant state. Here $\omega = 1$ is the state where choosing the developer is more socially beneficial; $\omega = 0$ is the state where choosing the outsider is more socially beneficial. I need two requirements on the bureaucrats' preference, (1) they prefer $a = 1$ at state $\omega = 1$ and $a = 0$ at state $\omega = 0$, and (2) their payoffs are always non-negative, so that delay is costly. I assume quadratic-loss social welfare, which is also the payoffs of bureaucrats:

$$u(a, \omega) = C - (a - \omega)^2,$$

where $C > \frac{1}{2}$, so that the expected payoff of a correct decision is strictly positive.

Both bureaucrats have the same objective, maximizing discounted social welfare, while the developer has a state-independent payoff: $v : \mathcal{A} \rightarrow \mathbb{R}$. The developer prefers $a = 1$, and I normalize the payoff from $a = 0$ to 0 and $a = 1$ to 1. Finally, players share a discount factor $\delta < 1$.

At the beginning of the game, players share the same prior $p_0 = Pr(\omega = 1)$, but in period 2, the priors may be different between B_2 and the developer, denoted as p_2 and p_2^d respectively. Moreover, to focus on the non-trivial case, I assume that the prior is in favor of the outsider ($p_0 \leq \frac{1}{2} - \epsilon$), where ϵ is a small positive number.³ The belief $Pr(\omega = 1) = \frac{1}{2}$ is the point where the bureaucrat is indifferent between $a = 1$ and 0.

In each period, the developer chooses an experiment: $\pi_t = (\pi_t(\cdot|\omega))_{\omega \in \Omega} \in \times_{\omega \in \Omega} \Delta(S) := \Pi$, where S is an unrestricted signal space. I assume that outcomes of experiments in different periods are independent conditional on the state (i.e. $Pr(s_1, s_2|\omega) = \pi_1(s_1|\omega) \cdot \pi_2(s_2|\omega)$). This is an important assumption – if I were to allow experiments to be correlated across periods, the developer could credibly disclose the outcome of the past experiment.

The assumption that the experiments presented to the two bureaucrats are independent conditional on the state is an important one and needs justification. One justification is as follows. Suppose that the developer has a facility that needs to be inspected by the bureaucrat in order to ascertain her suitability for fulfilling the contract. The developer may specify the length of the inspection, thereby determining the informativeness of the experiment, but the bureaucrat must choose a sample of aspects of the facility to inspect. It is plausible that the two bureaucrats independently select their samples, giving rise to

³If the initial belief p_0 is in favor of $a = 1$, the result is trivial: the developer provides an uninformative experiment.

experiments that are independent conditional on the state.

In period 1, observing the outcome of the experiment, beliefs are updated from p_0 to q_1 . Then B_1 chooses an action from $\{1, 0\} \cup \{hide, record\}$. If $a = 1$ or 0 is chosen, the game ends and payoffs are realized.

On the other hand, if *hide* or *record* is chosen, the game proceeds into the next period, and the current bureaucrat's payoff is decided by the action chosen by the future bureaucrat. By recording, B_1 records and discloses the experiment and the outcome in period 1 to B_2 , so B_2 has a prior $p_2 = q_1$. By hiding, no outcome or experiment is recorded, and B_2 observes nothing. I assume that B_2 is constrained by the bureaucratic norm that decisions can only depend on recorded evidence, so B_2 has the belief $p_2 = p_0$ observing no recorded evidence of the last experiment. The developer's prior in period 2 is $p_2^d = q_1$. Furthermore, in period 2, which is the deadline, the bureaucrat must choose from $\mathcal{A} = \{1, 0\}$.

The equilibrium concept in this paper is perfect Bayesian equilibrium: players maximize their expected payoffs given other players' strategies and the beliefs generated by the Bayes rule if possible. Notice that this PBE is with a restriction as in Seidmann (2005): if B_1 does not record the first experiment, B_2 can only base his decision on the prior and the experiment presented to him. Furthermore, as a tie-breaker, I assume that a bureaucrat will choose the action preferred by the developer if he is indifferent between two actions.

2.2 Benchmark

To look at the effect of the bureaucratic norm, I use the dynamic persuasion game without the norm as a benchmark. In this benchmark, B_2 can do Bayesian updating even when B_1

hides the information, and all other setups remain the same. I find that in this benchmark, we always have an equilibrium that replicates the outcome of the static persuasion. See Section 4 for the full discussion.

Proposition 0. *When B_2 is not restricted by the bureaucratic norm, there is an equilibrium where players get the same payoffs as under static persuasion.*

Before I go to the outcome of the static persuasion, I first define an important concept, *the value of persuasion* (see Kamenica and Gentzkow (2011)).

Definition 1. *The value of persuasion to a player is his expected payoff in equilibrium minus his expected payoff without any experiment.*

Denote B_1 's value of persuasion in the dynamic game as V_B and the developer's value of persuasion in the dynamic game as V_D .⁴ And denote players' values of persuasion in the static game as \bar{V}_B and \bar{V}_D . Based on this definition, I introduce the following useful lemma, which is the outcome of the static persuasion.

Lemma 1. *In static Bayesian persuasion with binary actions and a developer of state-independent preference, $\bar{V}_B = 0$ and $\bar{V}_D > 0$ when the prior is in favor of the outsider.*

Lemma 1 can be obtained from the results in Kamenica and Gentzkow (2011). Because my benchmark has the same outcome as the static persuasion, I will compare my results to the outcome in Lemma 1.

⁴Since in period 2, the game is the same as a static persuasion, I focus on B_1 's value of persuasion in the dynamic game.

2.3 Equilibrium results

The bureaucratic norm leads to different results from the benchmark, as stated in the following proposition.

Proposition 1. *When δ is close enough to 1, $\exists \alpha(\delta)$ and $\beta(\delta) \in (\alpha(\delta), \frac{1}{2})$ s.t. in the unique equilibrium:*

- (1) *if $p_0 \notin [\alpha(\delta), \beta(\delta)]$, $V_B > 0$ and the first bureaucrat always makes the decision;*
- (2) *if $p_0 \in [\alpha(\delta), \beta(\delta)]$, $V_B = 0$ and the first bureaucrat delays at the positive signal and chooses $a = 0$ at the negative signal;*
- (3) *V_D is strictly less than \bar{V}_D .*

If the developer conducts an experiment where the first bureaucrat only marginally prefers to award the contract to the developer after some signal realization, B_1 will defer the decision to B_2 . To avoid this costly delay, the developer may conduct a more informative experiment so that B_1 is enthusiastic enough about the developer after a positive signal and takes action immediately. B_1 strictly prefers $a = 1$ to $a = 0$ at this positive signal. In this case, the more informative experiment induces only immediate actions and gives B_1 a positive value of persuasion. As a result, the bureaucratic norm is socially beneficial here.

However, with other parameters, delay happens in equilibrium. The experiment in this case generates a positive signal and a negative signal, and B_1 delays at the positive signal. B_1 has a zero value of persuasion here.

This happens when the developer finds that providing a more informative experiment is so costly that it is better to provide a less informative experiment and let B_1 delay. With this less informative experiment, B_1 defers the decision to B_2 at the positive signal, and then

the developer provides a second experiment to B_2 in period 2. At the negative signal, B_1 immediately chooses $a = 0$. Though bureaucrats act on the basis of more information in this equilibrium, delay cost offsets the benefit. Thus, B_1 has $V_B = 0$ in this case, which is the same as the benchmark.

We can solve for these equilibrium results by first looking at period 2. In period 2, the situation is the same as a static Bayesian persuasion game with possibly heterogeneous priors. When priors are heterogeneous, the equilibrium can still be solved by the concave closure according to Alonso and Câmara (2016).

As for period 1, firstly we observe that *record* will not be chosen by B_1 .

Lemma 2. *B_1 will not choose record in equilibrium.*

When B_1 delays and records the information, he shares the same expected payoff as B_2 . However, Lemma 1 says that B_2 has zero value of persuasion in period 2 in this case, so the extra experiment in period 2 does not improve B_1 's payoff. Moreover, there is a discount factor δ , which makes an immediate action strictly better than *record* for B_1 .

As a result, to solve for B_1 's strategy, we only need to keep track of actions $a = 1$, $a = 0$, and $a = \textit{hide}$. The payoffs from $a = 1$ and $a = 0$ are $C - (1 - q_1)$ and $C - q_1$ respectively when the posterior is q_1 . If $a = \textit{hide}$ is chosen, the developer in period 2 chooses a posterior split between 0 and $\frac{1}{2}$ for B_2 .

I summarize B_1 's strategy under certain p_0 and q_1 in the following lemma (please see Appendix A.1 for proof).

Lemma 3. *For discount factor $\delta \in (0, 1)$, there exists $\hat{\delta}(p_0) \in (0, 1)$:*

(1) *If $\delta > \hat{\delta}(p_0)$, there exist cutoff points $\hat{\alpha}(\delta, p_0), \hat{\beta}(\delta, p_0)$ ($p_0 < \hat{\alpha}(\delta, p_0) < \hat{\beta}(\delta, p_0) < 1$)*

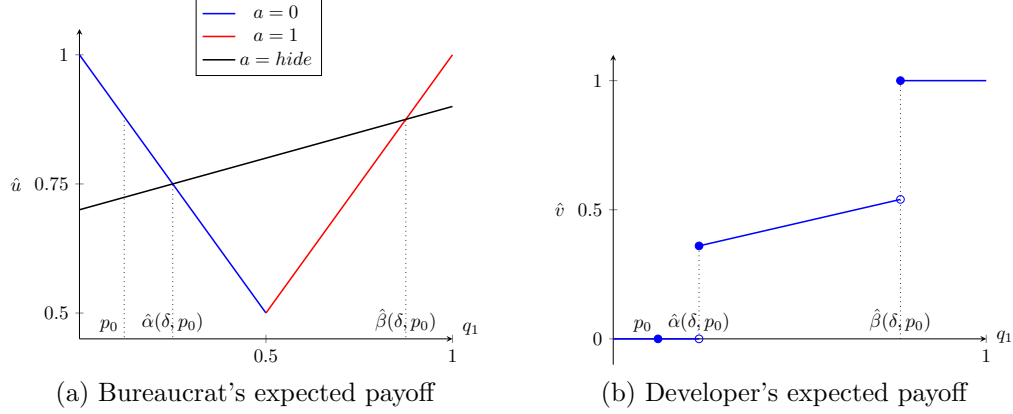


Figure 1: Period 1

such that B_1 chooses $a = 1$ immediately with $q_1 \geq \hat{\beta}_2(\delta, p_0)$; chooses $a = 0$ immediately with $q_1 < \hat{\alpha}(\delta, p_0)$; chooses *hide* with $\hat{\alpha}(\delta, p_0) \leq q_1 < \hat{\beta}_2(\delta, p_0)$.

(2) If $\delta \leq \hat{\delta}(p_0)$, B_1 always acts immediately.

For $\delta > \hat{\delta}(p_0)$, two players' expected payoffs (*y-axis*) given different posteriors after the first experiment (*x-axis*) can be summarized in Figure 1.⁵

Notice that Figure 1 is drawn for a specific prior p_0 . When p_0 changes, two graphs in Figure 1 change. More specifically, as p_0 goes up, the payoff line from *hide* in Figure 1 (a) moves downwards, making two cutoff points $\hat{\alpha}(\delta, p_0)$ and $\hat{\beta}_2(\delta, p_0)$ closer to each other.

From Figure 1 (b), it is easy to see that because we have $p_0 < \hat{\alpha}(\delta, p_0)$, the optimal experiment in period 1 has two possibilities: (1) the posterior split between 0 and $\hat{\alpha}(\delta, p_0)$ (*delaying experiment*), where B_1 chooses *hide* at $\hat{\alpha}(\delta, p_0)$; (2) the split between 0 and $\hat{\beta}_2(\delta, p_0)$ (*immediate experiment*), where B_1 takes action immediately at both posteriors. Which one is the optimal experiment will depend on the values of the initial priors and the discount factor, which is summarized in Proposition 1. Moreover, according to Figure 1 (a), *delaying experiment* gives $V_B = 0$ to B_1 , while *immediate experiment* gives $V_B > 0$ to B_1 . In other

⁵Graphs in this section are plotted with the parameter value $C = 1$, and serve the purpose of illustration.

words, when *immediate experiment* is the optimal experiment, the bureaucratic norm is beneficial to bureaucrats. However, this norm is not beneficial or harmful if *delaying experiment* is the optimal experiment.

Generically, the optimal experiment is unique — either *delaying experiment* or *immediate experiment*. The uniqueness is ensured when the origin and two cutoff points in Figure 1 (b) are not collinear. Only in the knife-edge cases where these three points are collinear, the optimal experiment is not unique. However, even in knife-edge cases, the equilibrium can still be unique with the help of the tie-breaker making the developer choose the less informative experiment when indifferent. A similar tie-breaker is also used in Bizzotto, Rüdiger and Vigier (2021).

When $\delta \leq \hat{\delta}(p_0)$, there are only immediate actions and the developer just chooses the static optimal experiment in period 1.

The discussion above shows how we get the two possibilities of the unique equilibrium in Proposition 1. For the conditions governing these possibilities, see the proof in Appendix A.2.

3 Multiple actions

The results in the example can be extended to a more general case, where bureaucrats have more than two actions.

Consider the situation where the government divides the project into N parts and decides on awarding how many parts to the local developer and the outsider. The decision set of bureaucrats is $\mathcal{A} := \{a_0, a_2, \dots, a_N\}$ now, where a_n means awarding $\frac{n}{N}$ of the project to the

developer and leaving the rest to the outsider.

The state space is still binary: $\Omega \in \{0, 1\}$. The state $\omega = 1$ means it is better to award the project to the developer, and $\omega = 0$ means it is better to award the project to the outsider.

I still assume the same quadratic-loss social welfare:

$$u(a, \omega) = C - (a - \omega)^2,$$

where $C > \frac{1}{2}$, so that the expected payoff of a correct decision is strictly positive, and thus the delay is costly for B_1 .

The bureaucrat's optimal action among the decision set \mathcal{A} is a_n if the belief $q = Pr(\omega = 1) \in [\frac{2n-1}{2N}, \frac{2n+1}{2N})$, $n = 1, 2, \dots, N-1$. a_0 is optimal with belief $q \in [0, \frac{1}{2N})$, and a_N is optimal with belief $q \in [\frac{2N-1}{2N}, 1]$. The bureaucrat is indifferent between two adjacent actions at these cutoff points. Denote indifferent cutoffs as $\bar{q}_n = \frac{2n-1}{2N}$, $n = 1, 2, \dots, N$, and $\bar{q}_0 = 0$, $\bar{q}_{N+1} = 1$.

The developer has a state-independent payoff: $v : \mathcal{A} \rightarrow \mathbb{R}$, and $v(a_i) < v(a_j)$, $\forall i < j$. She always prefers a larger share of the project. I normalize $v(a_N) = 1$. Moreover, the marginal benefit decreases in the share, i.e., $v(a_{n+1}) - v(a_n) < v(a_n) - v(a_{n-1})$. I also require $v(a_N) - v(a_{N-1}) < \frac{1}{2}(v(a_{N-1}) - v(a_{N-2}))$. With these assumptions, in the static game, $\bar{V}_B = 0$. Actually, $\bar{V}_B = 0$ in the static game is equivalent to that the static optimal experiment is inducing two adjacent actions, i.e., denote the initial prior as $p_0 = Pr(\omega = 1)$, if $p_0 \in [\bar{q}_{n-1}, \bar{q}_n)$, the static optimal experiment is the posterior split of \bar{q}_{n-1} and \bar{q}_n . The assumptions required here are to make the static optimal experiment always induce two adjacent actions, and thus $\bar{V}_B = 0$.

I assume that the prior p_0 is bounded away from indifferent points \bar{q}_n by a small positive

number ϵ , i.e., $p_0 \notin \mathcal{B}_\epsilon(\bar{q}_n)$, $n = 0, 1, \dots, N + 1$. Additionally, I assume that $p_0 < \bar{q}_N$.⁶

Finally, I apply tie-breakers that the bureaucrat chooses the action preferred by the developer when indifferent, and the developer in period 2 chooses the experiment that is better for the developer in period 1 when indifferent.

3.1 Equilibrium in the multiple-action case

The results in the binary-action case do not directly apply to the multiple-action case. When B_1 chooses *hide* at posterior $q \neq p_0$, the game in period 2 is a Bayesian persuasion with heterogeneous priors as in Alonso and Câmara (2016). In the binary-action case, no matter what q is, when B_1 chooses *hide*, the optimal experiment in period 2 is always the posterior split of 0 and $\frac{1}{2}$ for prior p_0 . As a result, B_1 's payoff from *hide* is linear in his posterior q .

However, in the multiple-action case, we have more than one indifferent cutoff point (\bar{q}_n). When $q > p_0$, after considering the heterogeneous priors in period 2, the point that is originally on the concave closure of the developer's payoff function can become inside the closure, which makes the optimal experiment in the case of heterogeneous priors different from the one with homogeneous priors (a more detailed argument is provided in Appendix A.3).

Despite the complexity of the multiple-action case, under our assumptions, we have similar results as the leading example when δ is large enough. In equilibrium, the developer is still worse off than static persuasion, and V_B can be zero or positive for B_1 .

⁶The developer will just choose an uninformative experiment when p_0 is equal to indifferent points or $p_0 \geq \bar{q}_N$.

Proposition 2. *When δ is close enough to 1, the unique equilibrium is either:*

- (1) $V_B > 0$, and the first bureaucrat always chooses immediate actions, or
- (2) $V_B = 0$, and the first bureaucrat delays at the positive signal of the experiment.
- (3) V_D is strictly less than \bar{V}_D

Please see Appendix A.3 for the proof.

3.2 Convex utility of the developer

In the binary-action case, when the prior favors the outsider, the static game invariably results in the optimal experiment being a posterior split between zero and the posterior making the bureaucrat indifferent between two actions. We always have $\bar{V}_B = 0$ in the static game.

However, in the multiple-action cases, the optimal static experiment can vary. For instance, consider the case with three actions, $\mathcal{A} = \{a_0, a_1, a_2\}$, and a prior p_0 that lies within the range $[0, \bar{q}_1)$. Depending on players' preferences, the optimal split can be zero and \bar{q}_1 , or zero and \bar{q}_2 . The former results in the bureaucrat's adjacent actions being induced, as depicted in Figure 2 (a), and here \bar{V}_B remains at zero. In contrast, the latter scenario, illustrated in Figure 2 (b), leads to the induction of extreme actions and is characterized by a positive \bar{V}_B .

As in Section 3, assuming a strictly concave utility function for the developer leads to an optimal static experiment characterized by inducing adjacent actions, as depicted in Figure 2 (a). On the other hand, a convex utility function for the developer yields an optimal static experiment that induces extreme actions, corresponding to Figure 2 (b).

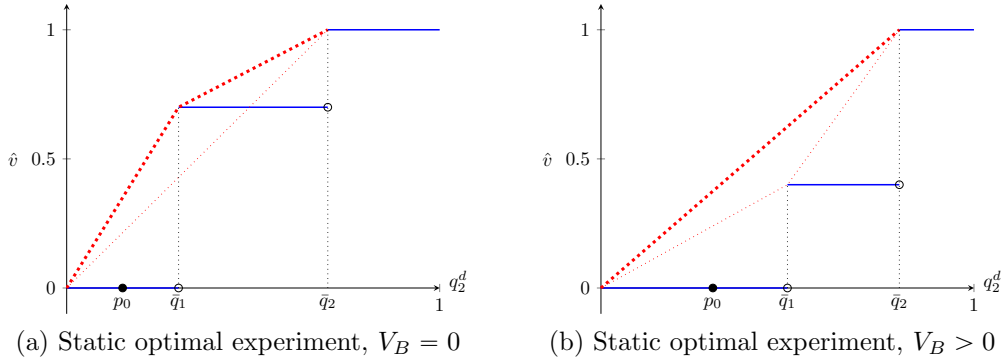


Figure 2: The developer's payoff in the static game

In the convex utility case, our analysis will change. To illustrate, consider a case with three possible actions for the bureaucrats ($\mathcal{A} = \{a_0, a_1, a_2\}$) and a convex utility for the developer, while the other assumptions from Section 3 remain in place.

I find that in this tri-action case, the first-period optimal experiment is uninformative when the discount factor δ is sufficiently large.

Proposition 3. *If $|\mathcal{A}| = 3$, when δ is close enough to 1, the developer chooses the uninformative experiment in period 1.*

In this convex utility case, B_1 finds it beneficial to postpone action and record the evidence when he is patient enough, because the value of persuasion in period 2 is positive here. This allows the developer to postpone all information disclosure to period 2. Different from the concave utility case, we see that B_1 may opt to delay even in the absence of period 1 disclosure.

Another difference from the concave utility case is that the equilibrium analysis must now account for the *record* action. Furthermore, *record* will happen in equilibrium when δ is smaller. For instance, consider the case where the prior belief p_0 falls within the range between \bar{q}_1 and \bar{q}_2 . For certain discount factors δ , the payoffs for B_1 and the developer,

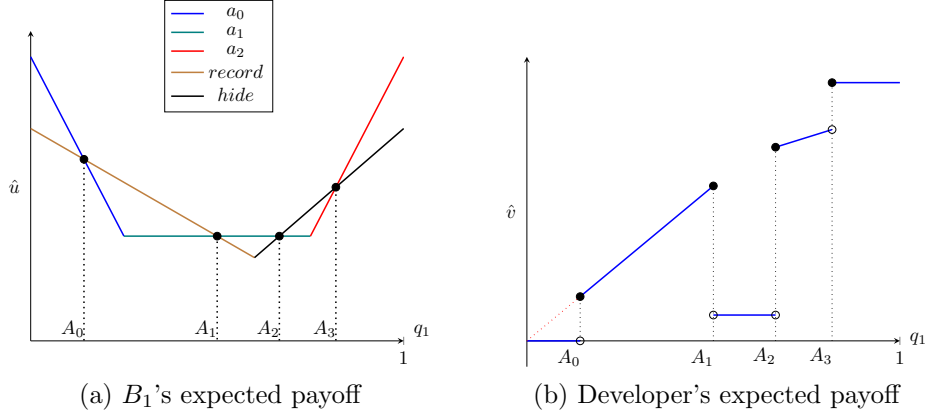


Figure 3: Players' payoffs with lower δ

varying with q_1 , are depicted in Figure 3. Some important cutoff points in the graph are denoted as A_0 , A_1 , A_2 , and A_3 . Since p_0 falls between the cutoff points A_1 and A_2 , by the concave closure method in Kamenica and Gentzkow (2011), there are several possible configurations that the unique optimal experiment could take, such as splits between A_1 and A_2 , and between A_1 and A_3 . Notice that B_1 chooses *record* at A_1 , *hide* at A_2 , and a_2 at A_3 . Consequently, the equilibrium can lead to outcomes where *record* is paired with either *hide* or a_2 .

Notice that with other parameter values, the unique optimal experiment in period 1 has other possibilities besides the ones above. Specifically, the optimal experiment may yield outcomes similar to those in the concave utility case: (1) it induces only immediate actions, and (2) it results in an immediate action at one signal and *hide* at the other signal.

4 Unrestricted bureaucrats

In this section, I look at bureaucrats not restricted by the bureaucratic norm, which is the benchmark. Specifically, B_2 will do Bayesian inference observing no information disclosure

from period 1, given other players' strategies in equilibrium. To align with the previous scenario where the developer cannot verifiably disclose past information to B_2 , I assume that if B_1 delays and hides at multiple posteriors on the equilibrium path, the developer chooses the same experiment for those posteriors in period 2 and maximizes the ex-ante expected payoff. Furthermore, I regard experiments as hard information, which means if B_1 chooses *record*, even when it is off the equilibrium path, B_2 will update according to the experiment outcome by Bayes rule.

In the game with an unrestricted B_2 , as stated in Proposition 0, we always have an equilibrium where the developer conducts a static optimal experiment in period 1, and B_1 acts immediately at both signals. Within this equilibrium, B_1 's action *hide* is off the equilibrium path, so we can assign any off-path belief to B_2 to support the equilibrium. Given the presence of a static optimal experiment in the first period and the immediate actions by B_1 , the resulting payoffs are the same as static persuasion. The following is a formal proof for Proposition 0.

Proof for Proposition 0. In such equilibrium, the developer conducts the static optimal experiment in period 1, which is the posterior split of \bar{q}_n and \bar{q}_{n+1} for the prior $p_0 \in [\bar{q}_n, \bar{q}_{n+1})$. Then B_1 chooses a_n at posterior \bar{q}_n and a_{n+1} at posterior \bar{q}_{n+1} . The off-path belief of B_2 when B_1 chooses *hide* is \bar{q}_{n+1} .

For B_1 with the posterior \bar{q}_n or \bar{q}_{n+1} , according to the arguments in Lemma 4, he will not choose *hide*. When B_1 chooses *record*, for the same reason as Lemma 2, it is worse than choosing an immediate action. Thus, B_1 will not deviate.

For the developer, because B_1 will choose immediate actions at posteriors \bar{q}_n and \bar{q}_{n+1} , there is no experiment better than the static optimal experiment due to the concavity of the

developer's utility. Thus, the developer does not deviate. ■

Furthermore, the equilibrium in Proposition 0 also survives a refinement in the spirit of the D1 criterion (henceforth D1 criterion). For this refinement, I assume that B_1 is infinitely more likely to deviate than the developer. Under this assumption, when we consider the off-path belief where B_1 chooses *hide*, we can regard B_1 's actions after the on-path experiment of the developer as his signaling and apply D1 criterion. For example, when the on-path experiment gives posteriors q and q' , I regard the game after this on-path experiment as a signaling game where B_1 has types q and q' .

Suppose after the on-path experiment in period 1, B_1 has a posterior q . I define B_1 's equilibrium payoff at this on-path posterior as $U^*(q)$.

Our objective is to determine B_2 's belief following B_1 's non-equilibrium action *hide* using D1 criterion. To do this, I represent the payoff obtained by B_1 when selecting *hide*, given B_1 's type is q and B_2 's belief is μ , as $U(q, \mu)$. Building on this, I introduce the following sets:

$$D(q) := \{\mu \in [0, 1] : U(q, \mu) > U^*(q)\}$$

and

$$D^0(q) := \{\mu \in [0, 1] : U(q, \mu) \geq U^*(q)\}$$

Here $D(q)$ and $D^0(q)$ are belief sets of B_2 where B_1 gains a higher payoff from *hide* than his equilibrium payoff. A type q is eliminated by D1 criterion if there exists another type q' such that $D^0(q) \subset D(q')$.

In the equilibrium in Proposition 0, the on-path experiment in period 1 is the posterior split of \bar{q}_n and \bar{q}_{n+1} . According to the proof for Proposition 2, when the belief assigned to

B_2 is some $q \in (\bar{q}_n, \bar{q}_{n+1})$, B_1 chooses *hide* at \bar{q}_{n+1} . Thus, $q \in D^0(\bar{q}_{n+1})$ and $q \notin D(\bar{q}_n)$ (since $q > \bar{q}_n$). As a result, the type \bar{q}_{n+1} survives D1 criterion.

Thus, the off-path belief \bar{q}_{n+1} of B_2 at the history where B_1 chooses *hide* survives D1 criterion, and this is the off-path belief I assign to B_2 in the equilibrium.

When there is no refinement applied, we can have other equilibria since we can assign arbitrary off-path beliefs. However, with unrestricted bureaucrats, the current bureaucrat cannot gain an informational advantage from the following bureaucrat in equilibrium, and delay is never on the equilibrium path.

Proposition 4. *When B_2 is not restricted by the bureaucratic norm, delay never happens in any equilibrium.*

The proof is in Appendix A.5.

5 Conclusion

This paper studies the bureaucratic norm that a bureaucrat can base the decision only on the evidence recorded by his predecessor and the evidence presented directly to him. With this norm, when the predecessor hides the evidence he has, the belief updating of the second bureaucrat is restricted. In a bureaucracy with revolving bureaucrats, the first bureaucrat may strategically hide the evidence from the successor due to this norm, even if they share the same preferences. This incentive to delay and hide can make the first bureaucrat have a positive value of persuasion by forcing the developer to provide more information. With other parameters, the bureaucrat has zero value of persuasion. Unlike bureaucrats, the developer is always worse off due to the bureaucratic norm.

References

- Alonso, Ricardo, and Odilon Câmara. 2016. “Bayesian persuasion with heterogeneous priors.” *Journal of Economic Theory*, 165: 672–706.
- Bénabou, Roland, and Jean Tirole. 2002. “Self-confidence and personal motivation.” *The Quarterly Journal of Economics*, 117(3): 871–915.
- Bizzotto, Jacopo, Jesper Rüdiger, and Adrien Vigier. 2021. “Dynamic persuasion with outside information.” *American Economic Journal: Microeconomics*, 13(1): 179–94.
- de Clippel, Geoffroy, and Xu Zhang. 2022. “Non-Bayesian Persuasion.” *Journal of Political Economy*, 130(10): 2594–2642.
- Ely, Jeffrey C. 2017. “Beeps.” *American Economic Review*, 107(1): 31–53.
- Galperti, Simone. 2019. “Persuasion: The art of changing worldviews.” *American Economic Review*, 109(3): 996–1031.
- Honryo, Takakazu. 2018. “Dynamic persuasion.” *Journal of Economic Theory*, 178: 36–58.
- Kamenica, Emir, and Matthew Gentzkow. 2011. “Bayesian persuasion.” *American Economic Review*, 101(6): 2590–2615.
- Levy, Gilat, Inés Moreno de Barreda, and Ronny Razin. 2018. “Persuasion with correlation neglect.” *Unpublished manuscript, London School Econ.*
- Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov. 2020. “Persuading the principal to wait.” *Journal of Political Economy*, 128(7): 2542–2578.

Seidmann, Daniel J. 2005. “The effects of a right to silence.” *The Review of Economic Studies*, 72(2): 593–614.

Seidmann, Daniel J, and Alex Stein. 2000. “The right to silence helps the innocent: a game-theoretic analysis of the fifth amendment privilege.” *Harvard Law Review*, 430–510.

Smolin, Alex. 2020. “Dynamic evaluation design.” *Available at SSRN 3051703*.

Wilson, James Q. 1989. *Bureaucracy : what government agencies do and why they do it*. New York:Basic Books.

A Appendix

A.1 Proof for Lemma 3

Firstly, with $q_1 = 0$ or 1 , B_1 acts immediately.

We then calculate B_1 's payoff from *hide* when the prior and the posterior are p_0 and $q_1 \in (0, 1)$.

According to Alonso and Câmara (2016), if two players with different priors $\mu_1, \mu_2 \in (0, 1)$ observe the same experiment, their posteriors μ'_1 and μ'_2 satisfy:

$$\mu'_1 = \frac{\mu'_2 \frac{\mu_1}{\mu_2}}{\mu'_2 \frac{\mu_1}{\mu_2} + (1 - \mu'_2) \frac{1 - \mu_1}{1 - \mu_2}}$$

I denote B_1 's payoff from immediate actions under belief q_1 as $\hat{u}(q_1) = \max_{a \in \mathcal{A}} \{q_1 u(a, \omega = 1) + (1 - q_1)u(a, \omega = 0)\}$, and his payoff from delay under belief q_1 as $\tilde{u}(q_1)$. According to

Lemma 2, we do not need to consider *record*, so here I use $\tilde{u}(q_1)$ to represent B_1 's payoff from *hide* under q_1 .

Given *hide*, when facing the experiment in period 2, B_1 has prior q_1 but B_2 has prior p_0 . Notice that the developer's optimal experiment in period 2 is the split of 0 and $q_2 = \frac{1}{2}$ for B_2 , so in B_1 's eyes, the split is 0 and $q_2' = \frac{q_2 \frac{q_1}{p_0}}{q_2 \frac{q_1}{p_0} + (1-q_2) \frac{1-q_1}{1-p_0}}$. Moreover, the probability of q_2' is $\frac{q_2 \frac{q_1}{p_0} + (1-q_2) \frac{1-q_1}{1-p_0}}{\frac{q_2}{p_0}}$. Since B_2 chooses $a = 0$ at 0 and $a = 1$ at q_2 , B_1 's expected payoff from *hide* is:

$$\tilde{u}(q_1) = \delta \cdot \left(C - \frac{p_0}{1-p_0} (1-q_1) \right)$$

When $q_1 < \frac{1}{2}$, the optimal immediate action is $a = 0$; when $q_1 \geq \frac{1}{2}$, the optimal immediate action is $a = 1$. Thus, We have:

$$\hat{u}(q_1) = \begin{cases} C - q_1, & \text{if } q_1 < \frac{1}{2} \\ C - (1 - q_1), & \text{if } q_1 \geq \frac{1}{2} \end{cases}$$

Comparing $\tilde{u}(q_1)$ and $\hat{u}(q_1)$, we can get that when $\delta > \frac{2C-1}{2C-\frac{p_0}{1-p_0}}$, there are $\hat{\alpha}(\delta, p_0) = \frac{C(1-\delta)-p_0(C(1-\delta)-\delta)}{1-(1-\delta)p_0}$ and $\hat{\beta}(\delta, p_0) = \frac{1-C(1-\delta)-(1+\delta-C(1-\delta))p_0}{1-(1+\delta)p_0}$, where $p_0 < \hat{\alpha}(\delta, p_0) < \hat{\beta}(\delta, p_0) < 1$, and B_1 chooses $a = 1$ with $q_1 \geq \hat{\beta}(\delta, p_0)$; chooses $a = 0$ with $q_1 < \hat{\alpha}(\delta, p_0)$; chooses *hide* in between.

When $\delta \leq \frac{2C-1}{2C-\frac{p_0}{1-p_0}}$, B_1 chooses $a = 1$ with $q_1 \geq \frac{1}{2}$ and $a = 0$ with $q_1 < \frac{1}{2}$.

A.2 Proof for Proposition 1

Knowing B_1 's strategy from Lemma 3, we can figure out the developer's optimal experiment in period 1.

When $\delta \leq \frac{2C-1}{2C-\frac{p_0}{1-p_0}}$, B_1 chooses $a = 1$ with $q_1 \geq \frac{1}{2}$ and $a = 0$ with $q_1 < \frac{1}{2}$. Since $p_0 < \frac{1}{2}$, the optimal experiment is the split between 0 and $\frac{1}{2}$ in B_1 's eyes.

When $\delta > \frac{2C-1}{2C-\frac{p_0}{1-p_0}}$, we have two cutoff points in B_1 's strategy $\hat{\alpha}(\delta, p_0)$, $\hat{\beta}(\delta, p_0)$. Since $p_0 < \hat{\alpha}(\delta, p_0) < \hat{\beta}_2(\delta, p_0)$, the optimal experiment is either the split between 0 and $\hat{\alpha}(\delta, p_0)$ or the split between 0 and $\hat{\beta}_2(\delta, p_0)$. For the developer, her expected payoff from the latter one is $\frac{p_0}{\hat{\beta}_2(\delta, p_0)}$. As for the former one, at 0 she gets 0, at $\hat{\alpha}(\delta, p_0)$ the game proceeds into the next period, and her expected payoff is $\delta \cdot (p_0 + \frac{p_0}{1-p_0} \frac{1-\hat{\alpha}(\delta, p_0)}{\hat{\alpha}(\delta, p_0)} p_0)$.

Comparing the payoffs from two experiments, the delaying experiment is the optimal one if and only if:⁷

$$\frac{\delta(1-C(1-\delta))p_0}{C(1-\delta)-p_0(C(1-\delta)-\delta)} - \frac{C(1-\delta)(1-p_0)}{1-C(1-\delta)-p_0(1+\delta-C(1-\delta))} \geq 1-\delta$$

$$\Leftrightarrow r_1 p_0^2 + r_2 p_0 + r_3 \geq 0 \tag{1}$$

Where:

$$r_1 = -\delta(1-C(1-\delta))(1+\delta-C(1-\delta)) - C(1-\delta)(C(1-\delta)-\delta)$$

$$-(1-\delta)(C(1-\delta)-\delta)(1+\delta-C(1-\delta)),$$

⁷When indifferent between two experiments, the tie-breaker makes the developer choose the delaying experiment, which is less informative.

$$\begin{aligned}
r_2 &= \delta(1 - C(1 - \delta))^2 + C(1 - \delta)(C(1 - \delta) - \delta) + C^2(1 - \delta)^2 \\
&+ C(1 - \delta)^2(1 + \delta - C(1 - \delta)) + (1 - \delta)(1 - C(1 - \delta))(C(1 - \delta) - \delta),
\end{aligned}$$

$$r_3 = -C^2(1 - \delta)^2 - C(1 - \delta)^2(1 - C(1 - \delta)).$$

Notice that for $p_0 = 0$ and $p_0 = \frac{1}{2}$, LHS of (1) is negative for δ close to 1. Moreover, when $\delta = 1$, we have $r_1 < 0$ and that the solution to (1) is $p_0 \in [0, \frac{1}{2}]$. So, when δ is close enough to 1, the solution to (1) would be $p_0 \in [\alpha(\delta), \beta(\delta)]$, where $0 < \alpha(\delta) < \beta(\delta) < \frac{1}{2}$.

It is easy to see that we can have $\bar{\delta} < 1$ such that for any $p_0 \leq \frac{1}{2} - \epsilon$ and $\delta > \bar{\delta}$, the optimal experiment choice depends on $p_0 \in [\alpha(\delta), \beta(\delta)]$ or not.

A.3 Proof for Proposition 2

In this proof, I assume $p_0 \in (\bar{q}_n, \bar{q}_{n+1})$.

Firstly notice that *record* does not happen in the equilibrium, due to $\bar{V}_B = 0$ in the static game. We only need to look at immediate actions and *hide*.

For the following proof, define $\hat{v}(\cdot)$ as the developer's payoff function in posteriors in the static game, i.e., $\hat{v}(q) = v(a_n)$ if $q \in [\bar{q}_n, \bar{q}_{n+1})$.

Lemma 4. *At any posterior $q < p_0$, hide is not the optimal action.*

Proof. Consider three points on the graph of the developer's payoff function in posteriors $(t_1, \hat{v}(t_1))$, $(t_2, \hat{v}(t_2))$, $(t_3, \hat{v}(t_3))$, $t_1 < t_2 < t_3$. If point 2 is above the line connecting points 1 and 3, we have

$$\frac{\hat{v}(t_2) - \hat{v}(t_1)}{t_2 - t_1} > \frac{\hat{v}(t_3) - \hat{v}(t_1)}{t_3 - t_1} \Leftrightarrow \frac{t_3 - t_1}{t_2 - t_1} > \frac{\hat{v}(t_3) - \hat{v}(t_1)}{\hat{v}(t_2) - \hat{v}(t_1)}.$$

According to Alonso and Câmara (2016), when B_2 with prior p_0 has a posterior t_i after an experiment, B_1 with prior q has the posterior $t'_i = \frac{t_i \frac{q}{p_0}}{t_i \frac{q}{p_0} + (1-t_i) \frac{1-q}{1-p_0}}$. Because $q < p_0$, we have

$$\frac{t'_3 - t'_1}{t'_2 - t'_1} = \frac{t_3 - t_1 \frac{t_2 \frac{q}{p_0} + (1-t_2) \frac{1-q}{1-p_0}}{t_2 - t_1} > \frac{t_3 - t_1}{t_2 - t_1}$$

Thus, $\frac{t'_3 - t'_1}{t'_2 - t'_1} > \frac{\hat{v}(t_3) - \hat{v}(t_1)}{\hat{v}(t_2) - \hat{v}(t_1)}$, and $(t'_2, \hat{v}(t_2))$ is still above the line connecting $(t'_1, \hat{v}(t_1))$ and $(t'_3, \hat{v}(t_3))$.

Then we conclude that \bar{q}_i ($i = 0, 1, \dots, N$), which is on the concave closure of the developer's utility $\hat{v}(\cdot)$, is still on the concave closure after considering heterogeneous priors when $q < p_0$. Thus, when the developer holds prior q (same as B_1) and B_2 holds prior p_0 ($p_0 > q$ and $p_0 \in (\bar{q}_n, \bar{q}_{n+1})$), the optimal experiment in period 2 is the split of \bar{q}_n and \bar{q}_{n+1} for prior p_0 .

When B_2 with prior p_0 has the posterior split \bar{q}_n and \bar{q}_{n+1} , B_1 with prior q has the split of $\bar{q}'_n < \bar{q}_{n-1}$ and $\bar{q}'_{n+1} < \bar{q}_n$. Notice that B_2 chooses a_n and a_{n+1} at B_1 's posteriors \bar{q}'_n and \bar{q}'_{n+1} . But with belief \bar{q}'_{n+1} , bureaucrats prefers a_n to a_{n+1} . As a result, for B_1 with prior q , the expected payoff from the experiment is smaller than his expected payoff by choosing a_n at belief q . Thus, when the first experiment gives posterior $q < p_0$, B_1 will not choose *hide*, which is worse than choosing a_n . ■

By Lemma 4, B_1 always chooses immediate actions for $q < p_0$, and *hide* can only happen at $q > p_0$. Also notice that when B_1 chooses *hide* at q , the point of q and the developer's payoff from *hide* is always inside the concave closure of her static utility function $\hat{v}(\cdot)$. Thus, if $p_0 \in [\bar{q}_n, \bar{q}_{n+1})$, for the optimal experiment chosen by the developer in period 1, which has two signals due to binary states, one signal induces the posterior \bar{q}_n .

Then we need to look at the other posterior of the optimal experiment, which is larger than p_0 . Denote this posterior as q^* .

When B_1 chooses *hide* at q , B_2 has prior p_0 . Let the optimal experiment in period 2 be the split of \bar{q}_i and \bar{q}_j for B_2 ($i < j$). The developer who has prior q will think the split is $\frac{\bar{q}_i \frac{q}{p_0}}{\bar{q}_i \frac{q}{p_0} + (1-\bar{q}_i) \frac{1-q}{1-p_0}}$ and $\frac{\bar{q}_j \frac{q}{p_0}}{\bar{q}_j \frac{q}{p_0} + (1-\bar{q}_j) \frac{1-q}{1-p_0}}$ instead, with B_2 choosing a_i and a_j at two posteriors. Her payoff from B_1 's action *hide* (without discount) is:

$$\begin{aligned} v_h(q) &:= \left[\frac{1 - \bar{q}_i}{1 - p_0} + \frac{\bar{q}_i - p_0}{p_0(1 - p_0)} q \right] \frac{\bar{q}_j - p_0}{\bar{q}_j - \bar{q}_i} v(a_i) + \left[\frac{1 - \bar{q}_j}{1 - p_0} + \frac{\bar{q}_j - p_0}{p_0(1 - p_0)} q \right] \frac{p_0 - \bar{q}_i}{\bar{q}_j - \bar{q}_i} v(a_j) \\ &\Rightarrow v'_h(q) = [v(a_j) - v(a_i)] \frac{\bar{q}_j - p_0}{p_0(1 - p_0)} \frac{p_0 - \bar{q}_i}{\bar{q}_j - \bar{q}_i} \end{aligned}$$

When the experiment chosen in period 2 stays unchanged, the developer's payoff when B_1 chooses *hide* is linear in posterior q . However, when q goes up, the experiment chosen in period 2 can change – the points originally on the concave closure of $\hat{v}(\cdot)$ can be inside the closure when $q > p_0$ (the opposite way compared to $q < p_0$).

By the similar argument as in Lemma 4, when $q > p_0$, if one point on the graph of developer's payoff $\hat{v}(\cdot)$ is below a line connecting two other points on the graph, it is still below the line after considering heterogeneous priors. Furthermore, as q grows even larger, the point still stays below the line. Also notice that the point that is originally above the line can become below the line as q goes up, by the similar argument as in Lemma 4. As a result, as q goes up, the new experiment chosen by the developer in period 2 is a mean-preserving spread of the previous experiment. In the new experiment, we will have larger $\bar{q}_j, v(a_j)$ and/or smaller $\bar{q}_i, v(a_i)$. Thus, $v'_h(q)$ becomes larger when q goes up and a new optimal experiment in period 2 occurs. Then we know $v_h(q)$ is a convex function in q when

$q > p_0$. By the convexity of the developer's payoff from *hide*, we have the following claim.

Claim 1. *If we have three posteriors s_1, s_2, s_3 where B_1 chooses *hide* and $s_3 > s_2 > s_1 > p_0$, the second point of the developer's optimal experiment in period 1 cannot be s_2 .*

With the claim, I will prove that when the second point of the optimal experiment in period 1 is $q^* > \bar{q}_{n+1}$, B_1 does not choose *hide* at q^* . Without loss, let $q^* \in (\bar{q}_{n+1}, \bar{q}_{n+2})$.

Suppose B_1 chooses *hide* at $q^* \in (\bar{q}_{n+1}, \bar{q}_{n+2})$, which is the second posterior of the optimal experiment. Then B_1 must choose *hide* at \bar{q}_{n+1} to make q^* the optimal choice, because \bar{q}_{n+1} is a better choice for the developer than q^* in period 1 if B_1 chooses the immediate action a_{n+1} at \bar{q}_{n+1} . This is because the point $(\bar{q}_{n+1}, v(a_{n+1}))$ is on the concave closure of the developer's payoff function in posteriors $\hat{v}(\cdot)$. Actually, when B_1 chooses the immediate action at \bar{q}_{n+1} , the optimal experiment in period 1 is the static optimal one.

If B_1 also chooses *hide* in $\mathcal{B}_\epsilon(q^*)$, the claim says q^* is not in the optimal experiment, contradiction.

If B_1 chooses *hide* at $q = q^* - \epsilon$ and a_{n+1} at $q^* + \epsilon$ for any small ϵ , we need to have $v_h(q^*) > v(a_{n+1})$ to keep q^* optimal for the developer. That means when the posterior is q^* in period 1 and B_1 chooses *hide*, in the optimal experiment in period 2, one of the posterior must induce a_k s.t. $k > n + 1$, to make $v_h(q^*) > v(a_{n+1})$. As a result, B_1 must choose *hide* at \bar{q}_{n+2} . Otherwise B_1 chooses the immediate action a_{n+2} at \bar{q}_{n+2} , and the split of \bar{q}_n and \bar{q}_{n+2} is a better experiment than the split of \bar{q}_n and q^* in period 1 (since $(\bar{q}_{n+2}, v(a_{n+2}))$ is on the concave closure of $\hat{v}(\cdot)$). Now we have $\bar{q}_{n+2} > q^*$ and $\bar{q}_{n+1} < q^*$ choosing *hide*, so q^* is not the optimal choice of the second posterior if B_1 chooses *hide* at q^* according to the claim, contradiction.

If B_1 chooses *hide* at $q = q^* + \epsilon$ and a_{n+1} at $q^* - \epsilon$ for any small ϵ , we have $\bar{q}_{n+1} < q^*$ and $q^* + \epsilon > q^*$ where B_1 chooses *hide*, so the claim says q^* is not the optimal choice, contradiction.

Thus, we conclude that when $q^* \in (\bar{q}_{n+1}, \bar{q}_{n+2})$ is induced in the optimal experiment in period 1, B_1 does not choose *hide* at q^* . The conclusion extends to $q^* \geq \bar{q}_{n+2}$.

So, only when the second posterior of the optimal experiment in period 1 is $q^* < \bar{q}_{n+1}$, *hide* possibly happens in the optimal experiment. When (1) δ is large enough, (2) $q^* < \bar{q}_{n+1}$ is in the optimal experiment, and (3) B_1 chooses *hide* at q^* , B_1 is indifferent between choosing *hide* and a_n at q^* . Notice that at the first posterior of the optimal experiment \bar{q}_n , B_1 chooses a_n . Thus, in this optimal experiment where *hide* happens at $q^* < \bar{q}_{n+1}$, B_1 has zero value of persuasion.

Also notice that when the second posterior of the optimal experiment is $q^* < \bar{q}_{n+1}$, B_1 does not choose an immediate action at q^* . This is because if B_1 chooses an immediate action at such q^* , the action will be a_n , and the developer's expected payoff from this experiment is $v(a_n)$. But the split of \bar{q}_n and 1 is a better choice.

When the second posterior of the optimal experiment is $q^* > \bar{q}_{n+1}$, B_1 will choose an immediate action at q^* and this gives B_1 positive value of persuasion.

In conclusion, if δ is large enough, when the second posterior of the optimal experiment $q^* < \bar{q}_{n+1}$, B_1 chooses *hide* at q^* and he has $V_B = 0$. When $q^* > \bar{q}_{n+1}$, B_1 chooses an immediate action at q^* and he has $V_B > 0$.

A.4 Proof for Proposition 3

To figure out the optimal experiment in period 1 in the convex case, we need to consider *record*, because it can be the optimal action of B_1 under some posteriors. This is because $\bar{V}_B > 0$.

As defined in the proof for Lemma 3, $\hat{u}(q_1)$ is B_1 's payoff by acting immediately and $\tilde{u}(q_1)$ is B_1 's payoff by delaying the decision. $\hat{u}(q_1)$ remains to be $\max_{a \in \mathcal{A}} \{q_1 u(a, \omega = 1) + (1 - q_1)u(a, \omega = 0)\}$ as in the proof for Lemma 3. However, in this case, we need to take *record* into consideration, so instead of representing payoff from *hide*, $\tilde{u}(q_1)$ is B_1 's payoff from *hide* or *record*, depending on which one is better for him.

By the similar argument as in Lemma 4, for $q_1 > p_0$, *hide* is a better choice than *record*; for $q_1 < p_0$, *record* is a better choice than *hide*. Now I fix the initial prior p_0 , if B_1 chooses to delay the decision at posterior q_1 , his payoff is

$$\tilde{u}(q_1) = \begin{cases} \delta(C - \frac{1}{3}q_1), & \text{if } q_1 < p_0 \\ \delta[C - \frac{1}{3}\frac{p_0}{1-p_0}(1-q_1)], & \text{if } q_1 \geq p_0 \end{cases}.$$

Thus, the payoff of the developer at q_1 when B_1 delays the decision, denoted as $\tilde{v}(q_1)$, is:

$$\tilde{v}(q_1) = \begin{cases} \delta \cdot \frac{4}{3}q_1, & \text{if } q_1 < p_0 \\ \delta[q_1 + \frac{1}{3}\frac{p_0}{1-p_0}(1-q_1)], & \text{if } q_1 \geq p_0 \end{cases}.$$

Also notice that B_1 's payoff by choosing an immediate action at q_1 is:

$$\hat{u}(q_1) = \begin{cases} C - q_1, & \text{if } q_1 < \bar{q}_1 \\ C - \frac{1}{4}, & \text{if } \bar{q}_1 \leq q_1 < \bar{q}_2 \\ C - (1 - q_1), & \text{if } q_1 \geq \bar{q}_2 \end{cases} .$$

When $\delta = 1$, we have $\tilde{u}(q_1) > \hat{u}(q_1)$, $\forall q_1 \in (0, 1)$. We also have that $\tilde{v}'(q_1) = \frac{4}{3}\delta$ when $q_1 < p_0$, $\tilde{v}'(q_1) = \delta(1 - \frac{1}{3}\frac{p_0}{1-p_0}) < \frac{4}{3}\delta$ when $q_1 > p_0$, and $\tilde{v}(1) = 1$ when $\delta = 1$. Thus, for any p_0 , there exists $\eta(p_0) < 1$ such that $\forall \delta > \eta(p_0)$, it is optimal to choose an uninformative experiment in period 1 when the initial prior is p_0 .

Moreover, since we are looking at p_0 such that $p_0 \notin \mathcal{B}_\epsilon(\bar{q}_i)$ and $p_0 < \bar{q}_2$, the supremum of $\eta(p_0)$ among these p_0 will be smaller than 1, denoted as $\tilde{\delta}$.

A.5 Proof for Proposition 4

Suppose B_1 delays at only one signal or delays at multiple signals but hides only one of them in an equilibrium, B_2 still has the same belief as B_1 . Then no value of persuasion in the static game makes delay worse than an immediate action. Thus, there is a profitable deviation.

Suppose B_1 delays and hides at multiple signals in an equilibrium, B_2 's prior belief is the weighted average of posteriors induced by these signals, denoted as \bar{p} , since the developer is using the same experiment for these delayed signals (assumed). Moreover, the developer using the same experiment for these delayed signals means that he is choosing the optimal experiment as if he has a prior of the weighted average \bar{p} . Suppose $\bar{p} \in [\bar{q}_{n-1}, \bar{q}_n)$, our

assumption says the experiment chosen by the developer in period 2 is the split of \bar{q}_{n-1} and \bar{q}_n , and B_2 chooses a_n at \bar{q}_{n-1} , a_{n+1} at \bar{q}_n .

Clearly, there is a hidden signal inducing belief $q' < \bar{p}$. If B_1 has belief q' , the experiment in period 2 is inducing $q'_1 < \bar{q}_{n-1}$ and $q'_2 < \bar{q}_n$ for him. The probabilities of getting two signals are denoted as x and $(1-x)$, where $xq'_1 + (1-x)q'_2 = q'$. And his expected payoff at these two posteriors are $\mathbb{E}_{q'_1}[u(a_n, \omega)]$ and $\mathbb{E}_{q'_2}[u(a_{n+1}, \omega)]$. Notice that a_n and a_{n+1} are indifferent to bureaucrats at belief \bar{q}_n , and bureaucrats have supermodular preference. Thus, $\mathbb{E}_{q'_2}[u(a_{n+1}, \omega)] < \mathbb{E}_{q'_2}[u(a_n, \omega)]$ for $q'_2 < \bar{q}_n$.

As a result, B_1 's expected payoff from the experiment is

$$\begin{aligned} x\mathbb{E}_{q'_1}[u(a_n, \omega)] + (1-x)\mathbb{E}_{q'_2}[u(a_{n+1}, \omega)] &< x\mathbb{E}_{q'_1}[u(a_n, \omega)] + (1-x)\mathbb{E}_{q'_2}[u(a_n, \omega)] \\ &= \mathbb{E}_{q'}[u(a_n, \omega)] \end{aligned}$$

Thus, B_1 's expected payoff from delay and hide is smaller than an immediate action, he has a profitable deviation.

As for $\bar{p} \geq \bar{q}_N$, there is an profitable deviation to an immediate action for B_1 because the developer chooses an uninformative experiment in period 2.