

Optimal Delay in Committees

ETTORE DAMIANO

University of Toronto

LI, HAO

University of British Columbia

WING SUEN

University of Hong Kong

March 20, 2012

Abstract. We consider a symmetric two-member, two-alternative, three-state committee problem of information aggregation. A tight bound on maximum delay that can be imposed limits what can be implemented with a single round of delay. In a dynamic delay mechanism, each round is a direct revelation mechanism with delay uniformly bounded from above. We characterize optimal sequences of delays for any given uniform upper bound on each delay. Any optimal sequence is finite, inducing in equilibrium both a “deadline play,” with no attempt made to reach the efficient decision just before the deadline, and “stop-and-start,” with maximum attempt alternating with no attempt before the deadline play kicks in. The latter is achieved by binding and slackening the bound on delay in turn.

JEL classification. C72, D7, D82

Keywords. delay mechanism, limited commitment, stop-and-start, localized variation method

1. Introduction

The committee problem is a prime example of strategic information aggregation.¹ The committee decision is public, affecting the payoff of each committee member; the information for the decision is dispersed in the committee and is private to committee members; and committee members have conflicting interests in some states and common interests in others. As a mechanism design problem, the committee problem has the following distinguishing features. First, there are no side transfers, unlike in Myerson and Satterthwaite (1983). Second, each committee member possesses private information about the state, in contrast with the strategic information communication problem of Holmstrom (1984). Third, the number of committee members is small, unlike the large election problem studied by Feddersen and Pesendorfer (1997).

It is well-known in the mechanism design literature that a universally bad outcome or a sufficiently large penalty can be useful toward implementing desirable social choice rules when agents have complete information about one another's preferences (Moore and Repullo 1990; Dutta and Sen 1991). In the absence of side transfers in a committee problem, costly delay naturally emerges as a tool to provide incentives to elicit private information from committee members. In a companion paper (Damiano, Li and Suen 2012), we show how introducing delay in committee decision-making can result in efficient information aggregation and ex ante welfare gain among committee members. In that paper, the cost incurred with each additional round of delay is fixed and is assumed to be small relative to the value of the decision at stake. If we drop the assumption of small delay, delay may even achieve the first best: under the threat of collective punishment, committee members would reach the Pareto efficient decision immediately with no delay incurred on the equilibrium path. However, achieving the first best requires delay to be sufficiently costly. This poses at least two problems for mechanism design. First, the mechanism is not robust in that a "mistake" made by one member will produce a bad outcome for all. More importantly, because imposing a lengthy delay is very costly ex post, the mechanism is not credible unless there is strong commitment power. This paper takes a limited commitment approach to mechanism design in committee problems.² Specifically we assume that the mechanism designer can commit to imposing a delay penalty and not renegotiating it away immediately upon a disagreement, but there is an upper bound on the amount of delay that he can commit to impose. That is, he

¹See Li, Rosen and Suen (2001) for an example and Li and Suen (2009) for a literature review.

²See Bester and Strausz (2001), Skreta (2006), and Kolotilin, Li and Li (2011) for other models of limited commitment.

can commit to “wasting” a small amount of money (or time) ex post, but not too much. The upper bound on the delay reflects the extent of his limited commitment power. In our model, a sufficiently tight bound on delay would imply that the efficient decision cannot be reached immediately, and that delay will occur in equilibrium. This gives rise to dynamic delay mechanisms in which committee members can make the collective decision in a number of rounds, punctuated by a sequence of delays between successive rounds, and with each delay uniformly bounded from above depending on the commitment power. When delays will be incurred in equilibrium, there is a non-trivial trade-off between raising the level of punishment through delay and the resulting improvement in quality of the collective decision. This framework allows us to ask questions that cannot be addressed in our companion paper: Does punishment (delay) work better if it is front-loaded or back-loaded? Is it optimal to maintain a constant sequence or delays between successive rounds? Do deadlines for agreements arise endogenously as an optimal arrangement? These questions are the subject of the present paper.

The model we adopt in this paper is a slightly simplified version of Damiano, Li and Suen (2012). In this symmetric, two-member committee problem, there are two alternatives to be chosen, with the two committee members favoring a different alternative ex ante. One can think of this as a situation in which each member derives some private benefit if his ex ante favorite alternative is adopted. The payoffs from the two alternatives also depend on the state. If it is known that the state is a “common interest state,” both members would choose the same alternative despite their ex ante preference. If it is known that the state is a “conflict state,” the two members would prefer to choose their own ex ante favorites. Therefore, the prior probability of the conflict state is an indicator of the degree of conflict in the committee. Information about the state, however, is dispersed among the two members. Each member cannot be sure about the state based on his private information alone, but they could jointly deduce the true state if they truthfully share their private information. This model is meant to capture the difficulties of reaching a mutually preferred collective decision when preference-driven disagreement (difference in ex ante favorites) is confounded with information-driven disagreement (difference in private information). Damiano, Li and Suen (2012) provide examples including competing firms choosing to adopt a common industry standard, faculty members in different specialties recruiting job candidates, and separated spouses deciding on child custody. For tractability, Damiano, Li and Suen (2012) adopts a model in which members choose their actions in continuous time. In this paper, since the focus is the optimal sequencing of delays, members move in discrete “rounds,” with variable delays between successive rounds.

We show an impossibility result in the committee problem that we study: in the absence of side transfers, there is no incentive compatible mechanism that Pareto-dominates flipping a coin if the degree of conflict in the committee is sufficiently high. This is a stark illustration of the difficulties of efficient information aggregation because the members would have agreed to make the same choice (in the common interest state) had they been able to share their information. We also show that introducing a collective punishment in the form of delay if the members disagree may improve decision-making. Indeed, committing to a sufficiently long and thus costly delay would achieve the first-best outcome of Pareto efficient decision—implement the agreed alternative in each common interest state and flip a coin in the conflict state—without actually incurring the delay.

This paper focuses on situations in which the first best is unachievable because there is a limit to how much time members can commit to wasting when both committee members persist with their own favorite alternatives. The members can attempt to reach an agreement repeatedly in possibly an infinite number of rounds, but the length of delay between successive rounds cannot exceed a fixed upper bound. In this framework, any given sequence of delays is a mechanism that induces a dynamic game between the members, and we examine the “optimal” sequence that maximizes the members’ ex ante payoffs subject to the uniform bound on the length of each delay.

The dynamic game induced by a delay mechanism resembles a war of attrition with incomplete information and interdependent values.³ In equilibrium of this game, an informed member (who knows that the state is a common interest state for his alternative) always persists with his own favorite. An uninformed member (who is unsure whether it is a conflict state or a common interest state for his opponent’s alternative) may randomize between persisting with his favorite and conceding with his opponent’s favorite. Because of the structure of this equilibrium strategy profile, an uninformed member’s belief that his opponent is also uninformed (i.e., the state is a conflict state) weakly decreases in the next round when both members are observed to be persisting with their favored alternatives in the current round. Given any fixed delay mechanism, finding the equilibrium of the dynamic game involves jointly solving the sequence of actions chosen by the uninformed, the sequence of beliefs, and the sequence of continuation payoffs. For an arbitrary sequence of delays, such an approach is not manageable and does not yield any particular insights. In this paper, we introduce a “localized variation method” to study the design of an optimal delay mechanism. Consider changing the delay at some

³See also Hendricks, Weiss and Wilson (1988), Cramton (1992), Abreu and Gul (2000), Deneckere and Liang (2006), and Damiano, Li and Suen (2012).

round t . We study its effect by simultaneously adjusting the delay at round $t - 1$ (through the introduction of an extra round if necessary) in such a way that keeps the continuation payoff for round $t - 2$ fixed, and adjusting the delay at round $t + 1$ (also through the introduction of an extra round if necessary) in such a way that keeps the equilibrium belief at round $t + 2$ constant. In this manner the effects of these variations are confined to a narrow window, so that there is no need to compute the entire sequence of equilibrium actions, equilibrium beliefs and continuation payoffs. It turns out that just by employing this local variation method, we can arrive at a complete characterization of optimal delay mechanisms.

The main result of this paper is a characterization of all delay mechanisms that have a perfect Bayesian equilibrium with the maximum ex ante expected payoff to each member. Such “optimal” delay mechanisms have interesting properties that we highlight in Section 3 and establish separately in Section 4. First, we show that any optimal delay mechanism is a finite sequence of delays. Thus, it is optimal to have a final round, or “deadline,” for making the decision; failing to make the decision in the final round would entail that the decision is made by flipping a coin after incurring the final delay. In an optimal delay mechanism, however, an informed committee member always persists with his favorite alternative while an uninformed member concedes to the favorite alternative of his fellow member with probability one at the final round if it is reached. The decision is thus always Pareto efficient in equilibrium. Second, we show that in any perfect Bayesian equilibrium of an optimal delay mechanism there is a “deadline play,” in which each member persists with his own favorite alternative for a number of rounds before the deadline. This means that it is optimal to have the committee make no attempt at reaching a decision just before the deadline arrives. Third, we show that an optimal delay mechanism induces a “stop-and-start” pattern of making concessions. At the first round, each uninformed member starts by adopting a mixed strategy with the maximum feasible probability to conceding to the favorite alternative of his fellow member. If the committee fails to reach an agreement, the uninformed types would make no concession in the next round or next few rounds. After one or more rounds of no concession, the uninformed types start making the maximum feasible concession again, and would stop making any concession for one or more rounds upon failure to reach an agreement. Thus equilibrium play under the optimal delay mechanism alternates between maximum concession and no concession, until the deadline play kicks in. Perhaps surprisingly, it is not optimal to induce an uninformed committee member to concede to the favorite alternative of his fellow member with a positive probability in each round of the dynamic delay mechanism. Instead, under an optimal delay mechanism, a round of maximum conces-

sion probability by an uninformed member is immediately followed by no concession. To achieve this “stop-and-start” pattern of equilibrium play, the length of delay between successive rounds cannot be constant throughout. Before the deadline play is reached, the delay is equal to the limited commitment bound in rounds when members are making concessions, and is strictly lower than the bound in rounds when they are not making concessions.

As the uniform upper bound on delay goes to zero, optimal delay mechanisms characterized in the present paper converge to the optimal deadline in the continuous-delay model of Damiano, Li and Suen (2012) after adjusting for the differences in the payoff structure in the latter model. This convergence is derived in Section 5. There we also briefly discuss two robustness issues regarding our main results. The first robustness issue concerns the implicit assumptions on the payoff structure made in the committee problem introduced in Section 2.1. In particular, we have assumed that the two members derive the same benefit of implementing the correct alternative in a common interest state, which is also equal to the private benefit of implementing one’s ex ante favorite alternative in the conflict state. We show that our characterization of optimal delay mechanisms remains qualitatively valid for general payoff structures. The second robustness issue has to do with the assumption made in the delay mechanism introduced in Section 2.2. Specifically, we assume that in each round a particular direct revelation mechanism is played: any agreement leads to the implementation of the agreed alternative without delay, and a disagreement caused by the two members conceding to each other’s favorite alternative leads to a coin flip without delay. It turns out that these restrictions on the mechanism used in each round are not without loss of generality. We discuss some implications of this issue.

2. Model

2.1. A simple committee problem

Two players, called LEFT and RIGHT, have to make a joint choice between two alternatives, l and r . There are three possible states of the world: L , M , and R . We assume that the prior probability of state L and state R is the same. The relevant payoffs for the two players are summarized in the following table:

In each cell of this table, the first entry is the payoff to LEFT and the second is the payoff to RIGHT. We normalize the payoff from making the preferred decision to 1 and let the payoff from making the less preferred decision be $1 - 2\lambda$. The parameter $\lambda > 0$

	L	M	R
l	$(1, 1)$	$(1, 1 - 2\lambda)$	$(1 - 2\lambda, 1 - 2\lambda)$
r	$(1 - 2\lambda, 1 - 2\lambda)$	$(1 - 2\lambda, 1)$	$(1, 1)$

is the loss from making the wrong decision relative to a coin flip. In state L both players prefer l to r , and in state R both prefer r to l . The two players' preferences are different when the state is M : LEFT prefers l while RIGHT prefers r . We refer to l as the ex ante favorite alternative for LEFT, and r as the ex ante favorite alternative for RIGHT. In this model there are elements of both common interest (states L and R) and conflict (state M) between these two players.

The information structure is such that LEFT is able to distinguish whether the state is L or not, while RIGHT is able to distinguish whether the state is R or not. Such information is private and unverifiable. When LEFT knows that the state is L , or when RIGHT knows that the state is R , we say they are "informed;" otherwise, we say they are "uninformed." Thus, an informed LEFT always plays against an uninformed RIGHT (in state L), and an informed RIGHT always plays against an uninformed LEFT (in state R). Two uninformed players playing against each other can only occur in state M . Without information aggregation, however, an uninformed LEFT does not know whether the state is M or R . Let $\gamma_1 < 1$ denote his belief that the state is M .⁴ We note that γ_1 can be interpreted as the ex ante degree of conflict. When γ_1 is high, an uninformed player perceives that his opponent is likely to have different preferences regarding the correct decision to be chosen.

In the absence of side transfers, if $\gamma_1 \leq 1/2$, the following simultaneous voting game implements the Pareto efficient outcome. Imagine that each player votes l or r , with the agreed alternative implemented immediately and any disagreement leading to an immediate fair coin flip between l and r and a payoff of $1 - \lambda$ to each player. It is a dominant strategy for an informed player to vote for his ex ante favorite alternative. Given this, because $\gamma_1 \leq 1/2$, it is optimal for an uninformed player to do the opposite. This follows because, regardless of the probability x_1 that the opposing uninformed player votes for his own favorite, the payoff from voting the opponent's favorite is higher than the payoff from voting one's own favorite:

$$\gamma_1 [x_1(1 - 2\lambda) + (1 - x_1)(1 - \lambda)] + 1 - \gamma_1 \geq \gamma_1 [x_1(1 - \lambda) + (1 - x_1)] + (1 - \gamma_1)(1 - \lambda).$$

⁴An uninformed RIGHT shares the same belief. The implied common prior beliefs are: the state is M with probability $\gamma_1/(2 - \gamma_1)$, and is L or R each with probability $(1 - \gamma_1)/(2 - \gamma_1)$.

The equilibrium outcome is Pareto efficient: the informed player gets the highest payoff of 1, and the uninformed player also gets 1 when the state is the common interest state in favor of his opponent or otherwise gets $1 - \lambda$ from a coin-flip in the conflict state.

In contrast, if $\gamma_1 > 1/2$, the unique equilibrium in the above voting game has both the informed and the uninformed types voting for their favorite alternatives. As a result, the decision is always made by a coin-flip in equilibrium, despite the presence of a mutually preferred alternative in a common interest state. Indeed, a standard application of the revelation principle can be used to generalize this negative result into an impossibility result: in the absence of side transfers, there is no incentive compatible mechanism that Pareto dominates flipping a coin when $\gamma_1 > 1/2$, and thus, our model provides a stark environment that illustrates the severe restrictions on efficient information aggregation in committees when side transfers are not allowed. To see this, let us write q_R , q_M , and q_L as the probabilities of implementing alternative r when the true states are R , M , and L , respectively, and let \tilde{q} be the probability of implementing r when the reports are inconsistent, that is, when both report that they are informed. The incentive constraint for an informed RIGHT is

$$q_R + (1 - q_R)(1 - 2\lambda) \geq q_M + (1 - q_M)(1 - 2\lambda).$$

From the above condition and a symmetric condition for informed LEFT we immediately get $q_R \geq q_M$ and $q_M \geq q_L$. The incentive constraint for an uninformed RIGHT is

$$\begin{aligned} & \gamma [q_M + (1 - q_M)(1 - 2\lambda)] + (1 - \gamma) [q_L(1 - 2\lambda) + (1 - q_L)] \\ & \geq \gamma [q_R + (1 - q_R)(1 - 2\lambda)] + (1 - \gamma) [\tilde{q}(1 - 2\lambda) + (1 - \tilde{q})], \end{aligned}$$

which together with the symmetric condition for an uninformed LEFT implies that $(1 - \gamma)(\tilde{q} - q_L) \geq \gamma(q_R - q_M)$ and $(1 - \gamma)(q_R - \tilde{q}) \geq \gamma(q_M - q_L)$, and thus

$$(1 - \gamma)(q_R - q_L) \geq \gamma(q_R - q_L).$$

The above is inconsistent with $\gamma > 1/2$ unless $q_R - q_L = 0$, and it follows that $q_R = q_M = q_L$ when $\gamma > 1/2$ in any incentive compatible outcome. There is no incentive compatible mechanism that Pareto dominates a coin flip.

2.2. Delay mechanisms

As suggested in our companion paper (Damiano, Li and Suen 2012), delay in making decisions can improve information aggregation and ex ante welfare in the absence of side transfers. We model delay by an additive payoff loss to the players, and denote it as $\delta_1 \geq 0$. Properly employed by a mechanism designer, delay helps improve information aggregation by “punishing” the uninformed player when he acts like the informed. Suppose we modify the voting game in subsection 2.1 by adding delay: when both players vote for their favorite alternatives, a delay δ_1 is imposed on the players before the decision is made by flipping a coin. It is straightforward to show that this modified game achieves the first-best outcome of Pareto efficient decision without the players having to incur delay. More precisely, for any $\gamma_1 > 1/2$ and $\delta_1 \geq \lambda(2\gamma_1 - 1)/(1 - \gamma_1)$, the unique equilibrium in the modified voting game is that the informed votes for his favorite alternative while the uninformed votes for his opponent’s favorite alternative.

Using delay to improve information aggregation in committees is both natural and, as a mechanism, simple to implement. However, as a form of collective punishment, such “delay mechanism” requires commitment. Furthermore, if the degree of conflict becomes larger, that is, if γ_1 becomes greater, the amount of delay required to achieve first best increases without bound. In this paper, we assume that there is “limited commitment” in the sense that the amount of delay δ_1 is bounded from above by some exogenous positive parameter Δ . Throughout the paper, we assume that

$$0 < \Delta < \lambda \tag{1}$$

This is admittedly a crude way of modeling the constraint on commitment power: the destruction of value on and off the equilibrium path is unlikely to be credible unless the amount involved is small relative to the decision at stake.

Of course, whether the bound Δ is binding or not depends on the initial degree of conflict γ_1 . Throughout the paper, we assume that $\Delta < \lambda(2\gamma_1 - 1)/(1 - \gamma_1)$, so that the first-best outcome cannot be achieved through a delay mechanism with $\delta_1 \leq \Delta$. Equivalently, this assumption can be written as:

$$\gamma_1 > \gamma_* \equiv \frac{\lambda + \Delta}{2\lambda + \Delta}. \tag{2}$$

Note that $\gamma_* > 1/2$. Under assumption 2, using delay to achieve the “second best” leads to a trade-off because a greater δ_1 makes the uninformed more willing to vote for his

opponent's favorite alternative but it raises the payoff loss whenever delay occurs.

It is straightforward to show that it is optimal to set $\delta_1 = \Delta$ if the initial degree of conflict γ_1 is not too high. In the unique equilibrium of the voting game the informed player always votes for his favorite alternative while the uninformed player votes for his favorite alternative with some probability x_1 , with x_1 decreasing in δ_1 . To see why the trade-off must result in a corner solution, note that whenever the uninformed player is randomizing ($x_1 < 1$), his payoff can be given by the payoff from voting the opponent's favorite:

$$\gamma_1 [x_1(1 - 2\lambda) + (1 - x_1)(1 - \lambda)] + 1 - \gamma_1.$$

Raising δ_1 lowers x_1 and therefore benefits the uninformed. Furthermore the same reduction in x_1 benefits an informed player more than it benefits an uninformed player, because a lower x_1 always improves the quality of the decision from the informed player's perspective, but it may or may not improve the quality of the decision for the uninformed player depending on the state. As a result, whenever it is desirable to have $x_1 < 1$, the optimal mechanism is to set δ_1 equal to the upper bound to induce the lowest possible x_1 .

When the initial belief γ_1 is close to γ_* , the one-round delay mechanism with $\delta_1 = \Delta$ implements an equilibrium outcome that approximates the first best. As γ_1 increases, however, the uninformed player votes for his favorite alternative with a greater probability, which leads to a greater payoff loss to the informed as well as to the uninformed player due to delay. At some γ_1 , the benefit of inducing the uninformed player to vote for the opponent's favorite alternative some of the time is exactly offset by the payoff loss due to delay. So at this belief a one-round delay mechanism with $\delta_1 = \Delta$ gives the same ex ante payoff to the players as simply flipping a coin (i.e., $\delta_1 = 0$). However we can show that a two-round mechanism can do strictly better in this case.

Imagine that we modify the original one-round mechanism by replacing the coin flip outcome after delay with a second one-round mechanism with some delay $\delta_2 \leq \Delta$. Suppose that we can choose δ_2 such that the uninformed player obtains the coin-flip payoff of $1 - \lambda$ from equilibrium randomization $x_2 < 1$ in the continuation round. Then, it remains an equilibrium for the uninformed player to vote for his favorite alternative in the first round with the same probability x_1 as in the original one-round mechanism. As both x_1 and the continuation payoff $1 - \lambda$ remain unchanged in the modified mechanism, the payoff to the uninformed player does not change. Furthermore, a small x_2 benefits the informed player more than it benefits the uninformed player. Whenever the uninformed is indifferent between a continuation round with $x_2 < 1$ and $\delta_2 > 0$ and no continuation round ($x_2 = 1$ and $\delta_2 = 0$), the informed is strictly better off with the former than with

the latter. Thus, this two-round mechanism delivers the same payoff to the uninformed player but improves the payoff of the informed player.

In the above two-round mechanism we have assumed that there is a continuation round with delay δ_2 such that the uninformed player would get the coin-flip payoff of $1 - \lambda$ in the equilibrium sub-game. This is a valid assumption if two conditions are met by the “continuation belief” γ_2 (i.e., the updated belief of the uninformed player that the state is the conflict state M at the end of the original one-round mechanism with delay Δ). First, the initial belief γ_1 is not too close to γ_* , so that x_1 is sufficiently large to give an updated belief γ_2 which is not too low. If γ_2 is too low, there is no way to support the coin-flip payoff of $1 - \lambda$ through randomization by the uninformed player in the continuation round.⁵ This is precisely the case where the original one-round delay mechanism performs poorly relative to a coin flip. Second, the initial degree of conflict γ_1 is not too high, so that the updated belief γ_2 is not too high and the required δ_2 in the continuation round does not exceed the bound Δ . This assumption is satisfied, for example, if the initial belief γ_1 is such that the original one-round mechanism gives the same payoff as $1 - \lambda$. At this γ_1 , a delay $\delta_1 = \Delta$ induces $x_1 < 1$ in the original one-round mechanism. Bayesian updating implies that $\gamma_2 < \gamma_1$. Therefore to induce $x_2 < 1$ in the continuation round would require a delay $\delta_2 < \Delta$.

That a two-round mechanism can do better than a one-round mechanism raises the question about what a general dynamic delay mechanism can achieve. A *delay mechanism* is formally defined as a sequence of delays, $(\delta_1, \dots, \delta_T)$, with each $\delta_t \in [0, \Delta]$. We allow T to be finite or infinite. Any delay mechanism defines an extensive game form. In each round $t < T$ conditional on the game having not ended, each player chooses between voting for his favorite alternative (*persist*) and voting against it (*concede*). If the two votes agree, the agreed alternative is implemented immediately and the game ends. If both players vote for their opponent’s favorite alternative (we call this a *reverse disagreement*), the decision is made by a coin flip without delay. If both vote for their own favorite (*regular disagreement*), the delay δ_t is imposed and the game moves on to the next round. If T is finite, the deadline round T differs from previous rounds only in that following a regular disagreement, the decision is made by a coin flip after the delay δ_T is imposed, which ends the game. To ensure that the game is well-defined for T infinite, we adopt the convention that the payoff to each player from never implementing an alternative is zero. Correspondingly, we normalize the payoffs by assuming that $\lambda < \frac{1}{2}$ so that it

⁵In the continuation round, if $x_2 < 1$ then the payoff to the uninformed is $\gamma_2[x_2(1 - 2\lambda) + (1 - x_2)(1 - \lambda)] + 1 - \gamma_2$. For this to equal $1 - \lambda$ with $x_2 < 1$, we need $\gamma_2 > 1/2$.

is not an equilibrium for the two players to persist forever even if the infinite sequence of delays is identically zero.

3. Results

Given the initial degree of conflict γ_1 and the upper-bound on delay Δ , we say that a delay mechanism, together with a symmetric perfect Bayesian equilibrium in the extensive-form game defined by the mechanism, is *optimal* if there is no delay mechanism with a symmetric perfect Bayesian equilibrium that gives a strictly higher ex ante payoff to each player. This definition of optimality allows for multiple symmetric perfect Bayesian equilibria in a given delay mechanism. When there is no danger of confusion, we will use “equilibrium” instead of symmetric perfect Bayesian equilibrium, and apply optimality and sub-optimality to the delay mechanism itself instead of to the combination of mechanism and equilibrium. We provide a characterization of optimal delay mechanisms and highlight their main properties in this section.

For any $\gamma_1 \in (\gamma_*, 1)$, let $r_*(\gamma_1)$ be the smallest integer r satisfying

$$\left(\frac{\lambda + \Delta}{\lambda}\right)^r \geq \frac{1 - \gamma_*}{1 - \gamma_1}. \quad (3)$$

We shall see in section 4.1 below that $r_*(\gamma_1)$ can be interpreted as the minimum number of “updating rounds” for the posterior belief to reach from γ_1 to γ_* . To avoid clutter, we sometimes write r_* to stand for $r_*(\gamma_1)$ when there is no danger of confusion. Define the “residue” η such that

$$\eta \left(\frac{\lambda + \Delta}{\lambda}\right)^{r_* - 1} = \frac{1 - \gamma_*}{1 - \gamma_1}. \quad (4)$$

By definition, $\eta \in (1, (\lambda + \Delta)/\lambda]$. When η is equal to its upper bound, (3) holds as an equation.

Main Result. *There exists $\underline{\gamma}$ and $\bar{\gamma}$, with $\gamma_* < \underline{\gamma} < \bar{\gamma} < 1$, such that*

(a) *for $\gamma_1 \leq \underline{\gamma}$, an optimal delay mechanism is the one-round mechanism with maximum delay:*

$$\delta_1 = \Delta, \quad \text{with } T = 1;$$

(b) *for $\gamma_1 \in (\underline{\gamma}, \bar{\gamma})$ such that $r_*(\gamma_1) = 1$, an optimal delay mechanism is given by:*

$$\begin{aligned} \delta_1 &= \Delta, \quad \delta_2 = \dots = \delta_{\tau+1} = (\gamma_1 \lambda - (1 - \gamma_1) \Delta) / \tau, \\ \delta_{2+\tau} &= \Delta, \quad \text{with } T = 2 + \tau, \end{aligned}$$

where τ is the smallest integer greater than or equal to $\gamma_1(\lambda + \Delta)/\Delta - 1$;

(c) for $\gamma_1 \in (\underline{\gamma}, \bar{\gamma})$ such that $r_*(\gamma_1) \geq 2$, an optimal delay mechanism is given by:

$$\begin{aligned} \delta_{2t-1} &= \Delta, & \delta_{2t} &= \lambda\Delta/(\lambda + \Delta), & \text{for } t &= 1, 2, \dots, r_* - 2, \\ \delta_{2(r_*-1)-1} &= \Delta, & \delta_{2(r_*-1)} &= \lambda(\eta - 1)/\eta, \\ \delta_{2r_*-1} &= \lambda(\eta - 1), & \delta_{2r_*} &= \delta_{2r_*+1} = \dots = \delta_{2r_*+\tau-1} = \gamma_*\lambda/\tau, \\ \delta_{2r_*+\tau} &= \Delta, & \text{with } T &= 2r_* + \tau; \end{aligned}$$

where τ is the smallest integer greater than or equal to $\gamma_*\lambda/\Delta$;

(d) for $\gamma_1 \geq \bar{\gamma}$, an optimal delay mechanism is flipping a coin:

$$\delta_1 = 0, \quad \text{with } T = 1.$$

Furthermore, the set of γ_1 for which case (b) applies is non-empty, and the set of γ_1 for which case (c) applies is non-empty if and only if $\Delta < (\sqrt{2} - 1)\lambda$.

The above characterization establishes that it is optimal to have a dynamic delay mechanism so long as the initial degree of conflict γ_1 is intermediate, that is, between $\underline{\gamma}$ and $\bar{\gamma}$. Otherwise, it is optimal to have either a one-round delay mechanism with maximum delay Δ is optimal when γ_1 is too close to γ_* (case (a) of Main Result), or a coin flip without delay when γ_1 is too close to 1 (case (d) of Main Result). When optimal delay mechanisms are dynamic (i.e, cases (b) and (c)), they induce intuitive properties of equilibrium play which highlight the logic of using delays dynamically to facilitate strategic information aggregation under the limited commitment constraint. We highlight the most interesting features of optimal delay mechanisms below:

- (i) Any optimal delay mechanism is finite with a *deadline* T .
- (ii) Any optimal dynamic delay mechanism induces *deadline play* with the *efficient deadline belief* in equilibrium: there is t with $2 \leq t \leq T - 1$ such that the uninformed player persists with probability one in round $t, \dots, T - 1$, and his belief γ_T entering the last round less than or equal to γ_* , so that the Pareto efficient decision is made at the deadline.
- (iii) Any optimal dynamic delay mechanism induces *stop-and-start* in equilibrium: for any two adjacent rounds before the deadline play, the uninformed player persists with probability one in one of them and randomizes in the other, starting with randomization in the first round.

Property (i) implies that in any equilibrium induced by an optimal mechanism the to-

tal delay is bounded from above by $T\Delta$. Intuitively, it follows from the optimality of the delay mechanism that there is a bound on the total delay such that an uninformed player concedes with probability one before it is incurred. However, this argument relies on the claim that an informed player persists with probability one regardless of the history of the play. We establish this claim and property (i) simultaneously in Section 4.5.⁶ Not surprisingly, the intuition behind the claim is that an informed player has a stronger incentive to persist with his favorite alternative than an uninformed player does.

Property (ii) implies that an optimal delay mechanism generally induces a “cooling-off” period before the deadline during which no attempt is made to reach a decision (all types of players persist with their favorite alternatives). The logic behind such deadline play is that neither too little nor too much concession by the uninformed player before the deadline arrives can be optimal. After explaining our methodology and presenting some preliminary results in Section 4.1, we show in Section 4.2 that any optimal delay mechanism induces the deadline belief γ_T of the uninformed player that is less than or equal to γ_* . A higher deadline belief would imply that the uninformed player would not concede with probability one in the last round, and as a result the Pareto efficient decision could not be achieved at the end. When $r_*(\gamma_1) \geq 2$ (i.e., case (c)), we show that the deadline belief γ_T must be exactly equal to γ_* . A lower deadline belief would imply that the uninformed player would concede in the last round even if in the last round the limited commitment bound is slack, and so the delays before the deadline can be reduced while still guaranteeing the Pareto efficient decision at the end.

Property (iii) refers to case (c) of the Main Result and is perhaps the most interesting insight of this paper. This property will be established in Section 4.3 below. When the upper bound on day is not too large ($\Delta < (\sqrt{2} - 1)\lambda$), it would take at least two “updating rounds” for the deadline belief γ_T to reach γ_* . It turns out that in this case inducing the uninformed player to concede with a positive probability in each round is not optimal. The logic behind this is that having the uninformed player randomize in the next round in fact fails to minimize the probability that the uninformed persists in the present round. This is because in the present round the probability of the uninformed player persisting is increasing in the expected payoff he obtains in the next round. If the probability of the uninformed player persisting in the next round is increased, then he will expect a lower payoff from the next round and is thus induced to persist with a lower probability in the present round. This is why an optimal mechanism has a stop-and-start feature: having

⁶The proof is relegated to the end of our main analysis section while we derive properties (ii) and (iii) assuming that (i) holds, because property (i) is less central to the insights of this paper.

the players stop conceding in the follow-on round provides a greater punishment in the event of a regular disagreement that induces them to concede more in the current round. Furthermore, the Main Result states that an optimal mechanism cannot have a delay equal to the upper bound in every round. Instead it is optimal to have delays which alternate between the maximum bound (to induce “start”) and a level strictly below the bound (to induce “stop”).

4. Analysis

4.1. Preliminary results

Until Section 4.5, we consider only finite delay mechanisms and restrict our analysis to symmetric perfect Bayesian equilibria in which the informed players always persist. Denote x_t as the equilibrium probability that the uninformed players persist in round t in a game induced by the mechanism $(\delta_1, \dots, \delta_T)$. Given the rules of our dynamic delay mechanism, the game would end immediately at round t whenever $x_t = 0$. Let γ_t be the equilibrium belief of the uninformed player that his opponent is uninformed at the beginning of round t . Given the initial belief γ_1 , the belief in subsequent rounds is derived from Bayes’ rule:

$$\gamma_{t+1} = \frac{\gamma_t x_t}{\gamma_t x_t + 1 - \gamma_t}. \quad (5)$$

We call γ_T the *deadline belief* of the game. Finally, we denote as U_t the equilibrium expected payoff of an uninformed player at the beginning of round t . This payoff is given by:

$$U_t = \begin{cases} \gamma_t [x_t(-\delta_t + U_{t+1}) + 1 - x_t] + (1 - \gamma_t)(-\delta_t + U_{t+1}) & \text{if } x_t > 0, \\ \gamma_t [x_t(1 - 2\lambda) + (1 - x_t)(1 - \lambda)] + 1 - \gamma_t & \text{if } x_t < 1. \end{cases} \quad (6)$$

In the above, the top expression is the payoff from persisting and the bottom expression is the payoff from conceding. The uninformed player is indifferent between these actions when $x_t \in (0, 1)$. We often write $U_t(\gamma_t)$ to acknowledge the relation between U_t and γ_t . We denote as V_t the equilibrium expected payoff of an informed player at the beginning of round t , and we write $V_t(\gamma_t)$ even though the informed player knows the state. The ex ante payoff of each player, before they learn their types, is given

$$W_1(\gamma_1) = \frac{1}{2 - \gamma_1} U_1(\gamma_1) + \frac{1 - \gamma_1}{2 - \gamma_1} V_1(\gamma_1). \quad (7)$$

An optimal delay mechanism is one that maximizes $W_1(\gamma_1)$ subject to the constraint that $\delta_t \leq \Delta$ for all t .

Given a delay mechanism $(\delta_1, \dots, \delta_T)$, an equilibrium of the induced game can be characterized by a sequence $\{\gamma_t, x_t, U_t\}_{t=1}^T$ that satisfies (5) and (6). The “boundary conditions” are provided by the initial belief γ_1 , and by the continuation payoff $-\delta_T + 1 - \lambda$ in the event that the players fail to reach an agreement at the last round T . Although it is possible to solve the equilibrium for some particular delay mechanism (such as one with constant delay), characterizing all equilibria for any given mechanism is neither feasible nor insightful. Instead we introduce a “local variation method” to derive necessary conditions on an equilibrium induced by an optimal delay mechanism. It turns out that these necessary conditions are sufficient to provide a full characterization of optimal mechanisms that gives our Main Result.

One interesting observation about our model is that equilibrium analysis depends only on the incentives of the uninformed players alone. We show that whenever the uninformed player is indifferent between persisting and conceding, the informed player strictly prefers to persist. Hence the incentive constraints for the informed player is not binding and does not play a part in the equilibrium analysis. The analysis of optimal mechanisms, however, requires studying the payoffs to both uninformed players and informed players. Therefore, in order to perform welfare analysis, it is necessary to link the equilibrium payoff of the informed $V_1(\gamma_1)$ to that of uninformed $U_1(\gamma_1)$. The following result proves to be important.

Lemma 1. (LINKAGE LEMMA) *Suppose a mechanism induces an equilibrium in which it is a weak best response for an uninformed player to persist through to the last round T from some round $t \leq T$ onward. Then*

$$U_t(\gamma_t) = \gamma_t V_t(\gamma_t) + (1 - \gamma_t) \left(1 - \lambda - \sum_{s=t}^T \delta_s \right). \quad (8)$$

Proof. Suppose an uninformed player persists in each round from round t onwards. With probability γ_t , his opponent is an uninformed player who persists with probability x_s for $s = t, \dots, T$. In this case his payoff would be identical to that of an informed player facing an uninformed opponent, who uses the same strategy as his own. With probability $1 - \gamma_t$, his opponent is an informed player who persists in every round. In this case his payoff would be $1 - \lambda - \sum_{s=t}^T \delta_s$. Since persisting from round t onwards is a weak best response,

the uninformed player's payoff at round t is given by equation (8). ■

Although simple, the Linkage Lemma has an important implication. By equation (8), if raising the total delay $\sum_{t=1}^T \delta_t$ does not lower $U_1(\gamma_1)$, and if persisting in each round remains a weak best response, then such a change strictly increases $V_1(\gamma_1)$ and hence the ex ante payoff $W_1(\gamma_1)$. The logic is that a greater delay keeps the expected payoff of an uninformed player unchanged only if it induces him to lower the probabilities of persisting. Since an informed player faces an uninformed opponent with probability 1, while an uninformed player faces an uninformed opponent with probability $\gamma_1 < 1$, the same reduction in probabilities of persisting by the opponent benefits an informed player by more than it benefits an uninformed player. We have already seen in section 2.2 that this logic leads to the conclusion that an optimal one-round mechanism has $\delta_1 = \Delta$. It remains useful when we use the local variation method to study general dynamic mechanisms.

The Linkage Lemma suggests that increasing the total delay may improve the ex ante payoff (7) of each player. Although each δ_t in a delay mechanism is bounded from above by Δ , the total delay can be increased by inserting additional rounds of delay. This is frequently used in our local variation approach.

Suppose that, in equilibrium of a delay mechanism, there are r rounds of randomization (with $x_t \in (0, 1)$) by the uninformed player prior to round T . Let $t(1) < \dots < t(r)$ represent such rounds. We call these *updating rounds* because, by Bayes' rule (5), the posterior belief that the state is M decreases after a regular disagreement in an updating round. On the other hand, if there is a round t before T such that $x_t = 1$, we call it a *stopped round* because the posterior belief does not change upon a regular disagreement. Clearly, the first round of an optimal delay mechanism must be an updating round (i.e., $t(1) = 1$), for otherwise the mechanism would be wasting time prior to round $t(1)$. When needed, we define $t(r+1)$ to be equal to T . For convenience of notation, we let

$$\sigma_{t(i)} \equiv \sum_{t=t(i)}^{t(i+1)-1} \delta_t$$

for each $i = 1, \dots, r$ be the *effective delay* in round $t(i)$.

Definition 1. *There is no slack in an updating round $t < T$ if $\delta_t = \Delta$ and $U_{t+1}(\gamma_{t+1}) = 1 - 2\gamma_{t+1}\lambda$. There is slack in updating round t if $\delta_t < \Delta$ or $U_{t+1}(\gamma_{t+1}) > 1 - 2\gamma_{t+1}\lambda$.*

Given fixed γ_t and the upper bound Δ on the delay in each round, the next result

shows that lowering x_t is feasible if and only if there is slack in round t . We can change an updating round t with slack to an updating round with no slack by raising δ_t if $\delta_t < \Delta$. If $U_{t+1}(\gamma_{t+1}) > 1 - 2\gamma_{t+1}\lambda$, we can increase delay by inserting one additional round s with delay δ_s between t and $t + 1$ and raising δ_s while maintaining $x_s = 1$ as long as $U_{t+1}(\gamma_{t+1}) - \delta_s \geq 1 - 2\gamma_{t+1}\lambda$. If the requisite δ_s to fill the slack exceeds the bound Δ , we can keep inserting additional rounds between t and $t + 1$ while ensuring the uninformed players persist in these extra rounds until there is no slack.

Lemma 2. (MAXIMAL CONCESSION LEMMA) *For any $t < T$ and $\gamma_t > \Delta/(\lambda + \Delta)$, the lowest feasible x_t is given by*

$$\chi(\gamma_t) \equiv \frac{\gamma_t\lambda - (1 - \gamma_t)\Delta}{\gamma_t(\lambda + \Delta)},$$

and the lowest feasible γ_{t+1} is given by

$$g(\gamma_t) \equiv \frac{\gamma_t\lambda - (1 - \gamma_t)\Delta}{\lambda}.$$

These lower bounds are attained if and only if there is no slack in round t .

Proof. If an uninformed player is randomizing at round t , the indifference condition (6) requires:

$$1 - \gamma_t\lambda - \gamma_t x_t \lambda = (\gamma_t x_t + 1 - \gamma_t)(-\delta_t + U_{t+1}(\gamma_{t+1})) + \gamma_t(1 - x_t).$$

At round $t + 1$, the uninformed player can guarantee a payoff of at least $1 - \gamma_{t+1}\lambda - \gamma_{t+1}x_{t+1}\lambda$ by conceding with probability one. Therefore, $U_{t+1}(\gamma_{t+1}) \geq 1 - 2\gamma_{t+1}\lambda$, with equality when $x_{t+1} = 1$. Using this bound and the bound on δ_t , we obtain

$$\gamma_t\lambda + \gamma_t x_t \lambda \leq (\gamma_t x_t + 1 - \gamma_t)(\Delta + 2\gamma_{t+1}\lambda),$$

with equality when there is no slack at round t . Using Bayes' rule for γ_{t+1} and solving for x_t , we obtain $x_t \geq \chi(\gamma_t)$. Since γ_{t+1} is increasing in x_t , plugging in the lowest value of x_t and using Bayes' rule give $\gamma_{t+1} \geq g(\gamma_t)$. ■

For fixed belief γ_t of the uniformed player, the probability x_t of him persisting in round t depends both on the round t delay δ_t and on the continuation payoff $U_{t+1}(\gamma_{t+1})$. In general, x_t is not necessarily decreasing in δ_t or increasing in $U_{t+1}(\gamma_{t+1})$ separately. Nonetheless, the Maximal Concession Lemma establishes a lower bound on x_t , and hence

a lower bound on the updated belief γ_{t+1} , which binds when both δ_t is maximized and $U_{t+1}(\gamma_{t+1})$ is minimized.

When there is maximal concession by the uninformed player, his belief that the state is M evolves according to

$$\frac{1 - g(\gamma_t)}{1 - \gamma_t} = \frac{\lambda + \Delta}{\lambda}. \quad (9)$$

Comparing this to equation (3), we see that it takes at least $r_*(\gamma_1)$ updating rounds for the belief to reach from γ_1 to γ_* , if the uninformed is making maximal concessions in each of these updating rounds.⁷ A tighter commitment bound Δ would mean that it requires more updating rounds for the initial degree of conflict γ_1 to reduce to the level γ_* , when the conflict can be efficiently resolved.

4.2. Efficient deadline belief

First we show that in any optimal delay mechanism the Pareto efficient decision is made with probability one. This is clearly the case if there is some round $N < T$ such that $x_N = 0$ given that the informed player always persists. If such round N does not exist, then we must have $x_T = 0$ in the last round for the decision to be Pareto efficient.

We prove this result using a local variation method. Suppose to the contrary that $x_T > 0$. We show that a modified mechanism which induces a smaller equilibrium deadline belief $\tilde{\gamma}_T$ will increase the ex ante payoff. Hence the original mechanism cannot not optimal.

Suppose there are r updating rounds under the original mechanism. A smaller deadline belief $\tilde{\gamma}_T$ can be obtained by introducing an extra updating round s , after round $t(r)$ but before round T , with an appropriately chosen delay $\tilde{\delta}_s$ to induce $\tilde{x}_s < 1$. The difficulty is that the equilibrium sequence $\{\gamma_t, x_t, U_t\}$ is jointly determined through (5) and (6). For example, introducing the extra round s to induce $\tilde{x}_s < 1$ would change the continuation payoffs at round $t(r)$ and before. Our local variation method bypasses this difficulty by introducing yet another extra round s' in between round $t(r)$ and round s . The delay $\tilde{\delta}_{s'}$ for this round is chosen in such a way to keep the continuation payoff in the event of a disagreement at round $t(r)$ fixed at the original value of

$$-\delta_{t(r)} + U_{t(r)+1}(\gamma_{t(r)+1}).$$

⁷Define $g^{(n)}(\gamma)$ be such that $g^{(1)}(\gamma) = g(\gamma)$ and $g^{(n)}(\gamma) = g(g^{(n-1)}(\gamma))$. Then $r_*(\gamma)$ is the smallest integer r such that $g^{(r)}(\gamma) \leq \gamma_*$.

Because both the initial belief γ_1 and the continuation payoff at round $t(r)$ are fixed, if $\{\gamma_t, x_t, U_t\}_{t=1}^{t(r)}$ is part of equilibrium under the original mechanism, then the same sequence constitutes part of equilibrium under the modified mechanism.⁸ In particular, the modified mechanism does not affect $U_1(\gamma_1)$. Furthermore, as long as $x_T > 0$, we can still have $\tilde{x}_T \geq 0$ in the modified mechanism after marginally lowering the deadline belief to $\tilde{\gamma}_T$. This means that persisting through to the end remains a weak best response for the uninformed player given the equilibrium strategy in the modified mechanism.

In this construction, to lower the probability of persistence from $x_s = 1$ to $\tilde{x}_s < 1$ requires raising the delay $\tilde{\delta}_s$ at round s . When \tilde{x}_s becomes smaller, the uninformed player's payoff increases, and therefore the effective delay $\tilde{\delta}_{s'}$ must also rise to keep the continuation payoff for round $t(r)$ fixed. As a result the total delay in the modified mechanism is higher than that in the original mechanism. It then follows from the Linkage Lemma that the ex ante payoff of the players must increase. The details of this construction are relegated to the Appendix.

Proposition 1. *In an optimal mechanism, $x_T = 0$.*

An immediate corollary to Proposition 1 is that the deadline belief must satisfy $\gamma_T \leq \gamma_*$. The intuition behind this result is that if the deadline belief is not sufficiently low to induce the Pareto efficient decision at the final round, then it is possible slightly modify the mechanism so as to increase the total delay without affecting the expected payoff of the uninformed. This is accomplished by inducing the uninformed to play another round of randomization before the final round. Since an informed player benefits more from concession by the uninformed than an uninformed player does, the Linkage Lemma implies that the ex ante payoff can be improved.

The next result establishes a counterpoint to Proposition 1. Although an optimal delay mechanism would drive down the degree of conflict from γ_1 to $\gamma_T \leq \gamma_*$, so that the Pareto efficient decision becomes achievable at the final round, it is generally too costly to use delay to drive γ_T to a level too much below what is needed to achieve efficiency. Specifically, if γ_T is strictly below γ_* then the payoff can be improved by ending the game "earlier," that is, by reducing the total delay. Thus, in an optimal mechanism the deadline belief of the uninformed player is *efficient*, in the sense that it is the highest possible belief that would ensure the Pareto efficient decision is made with probability one.

⁸In constructing this modified mechanism, we take the belief at the beginning of round s to be fixed at the original value of $\gamma_{t(r)+1}$ and find the $\tilde{\delta}_s$ that would induce $\tilde{x}_s < 1$. This is justified because the modified mechanism does not change the equilibrium play prior to round s , ensuring that the posterior belief at the beginning of round s is indeed fixed.

The proof of the proposition below makes the assumption that $\gamma_1 > g^{-1}(\gamma_*)$. This is equivalent to $r_*(\gamma_1) \geq 2$, so that there must be at least two updating rounds for the deadline belief to reach a level strictly below γ_* . This assumption is necessary because our local variation method relies on keeping unchanged the continuation payoff at the next-to-last updating round. If there is no such round, this method does not apply.⁹

Proposition 2. *Suppose $r_*(\gamma_1) \geq 2$. Then in an optimal mechanism, $\gamma_T = \gamma_*$.*

We outline the local variation argument used to establish Proposition 2; details of the calculations are given in the Appendix. Suppose by way of contradiction that $\gamma_T < \gamma_*$. Consider the belief at the last updating round $t(r)$. There are two possibilities: (i) $\gamma_{t(r)} \leq \gamma_*$; or (ii) $\gamma_{t(r)} > \gamma_*$.

In case (i), we show that the ex ante payoff can be improved by letting the game end with deadline belief $\gamma_{t(r)}$ instead of γ_T , through the following modification to the original mechanism:

1. Change $\sigma_{t(r)}$ in such a way to make the uninformed just indifferent between persisting and conceding at $\tilde{x}_{t(r)} = 1$. This is achieved by setting $\tilde{\sigma}_{t(r)} = \lambda\gamma_{t(r)}$.
2. Change $\sigma_{t(r-1)}$ in such a way to keep the continuation payoff for the uninformed at round $t(r-1)$ fixed. This is achieved by setting $\tilde{\sigma}_{t(r-1)} = \sigma_{t(r-1)} - \lambda\gamma_{t(r)}(1 - x_{t(r)})$.

In case (ii), we modify the original mechanism so that the game ends with deadline belief γ_* instead of γ_T . The approach is to choose $\tilde{x}_{t(r)}$ such that the posterior belief satisfies

$$\frac{\gamma_{t(r)}\tilde{x}_{t(r)}}{\gamma_{t(r)}\tilde{x}_{t(r)} + 1 - \gamma_{t(r)}} = \gamma_*,$$

and calculate the requisite effective delay needed to achieve this, while maintaining the continuation payoff for the uninformed at round $t(r-1)$ constant.

Because the continuation value at round $t(r-1)$ is kept constant, the equilibrium play up to and including round $t(r-1)$ remains part of equilibrium in the modified mechanism. Therefore $U_1(\gamma_1)$ is unchanged. However, since $\gamma_T < \gamma_*$, persistence is *not* a best response for the uninformed player at round T . This means we cannot use the Linkage Lemma. Instead, we calculate the payoff to the informed player at round $t(r-1)$, and show that $\tilde{V}_{t(r-1)} > V_{t(r-1)}$ in both cases (i) and (ii). By the recursion,

$$V_t(\gamma_t) = x_t [-\delta_t + V_{t+1}(\gamma_{t+1})] + 1 - x_t,$$

⁹When $r_*(\gamma_1) = 1$ the optimal delay mechanism is given by case (a) or (b) of the Main Result. In case (b), we have $\gamma_T < \gamma_*$ under an optimal delay mechanism. But even in this case, γ_T cannot be below $g(\gamma_*)$.

and by the fact that the sequence $\{\gamma_t, x_t, U_t\}_{t=1}^{t(r-1)}$ remains part of equilibrium after the modification, we conclude that $\tilde{V}_1(\gamma_1) > V_1(\gamma_1)$. So the original mechanism cannot be optimal.

In addition to its inherent interest, Proposition 2 is useful because it is a weak best response to persist throughout the game if $\gamma_T = \gamma_*$. This means that in our analysis of optimal mechanism, we can apply the Linkage Lemma to facilitate comparison of payoff to the informed whenever $r_*(\gamma_1) \geq 2$.

4.3. Stop-and-start

In this subsection, we provide a series of lemmas that characterize the presence or absence of slack in an optimal mechanism, leading to the main characterization result of Proposition 3. Throughout this series of lemmas, we maintain the assumption that $r_*(\gamma_1) \geq 2$, or equivalently, $\gamma_1 > g^{-1}(\gamma_*)$. The key result is the following lemma.

Lemma 3. *A mechanism with slack in both round $t(i)$ and round $t(i+1)$ ($i = 2, \dots, r-1$) is not optimal.*

The local variation argument leading to Lemma 3 is slightly more involved than before, because we have to make sure not only that the equilibrium play prior to round $t(i)$ remains unchanged, but also that the equilibrium play subsequent to round $t(i+1)$ is unchanged as well. We consider the following modifications to a mechanism for which there is slack in both round $t(i)$ and round $t(i+1)$:

1. Marginally change $x_{t(i)}$ by changing $\sigma_{t(i)}$.
2. Change $x_{t(i+1)}$ in such a way that $x_{t(i)}x_{t(i+1)}$ remains fixed. This is achieved by changing $\sigma_{t(i+1)}$.
3. Change $\sigma_{t(i-1)}$ in such a way to keep $-\sigma_{t(i-1)} + U_{t(i)}(\gamma_{t(i)})$ constant.

As before, step 3 of this construction ensures that the equilibrium sequence $\{\gamma_t, x_t, U_t\}$ from round 1 to round $t(i-1)$ is unchanged, which also implies that $\gamma_{t(i)}$ is unchanged. Step 1 changes $x_{t(i)}$, which changes $\gamma_{t(i+1)}$. But since step 2 also adjusts $x_{t(i+1)}$ in such a way to keep $x_{t(i)}x_{t(i+1)}$ constant, this means that the posterior belief $\gamma_{t(i+2)}$ under the modified mechanism is unchanged by the construction. As a result, the original equilibrium sequence $\{\gamma_t, x_t, U_t\}$ from round $t(i+2)$ to the deadline round T is also unchanged. By appealing to the Linkage Lemma, we only need to calculate the total delay between round $t(i-1)$ and round $t(i+2)$ in order to evaluate the effect of this modification on the ex ante payoff.

One can think of this local variation exercise as choosing $\gamma_{t(i+1)}$ to maximize the total

delay, while holding $\gamma_{t(i)}$ and $\gamma_{t(i+2)}$ (and the continuation payoff at round $t(i-1)$) fixed. From the Maximal Concession Lemma, the feasible set for the choice of $\gamma_{t(i+1)}$ is

$$\left[\max \left\{ \gamma_{t(i+2)}, g(\gamma_{t(i)}) \right\}, \min \left\{ \gamma_{t(i)}, g^{-1}(\gamma_{t(i+2)}) \right\} \right].$$

In the proof of Lemma 3, we show that the total delay is a convex function of $\gamma_{t(i+1)}$. When there is slack in both round $t(i)$ and round $t(i+1)$, $\gamma_{t(i+1)}$ is in the interior of the feasible set. So the mechanism cannot be optimal.

Since there cannot be slack in both round $t(i)$ and $t(i+1)$, the total delay as a function of $\gamma_{t(i+1)}$ in the problem considered above is maximized either at $\gamma_{t(i+1)} = g^{-1}(\gamma_{t(i+2)})$ (no slack at round $t(i+1)$) or at $\gamma_{t(i+1)} = g(\gamma_{t(i)})$ (no slack at round $t(i)$). It turns out that these two choices of $\gamma_{t(i+1)}$ entail the same total delay; therefore they are payoff-equivalent.

Lemma 4. *Holding $\gamma_{t(i)}$ and $\gamma_{t(i+2)}$ fixed, a mechanism with slack at round $t(i)$ but no slack at round $t(i+1)$ is payoff-equivalent to a mechanism with no slack at round $t(i)$ but slack at round $t(i+1)$.*

The next result shows that it is optimal to induce maximal concession by the informed player in the first round.

Lemma 5. *In an optimal mechanism, there is no slack in the first round.*

In the proof of this result, we assume that there is slack in the first round and consider the following modification to the original mechanism:

1. Marginally lower x_1 by raising σ_1 .
2. Raise $x_{t(2)}$ in such a way that $x_1 x_{t(2)}$ remains fixed. This is achieved by lowering $\sigma_{t(2)}$.

Step 2 ensures that the equilibrium play from round $t(2) + 1$ onward is not affected by the modification. However because we are changing x_1 , there is no prior round before round 1 to adjust for the resulting change in $U_1(\gamma_1)$. Instead we calculate the change in $U_1(\gamma_1)$ and $V_1(\gamma_1)$ directly and show that both changes are positive.

The following proposition summarizes the results from our series of lemmas.

Proposition 3. *Suppose $r_*(\gamma_1) \geq 2$. In an optimal mechanism, there can be at most one round with slack. If there is a round with slack, it cannot be the first round.*

Proof. Let $t(j)$ be the earliest round with slack and $t(j')$ be the latest round with slack. By Lemma 3, the mechanism cannot be optimal if $j' = j + 1$. Since there is slack in round $t(j)$

and no slack in round $t(j + 1)$, by Lemma 4, we can construct a mechanism for which the earliest round with slack is $t(j + 1)$ and the latest round with slack is $t(j')$, and which is payoff-equivalent to the original mechanism. If $j' = j + 2$, these two rounds with slacks will be adjacent and therefore the mechanism cannot be optimal by Lemma 3. If $j' > j + 2$, we proceed iteratively to a mechanism for which the earliest round with slack is $t(j + 2)$ and the latest round with slack is $t(j')$, and which is payoff-equivalent to the original mechanism. Since the number of rounds is finite, this construction eventually produces a mechanism with two adjacent rounds with slack, which by Lemma 3 contradicts the assumption that it is optimal. So an optimal mechanism can have at most one round with slack. Furthermore, by Lemma 5, this round (if it exists) cannot be the first round. ■

Proposition 3 establishes that all (but possibly one) updating rounds in an optimal mechanism are without slack. By the Maximal Concession Lemma, this means that the uninformed players adopt the largest feasible probabilities of conceding in these updating rounds. Moreover, since $U_{t(i)+1}(\gamma_{t(i)+1}) = 1 - 2\gamma_{t(i)+1}\lambda$ when there is no slack in round $t(i)$, this also means that the uninformed players are making no concession (i.e., $x_{t(i)+1} = 1$) right after the round when they have made the maximal concession. After that, in the next updating round $t(i + 1)$, they will make the maximal concession again provided there is no slack in round $t(i + 1)$. In this sense equilibrium play exhibits a stop-and-start pattern, alternating between maximal concession and no concession, until the deadline play kicks in after round $t(r)$, when they make no concession from round $t(r) + 1$ through round $T - 1$, followed by full concession (i.e., $x_T = 0$) at the deadline.

4.4. Optimal delay mechanism

For the time being assume that γ_1 is such that $r_*(\gamma_1) \geq 2$. When there is no slack in some round $t(i)$, the belief evolves according to $\gamma_{t(i+1)} = g(\gamma_{t(i)})$. Recall from our discussion of the Maximal Concession Lemma that the greatest extent of belief updating feasible occurs when there is no slack in an updating round. Such belief evolution is determined by equation (9); hence it takes at least $r_*(\gamma_1)$ updating rounds for belief to reach from γ_1 to γ_* . Since the initial belief γ_1 is given and the end belief γ_T must be γ_* (Proposition 2), the fact that there can be at most one round with slack (Proposition 3) implies that there are exactly $r_*(\gamma_1)$ updating rounds in an optimal mechanism.

Recall also that the “residue” η defined in (4) satisfies $1 < \eta \leq (\lambda + \Delta)/\lambda$. There are two cases to consider; the first case turns out to be a special case of the second. In the first case, $\eta = (\lambda + \Delta)/\lambda$. The belief reaches from γ_1 to γ_* with r_* rounds of randomization, with no slack in any of the r_* rounds. In the second case, $\eta < (\lambda + \Delta)/\lambda$. To satisfy

the restriction imposed by Proposition 3, there must be $r_* - 1$ rounds of randomization with no slack, and one round of randomization with slack. By Lemma 4, which round is given the slack is payoff-irrelevant; let us assume that the round with slack is r_* , the last updating round before the deadline.¹⁰

For each $i = 1, \dots, r_* - 2$, we have $\delta_{t(i)} = \Delta$. The equilibrium belief evolves according to $\gamma_{t(i+1)} = g(\gamma_{t(i)})$ and the equilibrium probability of persisting is $x_{t(i)} = \chi(\gamma_{t(i)})$. The condition that there is no slack at round $t(i)$ requires:

$$U_{t(i)+1}(\gamma_{t(i)+1}) = 1 - \gamma_{t(i)+1}\lambda - \gamma_{t(i)+1}\chi(\gamma_{t(i)+1})\lambda - \sum_{t=t(i)+1}^{t(i+1)-1} \delta_t = 1 - 2\gamma_{t(i)+1}\lambda,$$

which gives

$$\sum_{t=t(i)+1}^{t(i+1)-1} \delta_t = \frac{\lambda\Delta}{\lambda + \Delta}.$$

Note that the total delay between successive updating rounds is constant across i . How these sums are distributed across the intervening periods is immaterial. However, notice that the sum is less than Δ . Hence, the optimal mechanism can be implemented with just one intervening round between any two successive rounds of randomization. This corresponds to the first line of the delay mechanism described in case (c) of our Main Result.

For round $t(r_* - 1)$, no slack implies that $\delta_{t(r_*-1)} = \Delta$. Further, the equilibrium belief evolves according to $\gamma_{t(r_*)} = g(\gamma_{t(r_*-1)})$ and the equilibrium probability of persisting is $x_{t(r_*-1)} = \chi(\gamma_{t(r_*-1)})$. The condition that there is no slack at round $t(r_* - 1)$ requires:

$$1 - \gamma_{t(r_*)}\lambda - \gamma_{t(r_*)}x_{t(r_*-1)}\lambda - \sum_{t=t(r_*-1)+1}^{t(r_*)-1} \delta_t = 1 - 2\gamma_{t(r_*)}\lambda. \quad (10)$$

To solve this equation, note that $x_{t(r_*)}$ must be such that the posterior belief after round $t(r_*)$ is equal to γ_* . Using this and the fact that $(1 - \gamma_*)/(1 - \gamma_{t(r_*)}) = \eta$, we obtain:

$$x_{t(r_*)} = \frac{1 - (1 - \gamma_{t(r_*)})\eta}{\gamma_{t(r_*)}\eta}. \quad (11)$$

¹⁰Although it is optimal to choose any round $t(j)$, $j \in \{2, \dots, r_*\}$, to have slack, the choice of $j = r_*$ is special because such a mechanism would be “time-consistent.” If $j \neq r_*$, when the game reaches round $t(j)$, the mechanism would no longer be optimal because it violates Lemma 5.

Given the above expression, solving (10) gives:

$$\sum_{t=t(r_*)+1}^{t(r_*)-1} \delta_t = \frac{\lambda(\eta - 1)}{\eta}.$$

Since $\eta \leq (\lambda + \Delta)/\lambda$, the above is less than or equal to $\lambda\Delta/(\lambda + \Delta)$. So the optimal mechanism can be implemented with just one intervening round between $t(r_* - 1)$ and $t(r_*)$. This corresponds to the second line of the mechanism described in case (c) of our Main Result.

At round $t(r_*)$, there may or may not be slack depending on whether η is equal to its upper bound or not. The equilibrium probability of persisting $x_{t(r_*)}$ is given by equation (11) above, and by construction the posterior belief upon a regular disagreement would become γ_* . To find the effective delay that would implement these features, we use the fact that $U_T(\gamma_T) = 1 - \lambda\gamma_*$ and the indifference condition at round $t(r_*)$ to obtain:

$$\sum_{t=t(r_*)}^{T-1} \delta_t = \frac{\gamma_{t(r_*)}(1 - \gamma_*)\lambda}{1 - \gamma_{t(r_*)}} = \lambda(\eta - 1 + \gamma_*).$$

Although how the above effective delay is distributed is payoff-irrelevant, we can always choose $\delta_{t(r_*)} = \lambda(\eta - 1) \leq \Delta$ and

$$\sum_{t=t(r_*)+1}^{T-1} \delta_t = \gamma_*\lambda.$$

When $\eta = (\lambda + \Delta)/\lambda$ so that there is no slack in round $t(r_*)$, the above distribution is the only optimal way. The rounds $t(r_*) + 1$ through $T - 1$ constitute the deadline play in an optimal delay mechanism.¹¹ This corresponds to the third line of the mechanism described in case (c) of our Main Result.

Finally, at the last round T , since the belief is $\gamma_T = \gamma_*$, choosing $\delta_T = \Delta$ would induce the uninformed to play $x_T = 0$, which always ends the game with the Pareto efficient decision. This corresponds to the last line of the mechanism described in case (c) of our Main Result.

¹¹The effective delay can be larger than the bound Δ . In this case, the uninformed player persists for more than one round in the deadline play.

Summing over all rounds, the total delay is

$$\sum_{t=1}^T \delta_t = (r_* - 2) \left(\Delta + \frac{\lambda \Delta}{\lambda + \Delta} \right) + \left(\Delta + \frac{\lambda(\eta - 1)}{\eta} \right) + \lambda(\eta - 1 + \gamma_*) + \Delta. \quad (12)$$

In an optimal mechanism, the payoff to the uninformed is given by

$$U_1(\gamma_1) = 1 - 2\gamma_1\lambda + \frac{\lambda\Delta}{\lambda + \Delta}. \quad (13)$$

The payoff to the informed player $V_1(\gamma_1)$ can be obtained using equation (8).

The above analysis completes the derivation of the optimal delay mechanism assuming that $r_*(\gamma_1) \geq 2$ and that the ex ante payoff $W_1(\gamma_1)$ is greater than the coin flip payoff $1 - \lambda$. The next lemma deals with the case of small γ_1 , which corresponds to case (b) of our Main Result.

Lemma 6. *Suppose $r_*(\gamma_1) = 1$. If $T \geq 2$ in an optimal delay mechanism, then round 1 is the only updating round before the final round T and there is no slack in round 1.*

Proof. First, suppose the mechanism induces two or more updating rounds prior to the final round. Then logic of the proof of Proposition 2 still applies. This means that we can assume without loss of generality that the deadline belief is γ_* if there are two or more rounds of randomization. Note further that for such a mechanism, the logic of the proof of Proposition 3 remains valid. This means that there is no slack in all except one of these updating rounds, which contradicts the assumption that $r_*(\gamma_1) = 1$. This implies that when we deal with an optimal mechanism for the case $\gamma_1 \leq g^{-1}(\gamma_*)$, we can without loss of generality restrict attention to mechanisms that involve only one updating round prior to the final round.

Suppose x_1 for such a mechanism induces a deadline belief $\gamma_T \leq \gamma_*$. From the indifference condition for the uninformed, the delay σ_1 associated with such a mechanism satisfies $x_1\sigma_1 = \lambda\gamma_T$. Now, consider the payoff to the uninformed. Since $U_1(\gamma_1) = 1 - \lambda\gamma_1 - \lambda\gamma_1x_1$, his payoff is maximized when x_1 is minimized. The payoff to the informed is $V_1(\gamma_1) = 1 - x_1\sigma_1 = 1 - \lambda\gamma_T$, which is maximized when x_1 is minimized. By Lemma 2, x_1 is minimized at $x_1 = \chi(\gamma_1)$, with associated posterior deadline belief $\gamma_T = g(\gamma_1)$. An optimal sequence of delays that can implement this outcome is given by $\delta_1 = \Delta$,

$$\sum_{t=2}^{T-1} \delta_t = g(\gamma_1)\lambda = \gamma_1\lambda - (1 - \gamma_1)\Delta.$$

At the final round, choosing $\delta_T = \Delta$ induces $x_T = 0$. ■

To complete the derivation of optimal delay mechanisms given in the Main Result, we compute the ex ante payoffs for the dynamic delay mechanisms (cases (b) and (c)) and compare them with the payoffs from a coin flip (case (d)) and from the one-round delay mechanism with maximum delay (case (a)). The calculations are relegated to the Appendix.

Proposition 4. *There exist $\underline{\gamma}$ and $\bar{\gamma}$, with $\gamma_* < \underline{\gamma} < \bar{\gamma} < 1$, $r_*(\underline{\gamma}) = 1$, and $r_*(\bar{\gamma}) \geq 2$ if and only if $\Delta < (\sqrt{2} - 1)\lambda$, such that the one-round delay mechanism with maximum delay is optimal if and only if $\gamma_1 \leq \underline{\gamma}$, and flipping a coin is optimal if and only if $\gamma_1 \geq \bar{\gamma}$.*

4.5. Finite delay

So far we have restricted our analysis to finite delay mechanisms, with a deadline round T . In this subsection, we allow for *infinite* delay mechanisms, that is, any infinite sequence $(\delta_t)_{t=1}^{\infty}$. We first show that any infinite delay mechanism is *effectively finite* in that there exists N such that the induced game ends with probability one before round N in any equilibrium. Due to the convention that the payoff from never implementing an alternative is zero and the normalizing assumption of $\lambda < \frac{1}{2}$, persisting forever is not an equilibrium when the infinite sequence of delays is identically zero. This means that if an infinite delay mechanism $(\delta_t)_{t=1}^{\infty}$ is not effectively finite, we can assume that $\sum_{t \geq \tau} \delta_t > 0$ for any τ .

Lemma 7. *Any infinite delay mechanism is effectively finite.*

Lemma 7 follows from a series of claims. First, any infinite mechanism with $\sum_{t \geq \tau} \delta_t$ unbounded for all τ is effectively finite, because if not, the ex ante payoff to each player would be unbounded from below. Second, under any infinite mechanism with a bounded total delay, there is no equilibrium in which the probability y_t that the informed player persists in round t is strictly less than 1 infinitely often. This is because the informed player knows that he is playing against an uninformed player and can always obtain a higher payoff by imitating the latter, so he concedes with a positive probability infinitely often only if the uninformed player almost never concedes, which we know is impossible in equilibrium. Third, any infinite mechanism with a bounded total delay is effectively finite. This follows because, with the informed player eventually always persisting in any equilibrium, the belief γ_t of the uninformed player will have to converge if the mechanism is not effectively finite, but then the uninformed in equilibrium will almost never concede, which is again impossible.

By Lemma 7, we can restrict to finite mechanisms and infinite mechanisms that are effectively finite. In either case we can apply backward induction to show that $V_t \geq U_t$ for all t in any equilibrium, which leads to the following result.

Proposition 5. *In any delay mechanism, the informed player persists with probability one in any equilibrium.*

The above proposition holds for any delay mechanism and any equilibrium. The proof in the appendix makes it clear that the argument does not invoke properties that rely on optimality of the delay mechanism under consideration. This in turn validates our approach in this section of restricting to symmetric perfect Bayesian equilibrium under finite delay mechanisms.

5. Discussion

5.1. Continuous-delay limit

We have seen in section 2.1 that in the absence of costly delay, there is no incentive compatible mechanism that Pareto dominates coin-flip in our committee model. That is, if the uniform upper bound Δ on delay is identically zero, flipping a coin is the only symmetric incentive compatible outcome. From the continuous-delay model of Damiano, Li and Suen (2012), we already know that the set of incentive compatible outcomes is discontinuous at $\Delta = 0$. In this subsection, we derive the limit of optimal delay mechanisms characterized in the present paper as Δ converges to zero, and show that it coincides with the optimal deadline in the continuous-delay model of Damiano, Li and Suen (2012) (after adjusting the differences in the payoff structure in the latter model).

As Δ goes to 0, from (13) we have

$$\lim_{\Delta \rightarrow 0} U_1(\gamma_1) = 1 - 2\lambda\gamma_1.$$

The total number of rounds of course goes to infinity, but we have

$$\lim_{\Delta \rightarrow 0} \left(1 + \frac{\Delta}{\lambda}\right)^{r_*} = \lim_{\Delta \rightarrow 0} \frac{1 - \gamma_*}{1 - \gamma_1},$$

which implies

$$\lim_{\Delta \rightarrow 0} r_*\Delta = -\lambda \log(2(1 - \gamma_1)).$$

Furthermore, since $(\lambda + \Delta)/\lambda \geq \eta \geq 1$, η goes to 1 as Δ goes to 0. Therefore, the limit of

the total delay (12) is

$$\lim_{\Delta \rightarrow 0} \sum_{t=1}^T \delta_t = -\lambda \log(2(1 - \gamma_1)) + \frac{\lambda}{2}.$$

Using equation (8) for the relationship between $U_1(\gamma_1)$ and $V_1(\gamma_1)$, we obtain:

$$\lim_{\Delta \rightarrow 0} V_1(\gamma_1) = 1 - 2\lambda - 2\lambda \frac{1 - \gamma_1}{\gamma_1} \left(\frac{3}{4} - \log(2(1 - \gamma_1)) \right).$$

For the optimal deadline mechanism we consider in Damiano, Li and Suen (2012), the value function $V^*(\gamma)$ satisfies the differential equation:

$$\frac{dV^*(\gamma)}{d\gamma} = -\frac{1 - 2\lambda\gamma - V^*(\gamma)}{\gamma(1 - \gamma)}.$$

The solution to this differential equation is

$$V^*(\gamma) = \frac{1}{\gamma} (C(1 - \gamma) + 1 - 2\lambda - 2\lambda(1 - \gamma) \log(1 - \gamma)).$$

The optimal deadline mechanism entails the boundary condition $V^*(1/2) = 1 - \lambda/2$, which gives

$$C = -(1 - 2\lambda) + 2\lambda \left(\frac{3}{4} - \log 2 \right).$$

Substituting this value of C into $V^*(\gamma)$ gives $V^*(\gamma) = \lim_{\Delta \rightarrow 0} V_1(\gamma)$. This coincides with the optimal deadline mechanism of Damiano, Li and Suen (2012).

5.2. General payoff structures

The committee model of Damiano, Li and Suen (2012) has the same information structure as the present model, but with a slightly different payoff structure that in some sense is more general. The following table illustrates a general payoff structure that incorporates both models as special cases.

	L	M	R
l	$(\bar{v} + \bar{\beta}, \bar{v})$	$(\underline{v} + \beta, \underline{v})$	$(\underline{v} + \underline{\beta}, \underline{v})$
r	$(\underline{v}, \underline{v} + \underline{\beta})$	$(\underline{v}, \underline{v} + \beta)$	$(\bar{v}, \bar{v} + \bar{\beta})$

In the above table, \bar{v} is the common value component of each player's payoff when the correct decision is made in a common interest state, i.e., l is chosen in state L or r is chosen in state R . Likewise, $\underline{v} < \bar{v}$ is the common value component when the wrong decision is

made in a common interest state or when either alternative is chosen in the conflict state M . The non-negative private value component in each player's payoff is only obtained by the player whose ex ante favorite alternative is chosen, and it depends on the state and on the alternative chosen: β is obtained in the conflict state M , and $\bar{\beta}$ and $\underline{\beta}$ in the common interest states with the former corresponding to the correct alternative and the latter the wrong one. In the present committee model we have $\bar{\beta} = \underline{\beta} = 0$ and $\beta = \bar{v} - \underline{v}$ (with $\bar{v} = 1$ and $\underline{v} = 1 - 2\lambda$), while in the model of Damiano, Li and Suen (2012) we have $\bar{\beta} = \underline{\beta} = \beta > 0$.

By allowing the private value components $\bar{\beta}$, $\underline{\beta}$ and β to take on different values, we can use the above general payoff structure to capture a variety of mixtures with common values and private values. Of course, for L and M to be common interest states, we need $\underline{\beta} < \bar{v} - \underline{v}$. This assumption also ensures that in delay mechanisms the informed player always persist in equilibrium.

A full characterization of optimal delay mechanisms under the general payoff structure is cumbersome because the Linkage Lemma no longer holds. Unlike in the present payoff structure, the informed player gets $\bar{v} + \bar{\beta}$ when the uninformed opponent concedes, which is different from the payoff of $\underline{v} + \beta$ that the uninformed player gets. Furthermore, the coin-flip payoff for the informed player, $\frac{1}{2}(\bar{v} + \underline{v} + \bar{\beta})$, is different from the coin-flip payoff of for the uninformed player, which is $\underline{v} + \frac{1}{2}\beta$ in the conflict state and $\frac{1}{2}(\bar{v} + \underline{v} + \underline{\beta})$ in the corresponding common interest state. As a result, maximizing the total delay while keeping the payoff to the uninformed player unchanged and maintaining his willingness to persist through the game is not the same as maximizing the payoff to the informed player.

Despite the failure of the Linkage Lemma, the qualitative features of optimal delay mechanisms established in the Main Result, finite delay, efficient deadline belief and stop-and-start, are all robust with respect to the payoff structure. The critical feature of the model turns out to be the information structure, not the payoff structure. The dichotomy between the equilibrium analysis of the uninformed player and the welfare analysis of the informed player, repeated exploited in our local variational approach, is possible because the informed player knows the state and in any equilibrium under a delay mechanism always persists with probability one. To illustrate this point, we briefly explain how to establish Lemma 3 under the general payoff structure, which is the key to the stop-and-start feature of optimal delay mechanisms.

To show that a mechanism with slack in two consecutive rounds $t(i)$ and $t(i + 1)$ is suboptimal, we use the same local variation exercise of changing $\sigma_{t(i-1)}$, $\sigma_{t(i)}$ and $\sigma_{t(i+1)}$

to keep $\gamma_{t(i)}$ and $\gamma_{t(i+2)}$ unchanged. The expressions for $\sigma_{t(i-1)}$, $\sigma_{t(i)}$ and $\sigma_{t(i+1)}$ involving the free variable $\gamma_{t(i+1)}$ are the same as the ones in the proof of Lemma 3 in the appendix, except that λ is replaced with $\frac{1}{2}\beta$. Since the sequence $\{\gamma_t, x_t, U_t\}$ is unaffected in the variation for $t \geq t(i+2)$, so is $V_{t(i+2)}$. We can then write $V_{t(i-1)}$ forward as a function of the single variable $\gamma_{t(i+1)}$. The only non-linear term that depends on $\gamma_{t(i+1)}$ involves $1/(1 - \gamma_{t(i+1)})$, and has a positive coefficient. Thus, the expected payoff of the informed player, $V_{t(i-1)}$, is a convex function of $\gamma_{t(i+1)}$. As in the current proof of Lemma 3, since there is slack in both round $t(i)$ and round $t(i+1)$, $\gamma_{t(i+1)}$ is in the interior of the feasible set, which implies that the mechanism cannot be optimal.¹²

The reason that the coefficient of the term involving $1/(1 - \gamma_{t(i+1)})$ is positive is that an increase in $\gamma_{t(i+1)}$ requires a decrease in $\sigma_{t(i-1)}$ and an increase in $\sigma_{t(i+1)}$ to keep respectively $\gamma_{t(i)}$ and $\gamma_{t(i+2)}$ constant. The former change has a greater impact on $V_{t(i-1)}$ because the impact of the latter is “discounted” by the concessions made by the uninformed player during the intervening period of $t(i)$.

5.3. General delay mechanisms

As a mechanism design exercise, the optimal delay problem analyzed in the present paper has the novel feature that the designer has limited commitment power. We have explained in section 2.2 that this feature is responsible in our model for making the design problem a dynamic one. It is well-known that with limited or no commitment, the standard Revelation Principle no longer applies in the sense that there is no canonical class of mechanisms that the designer could restrict to without sacrificing implementation capacities. In our problem, the only private information that the two players have is at the outset. Requiring them to report it to the designer once for all would fail to exploit the possibility of using dynamic delay mechanisms to mitigate the impact of limited commitment power. The approach we take in this paper is to restrict to a class of mechanisms or game forms that is natural in the setting of committee decision-making, and try to provide insights that are robust in problems of strategic information aggregation. The class of mechanisms we consider is dynamic and has the interpretation that in each round a direct revelation mechanism is played, in spite of the fact that the private information in our model is perfectly persistent.

One implicit restriction in the delay mechanisms introduced in section 2.2 is that the

¹²Lemma 4 also holds. We still have the payoff-equivalence result for the same reason: it remains true that $(1 - g(\gamma_{t(i)}))(1 - g^{-1}(\gamma_{t(i+2)})) = (1 - \gamma_{t(i)})(1 - \gamma_{t(i+2)})$ for properly redefined function g . This is due to the fact that the terms in $V_{t(i-1)}$ that are involved in the local variation argument are independent of the payoff structure for the informed.

direct revelation mechanism played in each round has no delay once the state is revealed on the equilibrium path. This involves two assumptions. First, if one player concedes while the other player persists, in equilibrium the state is revealed to be the common interest state corresponding to the latter's ex ante favorite alternative, which is immediately implemented. Second, if both players concede, in equilibrium the state is revealed to be the conflict state and the alternative is chosen without delay by a coin flip. A natural question is whether these two assumptions are without loss of generality in that they follow from the optimality of delay mechanisms. The answer turns out to be negative.

To illustrate, we reconsider one-round delay mechanisms, allowing for possible delay (and inefficient decisions) even after the state is revealed on the equilibrium path. As in section 2.2, we assume that the initial belief of the uninformed player γ and the upper-bound on delay Δ satisfy $\gamma > \gamma_*$, so that the first-best outcome cannot be implemented. The Revelation Principle applies in this setting. To simplify, we consider only symmetric mechanisms, given by $\{(\delta^A, q^A), \delta^R, \delta^D\}$, where, after an agreement in which one player concedes while the other persists, $\delta^A \leq \Delta$ is the imposed delay, and q^A is the probability of implementing the agreed alternative (and $1 - q^A$ is the probability of implementing the other alternative); after a reverse disagreement in which both players concede, $\delta^R \leq \Delta$ is the imposed delay before a coin flip; and after a regular disagreement in which both players persist, $\delta^D \leq \Delta$ is the imposed delay before a coin flip. The truth-telling incentive condition for the informed player is given by

$$V^A \equiv -\delta^A + q^A + (1 - q^A)(1 - 2\lambda) \geq -\delta^R + 1 - \lambda \equiv V^R.$$

The truth-telling condition for the uninformed player is

$$\gamma V^R + (1 - \gamma)V^A \geq \gamma V^A + (1 - \gamma)(-\delta^D + 1 - \lambda) \equiv \gamma V^A + (1 - \gamma)V^D.$$

To maximize the ex ante payoff given by

$$\frac{1}{2 - \gamma}(\gamma V^R + (1 - \gamma)V^A) + \frac{1 - \gamma}{2 - \gamma}V^A,$$

we need to choose the minimum value of V^D , which is $1 - \lambda - \Delta$, by setting $\delta^D = \Delta$. This is the same outcome of maximum delay after a regular disagreement that we have assumed in the one-round delay mechanisms introduced in section 2.2. Furthermore, either V^A or V^R , or both, must attain their respective maximum value, for otherwise increasing both of them by the same amount would raise the ex ante payoff without violating the

two incentive conditions. The maximum value of V^A is 1, by setting $\delta^A = 0$ and $q^A = 1$; and the maximum value of V^R is $1 - \lambda$, by setting $\delta^R = 0$. However, if $V^A = 1$ then from the two incentive conditions it is optimal to have $V^R = 1 - \lambda$, which violates the incentive condition for the uninformed because $\gamma > \gamma_*$. The optimal one-round delay mechanism instead has $V^R = 1 - \lambda$, and V^A binds the incentive condition of the uninformed player and is given by

$$V^A = 1 - \lambda + \frac{1 - \gamma}{2\gamma - 1}\Delta,$$

which is strictly less than 1 because $\gamma > \gamma_*$. Thus, the one-round delay mechanism with maximum delay considered in section 2.2 is not optimal, because we have assumed that there is no delay in implementing the agreed alternative after only one player concedes. Intuitively, an appropriate delay after an agreement reduces the incentive for the uninformed player to act as if informed and helps information aggregation.¹³

We can enrich the delay mechanisms considered in the present paper by allowing possible delay, subject to the same limited commitment bound, after an agreement and after a reverse disagreement in the direct revelation mechanism played in each round. Moreover, the three sequences of delays can be potentially history-dependent. A complete characterization of optimal delay mechanisms in such a general setting is beyond the scope of this paper. We hope to return it in future research.

¹³Indeed, the optimal one-round delay mechanism characterized above gives a higher expected payoff to the uninformed player as well as the informed than the one-round delay mechanism with maximum delay in section 2.2, and unlike the latter, it Pareto dominates flipping a coin for all $\gamma < 1$.

Appendix

Proof of Proposition 1. Suppose by way of contradiction that $x_T > 0$ in an optimal mechanism. This implies

$$\gamma_T \geq \frac{\lambda + \delta_T}{2\lambda + \delta_T} > \frac{1}{2}.$$

Let t' be the last round before T for which conceding is weakly preferred to persisting. If there is slack in round $t(r)$, then $t' = t(r)$; if there is no slack in round $t(r)$, then $t' = t(r) + 1$. Such t' must exist because the delay mechanism would be wasting time in rounds 1 through $T - 1$ otherwise.

We first obtain some bounds for the total delay during the period of “deadline play.”

Claim 1.

$$\begin{aligned} \sum_{t=t'+1}^{T-1} \delta_t &< U_T(\gamma_T) - (1 - 2\gamma_T\lambda); \\ \delta_{t'} + \sum_{t=t'+1}^{T-1} \delta_t &\geq U_T(\gamma_T) - (1 - 2\gamma_{t'}\lambda). \end{aligned}$$

By definition of t' , if $t' + 1 < T$ the uninformed strictly prefers persisting to conceding at round $t' + 1$. This condition can be written as:

$$\sum_{t=t'+1}^{T-1} \delta_t < U_T(\gamma_T) - (1 - 2\gamma_T\lambda). \quad (14)$$

At round t' , since persisting is not strictly preferred to conceding by the uninformed, we have the indifference condition

$$1 - \gamma_{t'}\lambda - \gamma_{t'}x_{t'}\lambda = (\gamma_{t'}x_{t'} + 1 - \gamma_{t'}) \left(-\delta_{t'} - \sum_{t=t'+1}^{T-1} \delta_t + U_T(\gamma_T) \right) + \gamma_{t'}(1 - x_{t'}),$$

which can be rewritten as

$$\begin{aligned} &\gamma_{t'}x_{t'} \left(1 - \lambda + \delta_{t'} + \sum_{t=t'+1}^{T-1} \delta_t - U_T(\gamma_T) \right) \\ &= (1 - \gamma_{t'}) \left(-\delta_{t'} - \sum_{t=t'+1}^{T-1} \delta_t + U_T(\gamma_T) \right) - (1 - \gamma_{t'}) + \gamma_{t'}\lambda. \end{aligned}$$

We note that $1 - \lambda + \delta_{t'} + \sum_{t=t'+1}^{T-1} \delta_t - U_T(\gamma_T) > 0$: otherwise, the right-hand-side of the above is strictly positive because $\gamma_{t'} > \gamma_T > \frac{1}{2}$, which is a contradiction. Given that the left-hand-side expression is positive, it must be increasing in $x_{t'}$. Evaluating the indifference condition at $x_{t'} \leq 1$ gives the second inequality of Claim 1. If $t' + 1 = T$, $\sum_{t=t'+1}^{T-1} \delta_t$ is defined to be 0. The first inequality is obtained from the condition that uninformed player strictly prefers persisting to conceding at round $t' + 1$ if $t' + 1 < T$. If $t' + 1 = T$, the first inequality is not required to hold.

We establish that when $x_T > 0$, marginally reducing the equilibrium deadline belief would improve the ex ante payoff. This can be achieved by the following modification to the original mechanism:

1. Introduce two “extra rounds” s' and s such that $t' < s' < s < t' + 1$. Set $\delta_{s'} = 0$ and set $\delta_s = U_T(\gamma_T) - (1 - 2\gamma_T\lambda) - \sum_{t=t'+1}^{T-1} \delta_t \equiv \zeta$. Subtract the amount ζ from $\delta_{t'}$ in round t' .
2. Marginally raise δ_s from its original value of ζ in step 1, while at the same time raising $\delta_{s'}$ from its original value of 0 in such a way to keep $-\delta_{s'} + U_s(\gamma_s)$ fixed.

Step 1 of this construction is feasible because Claim 1 and the fact that $\gamma_{t'} \geq \gamma_T$ imply $\zeta \in [0, \Delta)$ and $\delta_{t'} - \zeta > 0$. Given this construction, the effective delay $\sigma_{t'}$ remains unchanged at round t' . Moreover, the uninformed is just indifferent between persisting and conceding with $x_s = 1$ at round s . At round s' , since $\delta_{s'} = 0$, $x_{s'} = 1$ is consistent with equilibrium. This step has no effect on the equilibrium play in all other rounds and it has no effect on the ex ante payoff.

Step 2 of the construction attempts to marginally lower γ_T by reducing x_s from its original value of 1 to a value below 1. Since $\delta_s = \zeta < \Delta$, there is slack at round s , so the modification is feasible. The effect on the ex ante payoff depends on whether the value of $U_T(\gamma_T)$ is changed by this step.

There are two cases to consider. In case (i), the uninformed player strictly prefers persisting to conceding at round T under the original mechanism. In this case, marginally changing δ_s would lower γ_T marginally but would leave $x_T = 1$ under the modified mechanism. The value of $U_T(\gamma_T)$ would be fixed at $1 - \lambda - \delta_T$. Since $x_{t'+1} = \dots = x_T = 1$, the uninformed is indifferent between persisting and conceding at round s if

$$1 - \gamma_s\lambda - \gamma_s x_s \lambda = (1 - \gamma_s + \gamma_s x_s) \left(1 - \lambda - \delta_s - \sum_{t=t'+1}^T \delta_t \right) + \gamma_s (1 - x_s).$$

For fixed γ_s , we can differentiate with respect to x_s to obtain:

$$(1 - \gamma_s + \gamma_s x_s) \frac{d\delta_s}{dx_s} = -\gamma_s \left(\delta_s + \sum_{t=t'+1}^T \delta_t \right).$$

Therefore $d\delta_s/dx_s < 0$. In other words, step 2 would require us to raise δ_s .

In case (ii), the uninformed player is indifferent between persisting and conceding at round T under the original mechanism. In this case, lowering γ_T would change the value of $U_T(\gamma_T)$. The indifference condition at round s is:

$$1 - \gamma_s \lambda - \gamma_s x_s \lambda = (1 - \gamma_s + \gamma_s x_s) \left(U_T(\gamma_T) - \delta_s - \sum_{t=t'+1}^{T-1} \delta_t \right) + \gamma_s (1 - x_s),$$

where $U_T(\gamma_T) = 1 - \lambda \gamma_T - \lambda \gamma_T x_T$, and γ_T depends on x_s through Bayes' rule. Differentiate with respect to x_s to obtain:

$$(1 - \gamma_s + \gamma_s x_s) \frac{d\delta_s}{dx_s} = \gamma_s \left(-1 + \lambda - \delta_s - \sum_{t=t'+1}^{T-1} \delta_t + U_T(\gamma_T) \right) + (1 - \gamma_s + \gamma_s x_s) \frac{dU_T(\gamma_T)}{d\gamma_T} \frac{d\gamma_T}{dx_s}.$$

We have already established the claim that the first term on the right-hand-side is negative. Furthermore, $d\gamma_T/dx_s > 0$ by Bayes' rule, and

$$\frac{dU_T(\gamma_T)}{d\gamma_T} = -(1 + x_T)\lambda - \gamma_T \frac{dx_T}{d\gamma_T} \lambda < 0,$$

because the indifference condition at round T ensures that $dx_T/d\gamma_T > 0$. Thus we again have $d\delta_s/dx_s < 0$ in this case.

For both cases (i) and (ii), when x_s falls, other things being equal, this would increase $U_s(\gamma_s)$. Our construction in step 2 raises $\delta_{s'}$ in such a way to keep the continuation value $-\delta_{s'} + U_s(\gamma_s)$ constant. Since $U_s(\gamma_s) = 1 - \gamma_s \lambda - \gamma_s x_s \lambda$, the change in $\delta_{s'}$ needed to keep the continuation value constant is $d\delta_{s'}/dx_s = -\gamma_s \lambda < 0$. In other words, step 2 would require us to raise $\delta_{s'}$.

In this construction, when we change $x_s = 1$ in step 1 to $\tilde{x}_s = 1 - \epsilon$ in step 2, the total delay changes by

$$-(d\delta_{s'}/dx_s + d\delta_s/dx_s)\epsilon > 0.$$

Moreover, since $x_T > 0$ in step 1, we have $\tilde{x}_T \geq 0$ in step 2 by choosing ϵ to be small. Therefore, it is a weak best response for the uninformed player to persist throughout in the modified mechanism. Finally, by construction, $U_1(\gamma_1)$ remains unchanged by our local variation method. Lemma 1 then implies that $V_1(\gamma_1)$ is increased. The original mechanism cannot be optimal. ■

Proof of Proposition 2. Suppose to the contrary that $\gamma_T < \gamma_*$. There are two cases: (i) $\gamma_{t(r)} \leq \gamma_*$; or (ii) $\gamma_{t(r)} > \gamma_*$. Since $r_*(\gamma_1) \geq 2$, there exist another updating round $t(r-1)$ before round $t(r)$.

Take case (i) first. We consider the following modification to the original mechanism:

1. Change $\sigma_{t(r)}$ in such a way to make the uninformed just indifferent between persisting and conceding at $\tilde{x}_{t(r)} = 1$. This is achieved by setting $\tilde{\sigma}_{t(r)} = \gamma_{t(r)}\lambda$.
2. Change $\sigma_{t(r-1)}$ in such a way to keep the continuation payoff for the uninformed at round $t(r-1)$ fixed. This is achieved by setting $\tilde{\sigma}_{t(r-1)} = \sigma_{t(r-1)} - \gamma_{t(r)}(1 - x_{t(r)})\lambda$.

With step 1, the new equilibrium belief after round $t(r)$ is $\tilde{\gamma}_{t(r)+1} = \dots = \tilde{\gamma}_T = \gamma_{t(r)}$. Note that the requisite delay $\tilde{\sigma}_{t(r)}$ can always be obtained by adding “extra rounds” between $t(r)$ and $t(r) + 1$ if necessary. Thus step 1 of the construction is feasible.

From the indifference condition at round $t(r-1)$ under the original mechanism,

$$\sigma_{t(r-1)} = \frac{\gamma_{t(r-1)}\lambda}{\gamma_{t(r-1)}x_{t(r-1)} + 1 - \gamma_{t(r-1)}} - \gamma_{t(r)}x_{t(r)}\lambda > \gamma_{t(r)}(1 - x_{t(r)})\lambda.$$

Thus $\tilde{\sigma}_{t(r-1)} > 0$, which means that step 2 is feasible.

Since the uninformed is persisting through rounds $t(r-1), \dots, T-1$ and concedes at round T (because $\tilde{\gamma}_T = \gamma_{t(r)} < \gamma_*$) in the equilibrium of the modified mechanism, the payoff to the informed at round $t(r-1)$ is

$$\tilde{V}_{t(r-1)}(\gamma_{t(r-1)}) = 1 - x_{t(r-1)} \left(\tilde{\sigma}_{t(r-1)} + \tilde{\sigma}_{t(r)} \right).$$

His payoff under the original mechanism is

$$V_{t(r-1)}(\gamma_{t(r-1)}) = 1 - x_{t(r-1)} \left(\sigma_{t(r-1)} + x_{t(r)}\sigma_{t(r)} \right).$$

Thus,

$$\tilde{V}_{t(r-1)}(\gamma_{t(r-1)}) - V_{t(r-1)}(\gamma_{t(r-1)}) = x_{t(r-1)}x_{t(r)} \left(\sigma_{t(r)} - \gamma_{t(r)}\lambda \right).$$

From the indifference condition for the uninformed player at round $t(r)$ under the original mechanism,

$$x_{t(r)} = \frac{\gamma_{t(r)}\lambda - (1 - \gamma_{t(r)})\sigma_{t(r)}}{\gamma_{t(r)}\sigma_{t(r)}}.$$

Thus, $x_{t(r)} < 1$ implies $\sigma_{t(r)} > \gamma_{t(r)}\lambda$. We conclude that $\tilde{V}_{t(r-1)}(\gamma_{t(r-1)}) > V_{t(r-1)}(\gamma_{t(r-1)})$.

Next, consider case (ii). Suppose we modify the mechanism so that the game ends with deadline belief $\tilde{\gamma}_T$ instead of γ_T , where $\tilde{\gamma}_T \in [g(\gamma_{t(r)}), \gamma_*]$. This is achieved by:

1. Change $\sigma_{t(r)}$ to induce equilibrium $\tilde{x}_{t(r)}$ that satisfies

$$\frac{\tilde{\gamma}_T}{1 - \tilde{\gamma}_T} = \frac{\gamma_{t(r)}}{1 - \tilde{\gamma}_{t(r)}} \tilde{x}_{t(r)}.$$

2. Change $\sigma_{t(r-1)}$ in such a way to keep the continuation payoff for the uninformed at round $t(r-1)$ fixed. This is achieved by setting $\tilde{\sigma}_{t(r-1)} = \sigma_{t(r-1)} - \gamma_{t(r)}(\tilde{x}_{t(r)} - x_{t(r)})\lambda$.

Using the fact that $U_T(\tilde{\gamma}_T) = 1 - \tilde{\gamma}_T\lambda$, the indifference condition at round $t(r)$ gives the requisite amount of effective delay in step 1 of the construction:

$$\tilde{\sigma}_{t(r)} = \frac{\gamma_{t(r)}\lambda}{\gamma_{t(r)}\tilde{x}_{t(r)} + 1 - \gamma_{t(r)}}.$$

Note also that this equation implies $\tilde{x}_{t(r)}\tilde{\sigma}_{t(r)} = \tilde{\gamma}_T\lambda$.

We have already shown in the analysis of case (i) that $\sigma_{t(r-1)} > \gamma_{t(r)}(1 - x_{t(r)})\lambda$. This implies that $\tilde{\sigma}_{t(r-1)} > 0$ for case (ii) as well. Thus step 2 of our modification is feasible.

The change in payoff to the informed, $\tilde{V}_{t(r-1)}(\gamma_{t(r-1)}) - V_{t(r-1)}(\gamma_{t(r-1)})$ is:

$$\begin{aligned} & x_{t(r-1)} \left(\sigma_{t(r-1)} + x_{t(r)}\sigma_{t(r)} \right) - x_{t(r-1)} \left(\tilde{\sigma}_{t(r-1)} + \tilde{x}_{t(r)}\tilde{\sigma}_{t(r)} \right) \\ &= x_{t(r-1)} \left(\gamma_{t(r)}(\tilde{x}_{t(r)} - x_{t(r)})\lambda + \gamma_T\lambda - \tilde{\gamma}_T\lambda \right) \\ &= x_{t(r-1)} (h(\tilde{\gamma}_T) - h(\gamma_T)) \lambda, \end{aligned}$$

where

$$h(\gamma) \equiv \frac{\gamma(\gamma - \gamma_{t(r)})}{1 - \gamma}.$$

Since $h(\tilde{\gamma}_T)$ is convex, it reaches a maximum when $\tilde{\gamma}_T$ is either γ_* or $g(\gamma_{t(r)})$. Evaluating

the value of the function at these two points, we obtain:

$$h(\gamma_*) - h(g(\gamma_{t(r)})) = (\gamma_* - g(\gamma_{t(r)})) \frac{\lambda}{\lambda + \Delta} > 0.$$

Since $h(\tilde{\gamma}_T)$ reaches a maximum at γ_* for any $\tilde{\gamma}_T$ in the interval $[g(\gamma_{t(r)}), \gamma_*]$, and since γ_T also belongs to that interval, we have $h(\gamma_*) > h(\gamma_T)$. We conclude that $\tilde{V}_{t(r-1)}(\gamma_{t(r-1)}) > V_{t(r-1)}(\gamma_{t(r-1)})$ when $\tilde{\gamma}_T$ is chosen to be equal to γ_* .

In both cases (i) and (ii), the modified mechanism does not change the uninformed player's strategy in rounds 1 through $t(r-1)$. Thus $U_1(\gamma_1)$ is not affected by the modification. But since $\tilde{V}_{t(r-1)}(\gamma_{t(r-1)})$ is increased, an induction argument back to round 1 shows that $\tilde{V}_1(\gamma_1)$ is increased by the modification. The original mechanism cannot be optimal. \blacksquare

Proof of Lemma 3. From the indifference condition at round $t(i)$, we have

$$\sigma_{t(i)} = \frac{\gamma_{t(i)}(1 - x_{t(i)}x_{t(i+1)})\lambda}{1 - \gamma_{t(i)} + \gamma_{t(i)}x_{t(i)}}.$$

We can use Bayes' rule to express this in terms of the belief $\gamma_{t(i+1)}$:

$$\sigma_{t(i)} = \frac{(\gamma_{t(i)} - \gamma_{t(i+2)})(1 - \gamma_{t(i+1)})\lambda}{(1 - \gamma_{t(i)})(1 - \gamma_{t(i+2)})}.$$

Likewise the indifference condition at round $t(i+1)$ is

$$\sigma_{t(i+1)} = \frac{\gamma_{t(i)}x_{t(i)}\lambda}{1 - \gamma_{t(i)} + \gamma_{t(i)}x_{t(i)}x_{t(i+1)}} + \frac{\gamma_{t(i)}x_{t(i)}x_{t(i+1)}\lambda}{1 - \gamma_{t(i)} + \gamma_{t(i)}x_{t(i)}x_{t(i+1)}} + U_{t(i+2)}(\gamma_{t(i+2)}) - 1.$$

All the terms except the first one are constant for fixed $\gamma_{t(i)}$ and $\gamma_{t(i+1)}$. We can use Bayes' rule again to transform the first term and obtain:

$$\sigma_{t(i+1)} = \frac{\gamma_{t(i+1)}(1 - \gamma_{t(i+2)})\lambda}{1 - \gamma_{t(i+1)}} + \text{constant}.$$

Finally to maintain $\gamma_{t(i)}$ fixed, we need to adjust $\sigma_{t(i-1)}$ to make sure that the continuation value $U_{t(i)}(\gamma_{t(i)}) - \sigma_{t(i-1)}$ is held constant. This implies that

$$\sigma_{t(i-1)} = -\gamma_{t(i)}x_{t(i)}\lambda + \text{constant} = -\frac{\gamma_{t(i+1)}(1 - \gamma_{t(i)})\lambda}{1 - \gamma_{t(i+1)}} + \text{constant}.$$

Summing up the three terms, we obtain

$$\sum_{t=t(i-1)}^{t(i+1)} \sigma_t = \frac{\lambda(\gamma_{t(i)} - \gamma_{t(i+2)})}{1 - \gamma_{t(i+2)}} \left(\frac{1 - \gamma_{t(i+1)}}{1 - \gamma_{t(i)}} + \frac{1 - \gamma_{t(i+2)}}{1 - \gamma_{t(i+1)}} \right) + \text{constant}.$$

This is a convex function of $\gamma_{t(i+1)}$, so it attains a maximum at the boundaries of the feasible set. If there is no slack in both round $t(i)$ and round $t(i+1)$, then $\gamma_{t(i+1)}$ is in the interior of the feasible set. This mechanism cannot be optimal. ■

Proof of Lemma 4. From the proof of Lemma 3, we can evaluate the total delay $\sum_{t=t(i-1)}^{t(i+1)} \sigma_t$ at the point $\gamma_{t(i+1)} = g^{-1}(\gamma_{t(i+2)})$ (i.e., no slack at round $t(i+1)$) and the point $\gamma_{t(i+1)} = g(\gamma_{t(i)})$ (i.e., no slack at round $t(i)$). They give exactly the same value of

$$\sum_{t=t(i-1)}^{t(i+1)} \sigma_t = \frac{\lambda + \Delta}{\lambda} + \frac{1 - \gamma_{t(i+2)}}{1 - \gamma_{t(i)}} \frac{\lambda}{\lambda + \Delta} + \text{constant}.$$

The two mechanisms are payoff-equivalent. ■

Proof of Lemma 5. Using the payoff from concession at round $t(2)$ to write the payoff to the uninformed player, the indifference condition at round 1 can be written as

$$\gamma_1(1 - x_1 x_{t(2)})\lambda = (1 - \gamma_1 + \gamma_1 x_1)\sigma_1.$$

Given our construction, γ_1 and $x_1 x_{t(2)}$ are both fixed. Therefore,

$$\frac{d\sigma_1}{dx_1} = -\frac{\gamma_1^2(1 - x_1 x_{t(2)})\lambda}{(1 - \gamma_1 + \gamma_1 x_1)^2} < 0.$$

This means that to lower x_1 requires raising σ_1 . Since there is slack in round 1, raising σ_1 is feasible.

Similarly, the indifference condition at round $t(2)$ can be written as

$$\gamma_1 x_1 \lambda + \gamma_1 x_1 x_{t(2)} \lambda = (1 - \gamma_1 + \gamma_1 x_1 x_{t(2)})(\sigma_{t(2)} + 1 - U_{t(2)}(\gamma_{t(2)})).$$

In the above equation, the values of γ_1 , $x_1 x_{t(2)}$, and $U_{t(2)}(\gamma_{t(2)})$ are held constant. Therefore,

$$\frac{d\sigma_{t(2)}}{dx_1} = \frac{\gamma_1 \lambda}{1 - \gamma_1 + \gamma_1 x_1 x_{t(2)}} > 0.$$

This means that to lower x_1 requires lowering $\sigma_{t(2)}$. Since $\sigma_{t(2)} > 0$ in the original mechanism, this step is also feasible.

The effect of this modification on the payoff to the uninformed is

$$\frac{dU_1(\gamma_1)}{dx_1} = -\gamma_1\lambda < 0.$$

Therefore, lowering x_1 (while holding $x_1x_{t(2)}$ constant) increases the payoff to the uninformed.

Since $\gamma_T = \gamma_*$ in both the original and the modified mechanism, persisting throughout the game is a weak best response to the equilibrium strategy of the uninformed. Therefore equation (8) holds. The effect of a change in x_1 (while holding $x_1x_{t(2)}$ constant) on the payoff of the informed can be calculated by:

$$\begin{aligned} \frac{dV_1(\gamma_1)}{dx_1} &= \frac{1}{\gamma_1} \frac{dU_1(\gamma_1)}{dx_1} + \frac{1-\gamma_1}{\gamma_1} \left(\frac{d\sigma_1}{dx_1} + \frac{d\sigma_{t(2)}}{dx_1} \right) \\ &= \frac{1-\gamma_1}{\gamma_1} \frac{d\sigma_1}{dx_1} + \lambda \left(-1 + \frac{1-\gamma_1}{1-\gamma_1 + \gamma_1x_1x_{t(2)}} \right) < 0. \end{aligned}$$

Hence, both the informed and the uninformed can be made better off by the modified mechanism if there is slack at round 1. ■

Proof of Proposition 4. Fix any $\gamma_1 \in (\gamma_*, (\lambda + \Delta)/(2\lambda))$. In the one-round delay mechanism with maximum delay, the probability that the uninformed player persists is given by

$$x_1 = \frac{\gamma_1\lambda - (1-\gamma_1)(\lambda + \Delta)}{\gamma_1\Delta},$$

which is strictly less than 1 because $\gamma_1 < (\lambda + \Delta)/(2\lambda)$. The expected payoff is

$$U_1 = 1 - \gamma_1\lambda - \gamma_1x_1\lambda$$

for the uninformed, and

$$V_1 = 1 - x_1(\lambda + \Delta)$$

for the informed. By (7), the difference between the ex ante payoff W_1 and the coin flip

payoff $1 - \lambda$ can be shown to have the same sign as

$$\frac{2\gamma_1\lambda\Delta}{1-\gamma_1} - \left(\frac{\gamma_1\lambda}{1-\gamma_1} + (\lambda + \Delta) \right) \left(\frac{\gamma_1\lambda}{1-\gamma_1} - (\lambda + \Delta) \right).$$

It is straightforward to verify that the above expression is positive at

$$\gamma_1 = g^{-1}(1/2) = \frac{\lambda + 2\Delta}{2(\lambda + \Delta)}.$$

Further we can show that the above expression is decreasing in γ_1 for all $\gamma_1 > \gamma_*$. Thus, the ex ante payoff W_1 under the one-round delay mechanism with maximum delay is strictly greater than the coin-flip payoff of $1 - \lambda$ for all $\gamma_1 \in (\gamma_*, g^{-1}(1/2))$.

Next, fix any $\gamma_1 \in (\gamma_*, g^{-1}(\gamma_*))$. Note that $g^{-1}(\gamma_*) \in (g^{-1}(1/2), (\lambda + \Delta)/(2\lambda))$. First, we compare the ex ante payoff $W_1(\gamma_1)$ under the mechanism derived in Lemma 6 with W_1 under the one-round delay mechanism with maximum delay derived above. By Lemma 6, the expected payoff is

$$U_1(\gamma_1) = 1 - \gamma_1\lambda - \gamma_1\chi(\gamma_1)\lambda$$

for the uninformed, and

$$V_1(\gamma_1) = 1 - g(\gamma_1)\lambda$$

for the informed. By (7), the difference between $W_1(\gamma_1)$ and W_1 under the one-round mechanism with maximum delay can be shown to have the same sign as

$$(\lambda + \Delta)^2 \left(\frac{\gamma_1\lambda}{1-\gamma_1} - (\lambda + \Delta) \right) - \frac{\gamma_1\lambda}{1-\gamma_1} \left((2\lambda + \Delta) + \frac{\gamma_1\lambda}{1-\gamma_1} \right) \left((\lambda + 2\Delta) - \frac{\gamma_1\lambda}{1-\gamma_1} \right).$$

The above expression is negative at $\gamma_1 = \gamma_*$, positive at $\gamma_1 = g^{-1}(1/2)$, and increasing in γ_1 for all $\gamma_1 > \gamma_*$. It follows that there exists a unique $\underline{\gamma} \in (\gamma_*, g^{-1}(1/2))$ such that the one-round delay mechanism with maximum delay dominates the mechanism derived in Lemma 6 if and only if $\gamma_1 < \underline{\gamma}$. Second, we compare the ex ante payoff $W_1(\gamma_1)$ with the coin-flip payoff of $1 - \lambda$. The difference between $W_1(\gamma_1)$ and $1 - \lambda$ can be shown to have the same sign as

$$2\lambda(\lambda + \Delta)(1 - \gamma_1) - (\lambda + (1 - \gamma_1)(\lambda + \Delta))(\gamma_1\lambda - (1 - \gamma_1)\Delta).$$

The above expression is strictly increasing in γ_1 . Further, at $\gamma_1 = g^{-1}(\gamma_*)$, we show below that the difference $W_1(\gamma_1) - (1 - \lambda)$ is positive if and only if $\Delta < (\sqrt{2} - 1)\lambda$.

Finally, fix any $\gamma_1 \in (g^{-1}(\gamma_*), 1)$. Using the expressions (13) for $U_1(\gamma_1)$ and $V_1(\gamma_1)$ from (8), we obtain from (7) that

$$W_1(\gamma_1) - (1 - \lambda) = \frac{1}{\gamma_1(2 - \gamma_1)} \left((1 - \gamma_1)^2 \sum_{t=1}^T \delta_t - (2\gamma_1 - 1)\lambda + \frac{\lambda\Delta}{\lambda + \Delta} \right).$$

The derivative with respect to γ_1 of the terms in the bracket on the right-hand-side of the above equation is given by

$$(1 - \gamma_1)^2 \frac{d}{d\gamma_1} \left(\sum_{t=1}^T \delta_t \right) - 2(1 - \gamma_1) \sum_{t=1}^T \delta_t - 2\lambda.$$

Using (12) and the definition of r_* from (3), we have

$$\frac{d}{d\gamma_1} \left(\sum_{t=1}^T \delta_t \right) = \frac{\lambda}{\eta^2} \frac{d\eta}{d\gamma_1} + \lambda < \frac{2 - \gamma_1}{1 - \gamma_1} \lambda,$$

as $\eta > 1$. Thus, the difference $W_1(\gamma_1) - (1 - \lambda)$ can cross zero only once and from above. As γ_1 approaches 1, we have $W_1(\gamma_1) - (1 - \lambda) < 0$. At $\gamma_1 = g^{-1}(\gamma_*)$, using (12) we have that

$$W_1(g^{-1}(\gamma_*)) - (1 - \lambda) = \frac{\lambda^2(\lambda^2 - (2\lambda + \Delta)\Delta)}{(2\lambda + \Delta)(3\lambda^2 + 3\lambda\Delta + \Delta^2)},$$

which is positive if and only if $\Delta < (\sqrt{2} - 1)\lambda$. The proposition follows immediately. ■

Proof of Lemma 7. Fix an infinite mechanism $(\delta_t)_{t=1}^{\infty}$ and suppose that it is not effectively finite. Due to the convention that the payoff from never implementing an alternative is zero and the normalization of $\lambda < \frac{1}{2}$, we can assume that $\sum_{t \geq \tau} \delta_t > 0$ for all τ .

First, we claim that there exists a positive number M such that $\sum_{t=1}^{\infty} \delta_t < M$. Suppose not. Since the mechanism is not effectively finite, it is a weak best response for the uninformed player to always persist. But the assumption of an unbounded total delay implies that the ex ante payoff to each player would be unbounded from below, contradicting equilibrium conditions.

Second, we claim that in any equilibrium $y_t < 1$ finitely often, where y_t is the probability that the informed player persists in round t . Suppose not. Since the mechanism is not effectively finite, there is no equilibrium in which $y_t = 0$ for some t . Then, for any τ such that $y_\tau < 1$, we have the indifference condition of the informed player

$$V_\tau = (1 - x_\tau) + x_\tau(-\delta_\tau + V_{\tau+1}) = (1 - x_\tau)(1 - \lambda) + x_\tau(1 - 2\lambda),$$

where $V_{\tau+1}$ is his continuation payoff. Since the informed player knows he is playing against the uninformed player and the latter can always concede to obtain a payoff of at least $1 - 2\lambda$, we have $V_{\tau+1} \geq 1 - 2\lambda$. It then follows from the above indifference condition that $x_\tau \geq \lambda/(\lambda + \delta_\tau)$. Since $\sum_{t=1}^{\infty} \delta_t < M$ by the first claim above, we δ_τ is arbitrarily close to 0 for sufficiently large τ , implying that x_τ is arbitrarily close to 1 and V_τ is arbitrarily close to $1 - 2\lambda$. Now, for any $\tau' > \tau$ the strategy of persisting from round τ through round $\tau' - 1$ and then conceding in round τ' yields an expected payoff at least as great as

$$1 - \prod_{t=\tau}^{\tau'-1} x_t + \prod_{t=\tau}^{\tau'-1} x_t (1 - 2\lambda) - \sum_{t=\tau}^{\tau'-1} \delta_t.$$

Since the last term in the above expression is arbitrarily close to 0 for sufficiently large τ and any $\tau' > \tau$ by the first claim above, for the strategy not to yield an expected payoff greater than $1 - 2\lambda$, we must have that $\prod_{t=\tau}^{\tau'-1} x_t$ is arbitrarily close to 1. As we take τ' to be arbitrarily large, the product of the equilibrium probabilities that the uninformed player concedes is arbitrarily close to 0. This is a contradiction, because by assumption that the payoff from never implementing an alternative is 0 to the informed player, and $y_t < 1$ infinitely often so always persisting is an equilibrium strategy for the informed player.

Third, we claim that $x_t = 0$ for some t , which establishes the lemma. By the second claim above, there is τ such that $y_t = 1$ for all $t \geq \tau$. The equilibrium belief γ_t of the uninformed player decreases in t for $t \geq \tau$ and is bounded from below by 0. Since the delay mechanism is not effectively finite, γ_t converges. It then follows from repeated applications of Bayes' rule that $\prod_{t=\tau}^{\tau'} x_t$ is arbitrarily close to 1 for all $\tau' > \tau$. This contradicts the equilibrium condition, because the payoff from never implementing an alternative is 0 to the uninformed player, and the mechanism is not effectively finite so always persisting is an equilibrium strategy for the uninformed player. ■

Proof of Proposition 5. Consider first a finite delay mechanism and an equilibrium that ends with probability 1 in the deadline round T . In this case, we write the continuation payoffs for the uninformed and the informed player after the last delay δ_T and before a coin-flip as $U_{T+1} = V_{T+1} = 1 - \lambda$.

Next, consider any finite delay mechanism that has an equilibrium ending with probability 1 in some round N before the deadline round T , or any infinite mechanism, which by Lemma 7 can only have an equilibrium ending in some round N . Since the the game ends with probability 1 in round N , in any such equilibrium the uninformed player concedes, with $x_N = 0$. Given this, in such equilibrium the informed player persists, with

$y_N = 1$. It follows that in this case the payoffs to the uninformed and the informed player are

$$U_N = 1 - \gamma_N \lambda \leq V_N = 1$$

regardless of the belief γ_N of the uninformed player.

Now, suppose that $U_{t+1} \leq V_{t+1}$ for some $t \leq \min\{T, N\}$. We first show that if the uninformed player weakly prefers persisting to conceding in round t then the informed player strictly prefers persisting to conceding. Let γ_t be the belief of the uninformed player. Then, for the uninformed player to weakly prefer persisting to conceding, we must have

$$\begin{aligned} & \gamma_t(1 - x_t + x_t(-\delta_t + U_{t+1})) + (1 - \gamma_t)((1 - y_t)(1 - 2\lambda) + y_t(-\delta_t + U_{t+1})) \\ & \geq \gamma_t((1 - x_t)(1 - \lambda) + x_t(1 - 2\lambda)) + (1 - \gamma_t)((1 - y_t)(1 - \lambda) + y_t), \end{aligned}$$

where x_t and y_t are the probability that the uninformed player and the informed player persist respectively. Since

$$(1 - y_t)(1 - 2\lambda) + y_t(-\delta_t + U_{t+1}) < (1 - y_t)(1 - \lambda) + y_t$$

for all y_t , we have

$$1 - x_t + x_t(-\delta_t + U_{t+1}) > (1 - x_t)(1 - \lambda) + x_t(1 - 2\lambda).$$

By assumption $U_{t+1} \leq V_{t+1}$, we have

$$1 - x_t + x_t(-\delta_t + V_{t+1}) > (1 - x_t)(1 - \lambda) + x_t(1 - 2\lambda),$$

implying that the informed player strictly prefers persisting to conceding in round t regardless of x_t . Thus, in any equilibrium we have $y_t = 1$. As a result, the expected payoffs for the uninformed and informed player in round t are

$$\begin{aligned} U_t &= \gamma_t(1 - x_t + x_t(-\delta_t + U_{t+1})) + (1 - \gamma_t)(-\delta_t + U_{t+1}); \\ V_t &= 1 - x_t + x_t(-\delta_t + V_{t+1}). \end{aligned}$$

It is straightforward to use the above expressions, and the assumption that $U_{t+1} \leq V_{t+1}$ to verify that $U_t \leq V_t$. The proposition follows immediately from induction. ■

References

- Abreu, D., and Gul, Faruk (2000). "Bargaining and Reputation." *Econometrica* 68, 85–117.
- Bester, Helmut, and Strausz, Roland (2001). "Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case." *Econometrica* 69(4), 1077–1098.
- Cramton, Peter (1992). "Strategic Delay in Bargaining with Two-Sided Uncertainty." *Review of Economic Studies* 59(1), 205–225.
- Damiano, Ettore, Li, Hao, and Suen, Wing (2012). "Optimal Deadline for Agreements." *Theoretical Economics*, forthcoming.
- Deneckere, Raymond, and Liang, Meng-Yu (2006). "Bargaining with Interdependent Values." *Econometrica* 74(5), 1309–1364.
- Dutta, B., and Sen, A. "A Necessary and Sufficient Condition for Two-Person Nash Implementation." *Review of Economic Studies* 58, 121–128.
- Feddersen, Timothy, and Pesendorfer, Wolfgang (1997). "Voting Behavior and Information Aggregation in Elections with Private Information." *Econometrica*, 65, 1029–1058.
- Hendricks, K., Weiss, A., and Wilson, C. (1988). "The War of Attrition in Continuous Time and Complete Information." *International Economic Review* 29(4), 663–680.
- Holmstrom, Bengt (1984). "On the Theory of Delegation," in *Bayesian Models of Economic Theory*, edited by Boyer and Kihlstrom. North-Holland, New York.
- Kolotilin, Anton, Li, Hao, and Li, Wei (2011). "Optimal Limited Authority for Principal." University of British Columbia working paper.
- Li, Hao, and Suen, Wing (2009). "Decision-making in Committees." *Canadian Journal of Economics* 42(2), 359–392.
- Li, Hao, Rosen, Sherwin, and Suen, Wing (2001). "Conflicts and Common Interests in Committees." *American Economic Review* 91(5), 1478–1497.
- Moore, John, and Repullo, R. (1990). "Nash Implementation: A Full Characterization." *Econometrica* 58, 1083–1099.

Myerson, Roger, and Satterthwaite, Mark (1983). "Efficient Mechanisms for Bilateral Trading." *Journal of Economic Theory* 29(2), 265–281.

Skreta, Vasiliki (2006). "Sequentially Optimal Mechanisms." *Review of Economic Studies* 73(4), 1085–1111.