Optimal Delay in Committees

ETTORE DAMIANO
University of Toronto

LI, HAO
University of British Columbia

WING SUEN
University of Hong Kong

November 2, 2013

Abstract. We consider a committee problem in which efficient information aggregation is hindered by differences in preferences. Sufficiently large delays could foster information aggregation but would require commitment. In a dynamic delay mechanism with limited commitment, successive rounds of decision-making are punctuated by delays that are uniformly bounded from above. Any optimal sequence of delays is finite, inducing in equilibrium both a "deadline play," in which a period of no activity before the deadline is followed by full concession at the end to reach the efficient decision, and "stop-and-start" in the beginning, in which the maximum concession feasible alternates with no concession. Stop-and-start results from simultaneously maximizing both the "static" incentives for truth-telling – by maximizing the immediate delay penalty – and the "dynamic" incentives – by minimizing continuation payoffs.

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. 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.[1] 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 an earlier 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, even the first best may be achievable: 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 can commit to

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

[2]See Bester and Strausz (2001), Skreta (2006), and Kolotilin, Li and Li (2013) for other models of limited commitment.

“wasting” a small amount of value (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 dynamic trade-off between imposing a greater collective punishment through delay and raising the probability of making the collectively desirable decision. This framework allows us to ask questions that cannot be addressed in our previous paper: Does punishment (delay) work better if it is front-loaded or back-loaded? Is it optimal to maintain a constant sequence of 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 preferences. 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 potentially variable delays between successive rounds.

We have deliberately chosen a simple model to capture the feature that, 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. Introducing a collective punishment in the form of delay if the members disagree may improve decision-making, and 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 to 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

[3] See also Hendricks, Weiss and Wilson (1988), Cramton (1992), Abreu and Gul (2000), and Deneckere and Liang (2006).

round t. We study its effect by simultaneously adjusting the delay in round $t-1$ (through the introduction an extra round if necessary) in such a way that keeps the continuation payoff for round $t-2$ fixed, and adjusting the delay in round $t+1$ (also through the introduction of an extra round if necessary) in such a way that keeps the equilibrium belief in 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 localized variation method, we can arrive at an essentially complete characterization of optimal delay mechanisms.

The main result of this paper is a characterization of all delay mechanisms that have a symmetric 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 1 in the final round if it is reached. The decision is thus always Pareto efficient in equilibrium. Second, we show that in 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 of 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.[4] To achieve this "stop-and-start" pattern

[4] A round of no concession following each round of maximal concession may be interpreted as temporary "cooling off" in a negotiation process. For negotiation practitioners, such cooling off is often seen as necessary to keep disruptive emotions in check and avoid break-downs, and sometimes as a useful negotia-

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). 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 common benefit of implementing the correct alternative in a common interest state is 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 an assumption made on the delay mechanisms 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. When more general dynamic delay mechanisms are allowed in which a delay up to the same limited commitment bound can be imposed after the two members concede to each other, we show how to improve the ex ante welfare of the committee over the optimal mechanism characterized in Section 3. In spite of the improvements, however, the design problem of the optimal general delay mechanism remains qualitatively similar, but a complete characterization is beyond the scope of this paper.

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

tion tactic (see, for example, Adler, Rosen and Silverstein, 1998) . Our characterization of the stop-and-start feature of optimal delay mechanism provides an alternative explanation.

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

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 $\lambda > 0$ is the loss from making the wrong decision relative to a fair 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 initial belief that the state is M.[5] 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 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

[5] An 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 $(1-\gamma_1)/(2-\gamma_1)$.

from voting one's own favorite:

$$\gamma_1 \left(x_1(1-2\lambda) + (1-x_1)(1-\lambda) \right) + 1 - \gamma_1 \geq \gamma_1 \left(x_1(1-\lambda) + (1-x_1) \right) + (1-\gamma_1)(1-\lambda).$$

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. In fact, using a standard application of the revelation principle one can show that a mechanism without transfers is incentive compatible if and only if the probability of implementing each alternative is independent of the true state.[6] Thus, in the absence of side transfers, there is no incentive compatible mechanism that Pareto dominates flipping a coin when $\gamma_1 > 1/2$. Thus, our model provides a stark environment that illustrates the severe restrictions on efficient information aggregation in committees when side transfers are not allowed.

2.2. Delay mechanisms

As suggested in our previous work (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 Section 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, which we refer to as a *one-round delay mechanism*, 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

[6] To see this, let q_R, q_M, and q_L denote the probabilities of implementing alternative r when the true states are R, M, and L, respectively, and let $\mathcal{Q}$ be the probability of implementing r when the reports are inconsistent, that is, when both players report that they are informed. The incentive constraints for an informed RIGHT and an informed LEFT, together imply $q_R \geq q_M$ and $q_M \geq q_L$. The incentive constraint for an uninformed RIGHT, together with the incentive constraint for an uninformed LEFT, imply that $(1-\gamma)(\mathcal{Q} - q_L) \geq \gamma(q_R - q_M)$ and $(1-\gamma)(q_R - \mathcal{Q}) \geq \gamma(q_M - q_L)$. Thus, incentive compatibility requires that $(1-\gamma)(q_R - q_L) \geq \gamma(q_R - q_L)$. This is inconsistent with $\gamma > 1/2$ unless $q_R - q_L = 0$, and so $q_R = q_M = q_L$ in any incentive compatible outcome.

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, but it captures the essential idea that 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 player more willing to vote for his opponent's favorite alternative but it increases the payoff loss whenever delay occurs.

It is straightforward to show that the optimal mechanism that maximizes the agents' ex-ante welfare is given by $\delta_1 = \Delta$ if the initial degree of conflict γ_1 is low and otherwise $\delta_1 = 0$. A larger cost reduces the equilibrium probability that the uninformed votes for his favorite alternative. The welfare gains from such reduction more than compensate the increased delay penalty in the case of a disagreement, and the net benefits are positive for both the uninformed and the informed. As a result, whenever it is possible and desirable to induce the uninformed to "concede" with positive probability, i.e. $x_1 < 1$, it is optimal to set δ_1 equal to the upper bound to induce the lowest possible x_1. Indeed, 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 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. For higher values of γ_1, the best a one-round delay mechanism can do is to set $\delta_1 = 0$, which is of course the same as flipping a coin.

If the binding constraint in using delay to help aggregate information is the maximum delay that can be imposed when the two players attempt to reach an agreement, can they do better by committing ex ante to repeated delay when they disagree? Imagine that we modify the original one-round mechanism by replacing the coin flip outcome after delay with a continuation one-round mechanism with some delay $\delta_2 \leq \Delta$. Suppose that in this *two-round delay mechanism* we can choose δ_2 such that the uninformed player obtains a continuation payoff in the second round exactly equal to the coin-flip payoff of $1 - \lambda$, but through equilibrium randomization with probability $x_2 < 1$ of voting for his ex ante favorite. 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 equilibrium payoff to the uninformed player is the same as in the one-round mechanism. However, because a smaller 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 a coin flip with $x_2 = 1$ 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 as in the original one-round mechanism but improves the payoff of the informed player.

That a two-round mechanism can improve over a one-round mechanism raises the question about what a general dynamic delay mechanism can achieve. We formally define a general dynamic delay mechanism as an extensive form game below. In the definition we denote as p the action profile in which each player votes for his favorite alternative (*persist*) in a given round, and with p^t the sequence of p's for round 1 through t. Further, let c be the action profile in a given round in which each player votes against his favorite alternative (*concede*), l (respectively, r) the action profile where LEFT (respectively, RIGHT) persists and RIGHT (respectively, LEFT) concedes, and a_t any action profile in round t. In the definition below, we assume there is a final round T. Later on in section 4.5 we remove this assumption and show that an optimal delay mechanism is necessarily finite.

Definition 1. *A sequence* $(\delta_1, \ldots, \delta_T)$*, with each* $\delta_t \in [0, \Delta]$*, defines a* **delay mechanism** *as the following extensive form game with imperfect information.*

1. *The players are* LEFT *and* RIGHT.
2. *The terminal histories are all sequences* (S, p^{t-1}, a_t)*, such that: i)* $S \in \{L, M, R\}$*; ii)* $t \leq T$*; and iii)* $a_t \neq p$ *whenever* $t < T$.
3. *The player function assigns the empty history to "chance" and every other proper subhistory to* {LEFT, RIGHT}.
4. *Chance chooses* L, M *or* R *with probability* $(1-\gamma_1)/(2-\gamma_1)$*,* $\gamma_1/(2-\gamma_1)$ *and* $(1-\gamma_1)/(2-\gamma_1)$ *respectively.*
5. *For each* $t \leq T$*, the information partition of* LEFT *is* $\{(L, p^{t-1})\}, \{(M, p^{t-1}), (R, p^{t-1})\}$*, and the information partition of* RIGHT *is* $\{(R, p^{t-1})\}, \{(M, p^{t-1}), (L, p^{t-1})\}$.
6. *Given a terminal history* (S, p^{t-1}, a_t) *the payoffs are as follows. If* $S = L$ *and* $a_t = l$ *or* $S = R$ *and* $a_t = r$*, the payoff is* $1 - \sum_{s \leq t-1} \delta_s$ *to both players. If* $a_t = c$*, the payoff is* $1 - \lambda - \sum_{s \leq t-1} \delta_s$ *to both players. If* $a_t = p$*, the payoff is* $1 - \lambda - \sum_{s \leq t} \delta_s$ *to both players. For all other cases, if* $a_t = l$ *(*$a_t = r$*) the payoff to* LEFT *(*RIGHT*) is* $1 - \sum_{s \leq t-1} \delta_s$ *and the payoff to* RIGHT *(*LEFT*) is* $1 - 2\lambda - \sum_{s \leq t-1} \delta_s$.

A delay mechanism is a simple multi-round voting game where in each round $t \leq T$, conditional on the game having not ended, each player chooses between voting for his favorite alternative and voting against it. 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; the game moves on to the next round if $t < T$, or ends with a coin flip if $t = T$.

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.[7]

[7]The main restriction we impose is symmetry. Generally there are asymmetric equilibria in which only one informed type concedes with a positive probability. Our approach is to impose symmetry and establish (in Section 4.5) that both informed types persist with probability 1 in any equilibrium. This is a more natural approach given the underlying committee problem set up in Section 2.1.

3. Results

For any $\gamma_1 \in (\gamma_*, 1)$, denote as $r_*(\gamma_1)$, or simply r_*, the smallest integer r satisfying

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

Define the "residue" η such that

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

By definition, $\eta \in (1, (\lambda + \Delta)/\lambda]$, and (3) holds as an equality if η is equal to the upper bound.

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

(a) for $\gamma_1 \leq \underline{\gamma}$, an optimal delay mechanism is:

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

(b) for $\gamma_1 \in (\underline{\gamma}, \overline{\gamma})$ such that $r_*(\gamma_1) = 1$, which always exists, an optimal delay mechanism is:

$$\delta_1 = \Delta, \quad \delta_2 = \ldots = \delta_{\tau+1} = (\gamma_1 \lambda - (1 - \gamma_1)\Delta)/\tau,$$
$$\delta_{2+\tau} = \Delta, \quad \textit{with } T = 2 + \tau,$$

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

(c) for $\gamma_1 \in (\underline{\gamma}, \overline{\gamma})$ such that $r_*(\gamma_1) \geq 2$, which exists if and only if $\Delta < (\sqrt{2} - 1)\lambda$, an optimal delay mechanism is:

$$\delta_{2t-1} = \Delta, \quad \delta_{2t} = \lambda\Delta/(\lambda + \Delta), \quad \textit{for } t = 1, 2, \ldots, r_* - 2,$$
$$\delta_{2(r_*-1)-1} = \Delta, \quad \delta_{2(r_*-1)} = \lambda(\eta - 1)/\eta,$$
$$\delta_{2r_*-1} = \lambda(\eta - 1), \quad \delta_{2r_*} = \delta_{2r_*+1} = \ldots = \delta_{2r_*+\tau-1} = \gamma_*\lambda/\tau,$$
$$\delta_{2r_*+\tau} = \Delta, \quad \textit{with } T = 2r_* + \tau;$$

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

(d) for $\gamma_1 \geq \overline{\gamma}$, an optimal delay mechanism is:

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

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

The most interesting features of optimal delay mechanisms are:

(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 1 in rounds $t, \ldots, T-1$, and his belief γ_T entering the last round is 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 1 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 total delay is bounded from above. 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 1 before it is incurred. However, this argument relies on the claim that an informed player persists with probability 1 regardless of the history of the play. We establish this claim and property (i) simultaneously in Section 4.5. 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. We relegate the proof of this property to the end of the analysis section (Section 4) because it does not rely on the localized variation method we use to characterize the next two properties of optimal delay mechanisms. Until then, we will focus on finite delay mechanisms with a deadline round T.

Property (ii) implies that an optimal delay mechanism generally induces a "stalling tactic" adopted by the uninformed types before the deadline, during which no attempt is made to reach a decision, until they make a "last minute concession" to take the opponent's favorite alternative when the deadline arrives. These two tactics are not contradictory, because the expectation of a high payoff from last minute concession at the deadline causes the uninformed types to stop any concession prior to the deadline. 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 a deadline belief γ_T of the uninformed player that is less than or equal to γ_*. Intuitively dynamic delay mechanisms work by driving down the uninformed player' belief that the state is a conflict state. Inducing a deadline belief greater than γ_* would imply that the uninformed player does not concede with probability 1 in the last round. As a result the Pareto efficient decision could not be achieved at the end, which cannot be optimal because adding more rounds for the uninformed player to have an opportunity to concede would improve the payoff of the informed. But driving down the uninformed player's belief through delay is costly. It does not pay to induce a deadline belief too much below γ_*. When $r_*(\gamma_1) \geq 2$ (i.e., in 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. It turns out that in general inducing the uninformed player to concede with a positive probability in two successive rounds is not optimal. In a delay mechanism, there are two ways to reduce the probability that the uninformed persists in a given round, by increasing the immediate delay and by decreasing the equilibrium continuation payoff after a disagreement. The first are "static" incentives, while the latter are "dynamic" ones only available in multi-round mechanisms. The logic behind the stop-and-start property is that maximizing the dynamic incentives requires the uninformed to persist for sure in the next round. This provides the greatest punishment in the event of a regular disagreement that induces the uninformed player to concede more in the current round. In an optimal dynamic mechanism, an "active round" of voting (in which the uninformed type concedes with positive probability) is always followed by an "inactive round" (in which the uninformed type concedes with zero probability). We show that this alternating pattern of start-and-stop drives the belief from γ_1 to γ_* in the smallest possible number of steps. Furthermore, the Main Result states that an optimal mechanism cannot have a delay equal to the upper bound in every round. The optimal delays should alternate between the limited commitment bound (to induce an "active round") and a level strictly below the bound (to induce an "inactive round").

4. Analysis

4.1. Preliminary results

In this subsection, we present two preliminary results that are the essential ingredients in our local-variation approach to characterizing the optimal delay mechanism. Both results restrict to equilibria in which the informed player persists after any history, and the first one in addition restricts to finite delay mechanisms. In Section 4.5, the restrictions are shown to be without loss of generality.

Denote x_t as the equilibrium probability that the uninformed player persists in round t in a game induced by a finite or infinite delay mechanism $(\delta_1, \ldots, \delta_T)$. Under the rules of our game, it ends immediately in round t whenever $x_t = 0$. Let γ_t be the equilibrium belief of the uninformed player that his opponent is uninformed (i.e., that the state is M) at the beginning of round t. Given the initial belief γ_1, in any equilibrium in which the informed player persists after any history, the belief in subsequent rounds is derived from Bayes' rule:

$$\gamma_{t+1} = \frac{\gamma_t x_t}{\gamma_t x_t + 1 - \gamma_t}. \tag{5}$$

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:

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

In the above, the top expression is the expected payoff from persisting and the bottom expression is the expected 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

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

An optimal delay mechanism maximizes $W_1(\gamma_1)$.

Given a finite delay mechanism $(\delta_1, \ldots, \delta_T)$, 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 $U_{T+1} = 1 - \lambda$ in the event that both players persist in 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 "localized variation method" to derive necessary conditions on an equilibrium induced by an optimal delay mechanism. This method presumes that equilibrium analysis depends on the incentives of the uninformed player alone. The analysis of optimal mechanisms, however, requires studying the payoffs to both uninformed and informed players. The following Lemma provides a link between the equilibrium payoffs of the informed and uninformed players. It is the first essential ingredient in our localized variation method.

Lemma 1. (LINKAGE LEMMA) *Suppose that a finite delay mechanism with deadline T induces an equilibrium in which the informed player persists after any history. If in the equilibrium it is a best response for an uninformed player to persist from some round $t \leq T$ onward, then*

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

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 $s = t, \ldots, T$. 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 best response, the uninformed player's payoff in 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 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.

The second ingredient in our localized variation method is a tight upper bound on how much concession in a given round t before the deadline that the uninformed player with belief γ_t can make in equilibrium. We say that round $t < T$ is *inactive* if $x_t = 1$; there is no updating of the belief of the uninformed player as neither the uninformed nor the informed makes any concession. Round $t < T$ is *active* whenever $x_t \in (0, 1)$, with the uninformed player updating his belief from γ_t to γ_{t+1} according to Bayes rule (5). The following lemma imposes lower bounds on x_t and γ_{t+1} in equilibrium as functions of γ_t only.

Lemma 2. (MAXIMAL CONCESSION LEMMA) *In any equilibrium in which the informed player persists after any history, for any $t < T$ and $\gamma_t > \Delta/(\lambda + \Delta)$, the lowest 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

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

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

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

In 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 on U_{t+1} and the bound Δ on δ_t, we obtain

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

Using Bayes' rule for γ_{t+1} and solving for x_t, we obtain $x_t \geq \chi(\gamma_t)$, which is positive if $\gamma_t > \Delta/(\lambda + \Delta)$. Since γ_{t+1} is increasing in x_t, plugging in the lowest value of x_t and using Bayes' rule give $\gamma_{t+1} \geq g(\gamma_t)$. ■

Maximal concession is attained in round t with belief γ_t in an equilibrium of a delay mechanism if and only if there is *no slack*, defined as follows:

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

Under the above definition, there is no slack in the static incentives in round t for truth-telling for the uninformed player, i.e. delay after regular disagreement is maximized at $\delta_t = \Delta$. Further, there is no slack in the dynamic incentives for truth-telling if $U_{t+1}(\gamma_{t+1}) = 1 - 2\gamma_{t+1}\lambda$, as the continuation payoff $U_{t+1}(\gamma_{t+1})$ for the uninformed player is minimized. The latter occurs in equilibrium if, after the regular disagreement in round t, the uninformed player persists with probability 1 in round $t+1$ with $x_{t+1} = 1$ but is indifferent between conceding and persisting. Thus minimizing the continuation payoff for the uninformed player entails that the probability of concession in the following round be zero. A delay mechanism with maximal concession in some round t necessarily results in stop-and-start equilibrium behavior.

Maximal concession as characterized in Lemma 2, or equivalently no slack in Definition 2, is an equilibrium property. Nonetheless, the lemma is suggestive of how to increase concession in a localized variation of an equilibrium of some delay mechanism. In particular, it suggests that the slack can be reduced by raising the *effective delay* incurred in round t after a regular disagreement. More precisely, if we denote the active arounds as $t(1) < t(2) < \ldots$, then

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

is the *effective delay* in round $t(i)$, which includes not only the immediate delay $\delta_{t(i)}$ but also all future delays in the inactive rounds between $t(i)$ and $t(i+1)$, which are incurred with certainty after the regular disagreement at round $t(i)$. In an active round t, if $\delta_t < \Delta$, we can raise δ_t to induce another equilibrium with more concession from the uninformed player, that is, a lower equilibrium x_t. If $\delta_t = \Delta$ but $U_{t+1}(\gamma_{t+1}) > 1 - 2\gamma_{t+1}\lambda$, we can try to achieve the same outcome by inserting an inactive round s with a sufficiently small delay δ_s between t and $t+1$ so as to raise the effective delay. This is possible because the continuation payoff $U_{t+1}(\gamma_{t+1})$ of the uninformed is not minimized in the original equilibrium. Indeed, adding inactive rounds is one of main localized variations we use in the following analysis.

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

$$\frac{1 - \gamma_{t+1}}{1 - \gamma_t} = \frac{\lambda + \Delta}{\lambda}. \tag{9}$$

Comparing this to equation (3), we see that it takes $r_*(\gamma_1)$ active rounds for the belief to reach from γ_1 to γ_* or below, if the uninformed is making maximal concessions in each active round.[8] A tighter commitment bound Δ would mean that it requires more active 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

In this subsection we show that in any optimal delay mechanism with at least two rounds 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 localized variation method. The intuition is that if a delay mechanism does not induce the Pareto efficient decision in the final round, then it is possible to slightly modify it 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 modified mechanism we consider first introduces an additional round, s, between the last active round $t(r)$ and the deadline T. The delay in round s, δ_s can be initially chosen so that, provided $\gamma_s = \gamma_T$, in the continuation equilibrium starting in round s, the payoff of the uninformed is $1 - 2\lambda\gamma_T$ and in equilibrium $x_s = 1$. Simultaneously, the delay in the last active round, $t(r)$, can be reduced by an amount equal to δ_s, so that the effective delay after a regular disagreement in round $t(r)$ remains unchanged. Thus, the strategy of the uninformed that maintains the same probability of persistence in every round other than s, and persists with probability 1 in the new round s is an equilibrium of the modified mechanism. By construction, in round s both there is slack and the uninformed is indifferent between persisting and conceding with $x_s = 1$. By the Maximal Concession Lemma, we can marginally reduce x_s from 1 by a small increase in δ_s in the continuation equilibrium starting in round s with $\gamma_s = \gamma_T$. To prevent this modification from increasing the payoff after a regular disagreement in round $t(r)$, we compensate by introducing yet another extra round s' with an appropriate delay $\delta_{s'}$ in between round $t(r)$ and round s. After the modification, the payoff in the event of a regular disagreement

[8]Define $g^{(n)}(\gamma)$ to 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 $g^{(r)}(\gamma) \leq \gamma_*$.

in round $t(r)$ remains fixed at the original value of $-\delta_{t(r)} + U_{t(r)+1}(\gamma_{t(r)+1})$. Because both the initial belief γ_1 and the continuation payoff in 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.[9] In particular, the modified mechanism does not affect $U_1(\gamma_1)$.

Since $x_T > 0$, as long as the uninformed persists with probability close to 1 in round s in the modified mechanism, in equilibrium the probability of persistence in the last round will still be strictly positive, and it is a best response for the uninformed to persist in every round. In this localized variation, to lower the probability of persistence below 1 in round s requires raising the delay in round s. As the probability of persistence in round s decreases, the uninformed player's payoff increases, and therefore the delay in round 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 greater 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 with at least two rounds, $x_T = 0$.*

An immediate corollary to Proposition 1 is that the deadline belief satisfies $\gamma_T \leq \gamma_*$. The next result establishes a counterpoint to Proposition 1. Although the Pareto efficient decision becomes achievable in the final round so long as a delay mechanism drives the degree of conflict from γ_1 to $\gamma_T \leq \gamma_*$, it is generally too costly to force γ_T to a level too much below what is needed to achieve efficiency. Specifically, if it takes two or more rounds to drive γ_T 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 1.

Proposition 2. *In an optimal mechanism with at least two active rounds, $\gamma_T = \gamma_*$.*

The argument used to establish Proposition 2 starts with a mechanism that yields $\gamma_T < \gamma^*$ in equilibrium, and modifies it so that the belief in the last round increases but without going above γ_*, with the uninformed still conceding in the last round. The Linkage Lemma cannot be used for an evaluation of welfare changes because when the deadline belief is below γ_*, the uninformed might strictly prefer concession in the last round.

[9] In constructing this modified mechanism, we take the belief at the beginning of round s to be fixed at the original value of $\gamma_{t(r)+1} = \gamma_T$ and find the delay in round s that would induce $x_s < 1$. This is justified because the modified mechanism does not change the equilibrium play prior to round s, ensuring that the belief at the beginning of round s is indeed fixed.

To obtain our welfare comparison, we keep the expected payoff of the uninformed constant throughout the modification. This requires the fewer concessions in the equilibrium of the modified mechanism to be compensated by less delay after regular disagreements for the uninformed player. We then directly establish that the savings in delay that maintains the payoff constant for the uninformed more than compensate the informed for the fewer concessions before an agreement is reached.

The formal proof of the proposition, provided in the Appendix, distinguishes between the case when in the original mechanism the belief goes below γ_* before T and the case when it does so in round T. In the first case, in the modified mechanism the effective delay in the last active round $t(r)$ is reduced to the level so that the uninformed is indifferent between persisting and conceding at $\tilde{x}_{t(r)} = 1$. This is achieved by setting $\tilde{\sigma}_{t(r)} = \gamma_{t(r)}\lambda$. The effective delay of the previous active round $t(r-1)$ is also reduced to keep the continuation payoff in round $t(r-1)$ unchanged.[10] In the second case, the effective delay in the last active round is reduced enough so that the equilibrium probability of concession brings the deadline belief to exactly γ_*. Again, the effective delay of the previous active round is also reduced to keep the continuation payoff of the uninformed in round $t(r-1)$ unchanged. In both cases, the modified mechanism has an equilibrium where the uninformed behavior only changes in round $t(r)$ and, by construction, the expected payoff of the uninformed is left unchanged.

The payoff changes for the informed player occur only to the continuation payoff in the next-to-last active round (i.e., round $t(r-1)$). The tradeoff is between a lower probability of persistence $x_{t(r)}$ in the last active round $t(r)$ in the original mechanism, versus a smaller delay both in round $t(r)$ and in round $t(r-1)$ for the modified mechanism. Using the equilibrium conditions that the uninformed is indifferent between conceding and persisting in round $t(r)$ and that the change in effective delay leaves his continuation payoff in round $t(r-1)$ unchanged, we can show that the informed is strictly better off in the modified mechanism.

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

[10] This is why the proposition requires at least two active rounds. In case (b) of Main result, there is a single active round before T and we have $\gamma_T < \gamma_*$ under an optimal delay mechanism. But even in this case, γ_T cannot be below $g(\gamma_*)$.

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. The key result is the following lemma.

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

The presence of slack in two consecutive active rounds $t(i)$ and $t(i+1)$ means that it is possible, by appropriately changing the effective delays, to both increase and decrease the probabilities of persistence, $x_{t(i)}$ and $x_{t(i+1)}$, of the uninformed in the two rounds. In the localized variation argument used to proof Lemma 3 (the details are provided in the Appendix,) we modify the effective delays in both round $t(i)$ and $t(i+1)$ while maintaining $x_{t(i)}x_{t(i+1)}$ constant. From repeated applications of Bayes' rule (5), this guarantees that the belief $\gamma_{t(i+2)}$ after the regular disagreement in round $t(i+1)$, and hence the continuation equilibrium, is left unchanged (we set $t(i+2) = T$ if $i = r-1$) . Finally, an appropriate change in the effective delay in round $t(i-1)$, keeps the continuation payoff after a regular disagreement in round $t(i-1)$ constant; this explains why we assume that $t(i)$ is not the first active round. As a result, the original equilibrium sequence $\{\gamma_t, x_t, U_t\}$ remains unchanged in the modified mechanism from round 1 to $t(i-1)$ and from round $t(i+2)$ to T. By the Linkage Lemma, we only need to calculate the total delay between round $t(i-1)$ and $t(i+2)$ to evaluate the effect of this modification on the ex ante payoff.

One can think of this localized variation exercise as choosing $\gamma_{t(i+1)}$ to maximize the total delay, while holding fixed $\gamma_{t(i)}$ and $\gamma_{t(i+2)}$ (as well as the continuation payoff in round $t(i-1)$). From the Maximal Concession Lemma, the feasible set for $\gamma_{t(i+1)}$ 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 $\gamma_{t(i+1)} = g(\gamma_{t(i)})$ (no slack in round $t(i)$) or at $\gamma_{t(i+1)} = g^{-1}(\gamma_{t(i+2)})$ (no slack in round $t(i+1)$). It turns out that these two choices of $\gamma_{t(i+1)}$ entail the same total delay, and are therefore payoff-equivalent.

Lemma 4. *Let $i \geq 2$. Holding fixed $x_{t(i)} x_{t(i+1)}$, and $x_{t(s)}$ for all $t(s) \neq t(i), t(i+1)$, a mechanism with slack in round $t(i)$ but no slack in round $t(i+1)$ is payoff-equivalent to a mechanism with no slack in round $t(i)$ but slack in round $t(i+1)$.*

Clearly, in any optimal delay mechanism, the first active round is round 1. In fact, the next result shows that it is optimal to have no slack and to induce maximal concession by the uninformed player in the first round, if there are at least two active rounds.

Lemma 5. *In an optimal mechanism with at least two active rounds, there is no slack in the first round.*

The localized variation construction used in the proof of Lemma 5 is similar to that used for Lemma 3. Suppose that there is slack in the first round. In the modified mechanism effective delay in the first round is raised to lower the probability of persistence x_1. Simultaneously, the effective delay in the second active round is reduced to increase the probability of persistence $x_{t(2)}$ to keep the belief after two active rounds constant. However, unlike the localized variation in the proof of Lemma 3, in the modified mechanism the expected payoff of the uninformed necessarily increases. This is because lowering x_1 raises the expected payoff of the uninformed in round 1, $U_1(\gamma_1)$, but there is no previous round for which we can increase the effective delay to compensate for the improvement in the uninformed payoff. Also, the Linkage Lemma does not apply. Instead we calculate the changes 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. *In an optimal mechanism, there can be at most one active round with slack. If there are at least two active rounds, there is no slack in the first round.*

Proof. Let $t(j)$ be the first active round with slack and $t(j')$ be the last one with slack. By Lemma 5, $j \geq 2$. 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 first round with slack is $t(j+1)$ and the last 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 first round with slack is $t(j+2)$ and the last 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, it cannot be the first round. ■

Proposition 3 establishes that all (but possibly one) active rounds in an optimal mechanism are without slack. By the Maximal Concession Lemma, this means that the probability of concession by the uninformed player is maximized in these active rounds. The efficient deadline belief γ_* is reached in the least possible number of active 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 player is making no concession (i.e., $x_{t(i)+1} = 1$) right after the round when he made the maximal concession. In other words, an active round must be followed by an inactive round. After that, in the next active 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. Slack in one round might be needed to prevent the deadline belief from becoming inefficiently low.

4.4. Optimal delay mechanism

In this subsection we use the properties of optimal delay mechanisms established so far to derive a complete characterization of all optimal mechanisms. As the initial belief varies, the minimum number of active rounds needed for the belief to reach γ_* changes. We distinguish two cases.

First, consider the case where $r_*(\gamma_1) \geq 2$, and assume that $T \geq 2$. 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 active round. Such belief evolution is determined by equation (9); hence it takes at least $r_*(\gamma_1)$ active 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 $r_*(\gamma_1)$ active rounds in an optimal mechanism.

Recall also that the "residue" η defined in (4) satisfies $1 < \eta \leq (\lambda + \Delta)/\lambda$. If $\eta = (\lambda + \Delta)/\lambda$, the belief reaches from γ_1 to γ_* with r_* rounds of randomization, with no slack in any of the r_* rounds. If $\eta < (\lambda + \Delta)/\lambda$, to satisfy the restriction imposed by Proposition 3, we need $r_* - 1$ 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 active round before the deadline.[11]

[11] Although it is optimal to have slack in any round $t(j)$, $j = 2, \ldots, r^*$, choosing $j = r^*$ 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.

For each $i = 1, \ldots, r_* - 2$, 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 in 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 active rounds is constant across i. How this sum is distributed across the intervening rounds is immaterial. However, since the sum is less than Δ, the optimal mechanism can be implemented with just one intervening inactive round between any two successive active rounds. 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 in round $t(r_* - 1)$ requires:

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

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

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

Given the above expression, solving (10) gives:

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

Since $\eta \leq (\lambda + \Delta)/\lambda$, the above is less than or equal to $\Delta/(\lambda + \Delta)\lambda$. So the optimal mechanism can be implemented with just one intervening inactive 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.

In round $t(r_*)$, there is generally slack, unless η is equal to its upper bound. The equilibrium probability of persisting $x_{t(r_*)}$ is given by equation (11) above, and by construction the belief after a regular disagreement would become γ_*. To find the effective delay in round $t(r_*)$, we use the fact that $U_T(\gamma_T) = 1 - \gamma_* \lambda$ and the indifference condition in round $t(r_*)$ to obtain:

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

Although how the above effective delay is distributed is payoff-irrelevant, we can always choose $\delta_{t(r_*)} = (\eta - 1)\lambda \leq \Delta$ 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, with no concession from the uninformed.[12] This corresponds to the third line of the mechanism described in case (c) of our Main Result.

Finally, in 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 in case (c) of our Main Result.

Summing over all rounds, the total delay is

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

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

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

The payoff to the informed player $V_1(\gamma_1)$ can be obtained using equation (8). This completes the derivation of the optimal delay mechanism assuming that $r_*(\gamma_1) \geq 2$ and $T \geq 2$ so that the ex ante payoff $W_1(\gamma_1)$ is greater than the coin flip payoff $1 - \lambda$ (i.e., case (d) of Main Result does not apply.)

[12] The 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.

Next, consider the second case of $r_*(\gamma_1) = 1$, and again assume that $T \geq 2$ so that the ex ante payoff $W_1(\gamma_1)$ is greater than the payoff from the one-round delay mechanism with maximal delay Δ (i.e., case (a) of Main Result does not apply.) In this case, the first round is the only active round. To see this, suppose that an optimal mechanism induces two or more active rounds prior to the final round. By Proposition 2, $\gamma_T = \gamma_*$, and by Proposition 3, there is no slack in all except one of these active rounds, which contradicts the assumption that $r_*(\gamma_1) = 1$.

Since the first round is the only active round, and since $r_*(\gamma_1) = 1$, we have $\gamma_T \leq \gamma_*$. From the indifference condition for the uninformed, the effective delay σ_1 associated with such a mechanism satisfies $x_1\sigma_1 = \gamma_T\lambda$. Now, consider the payoff to the uninformed. Since $U_1(\gamma_1) = 1 - \gamma_1\lambda - \gamma_1 x_1\lambda$, his payoff is maximized when x_1 is minimized. The payoff to the informed is $V_1(\gamma_1) = 1 - x_1\sigma_1 = 1 - \gamma_T\lambda$, which is maximized when x_1 is minimized. Thus, there is no slack in round 1. By the Maximal Concession Lemma, an optimal sequence of delays that can implement this outcome is given by $\delta_1 = \Delta$, and

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

The above corresponds to the first line in case (b) of Main Result, and is the *deadline play* in this case. At the final round, choosing $\delta_T = \Delta$ induces $x_T = 0$, corresponding to the second line in case (b) of Main Result.

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 (c) and (b)) and compare them with the payoffs from a coin flip (case (d)) and from the one-round delay mechanism with maximum delay (case (a)). By construction, there is no overlapping between cases (c) and (b), which require $r_*(\gamma_1) \geq 2$ and $r_*(\gamma_1) = 1$ respectively. Intuitively, case (d) applies when γ_1 is sufficiently close to 1 so that the benefit from a Pareto optimal decision through a delay mechanism given in case (c), with start-and-stop and deadline play, is overwhelmed by the cost of delay incurred. Likewise, case (a) applies when γ_1 is sufficiently close to γ_* so that the benefit from a dynamic delay mechanism given in case (b), with a Pareto optimal decision and deadline play, does not justify the additional cost of delay relative to a static delay mechanism. The calculations are relegated to the Appendix.

Proposition 4. *There exist $\underline{\gamma}$ and $\overline{\gamma}$, with $\gamma_* < \underline{\gamma} < \overline{\gamma} < 1$, $r_*(\underline{\gamma}) = 1$, and $r_*(\overline{\gamma}) \geq 2$ 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 $\gamma_1 \geq \overline{\gamma}$.*

4.5. Finite delay

Our localized variation approach assumes that the informed player persists after each history and thus equilibrium analysis depends on the incentives of the uninformed player alone. Intuitively, whenever the uninformed player weakly prefers persisting to conceding, the informed player strictly does so. In this subsection, we validate this intuition and show that this is in fact the case in any equilibrium of any delay mechanism. We achieve this using a backward induction argument, which requires us to establish the first property of the optimal delay mechanism characterized in Main Result, that it is finite.

An *infinite* delay mechanism may be represented with any infinite sequence $\{\delta_t\}_{t=1}^{\infty}$. Assuming that the payoff from never implementing either alternative is worse than the payoff from implementing any alternative, we show that any infinite delay mechanism is *effectively finite*. That is, there exists an N such that the induced game ends with probability 1 before round N in any equilibrium.

Lemma 6. *Any infinite delay mechanism is effectively finite.*

Lemma 6 follows from a series of claims. First, we show that in any equilibrium that induces a distribution of end times with unbounded support, the informed is persisting with probability one at all times. This is done by showing that at any time when it is optimal for the uniformed to persist it is uniquely optimal for the informed to do the same. This property implies that in any such equilibrium, the belief γ_t of the uninformed player is eventually decreasing and hence will converge. The conclusion that the mechanism is effectively finite follows from demonstrating that γ_t must converge to zero, which makes it optimal eventually for the uninformed to concede with probability one.

By Lemma 6, 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 1 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. Instead, it is driven by the feature in our model that the informed player knows the state is a common interest state and that his opponent is uninformed and should concede.

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 a 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 for the differences in the payoff structure in the latter model).

As Δ goes to 0, from (13) we have

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

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

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

which implies

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

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

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

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

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

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

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

Furthermore, the optimal deadline mechanism entails the boundary condition $V^*(1/2) = 1 - \lambda/2$. Solving this differential equation gives $V^*(\gamma) = \lim_{\Delta\to 0} V_1(\gamma)$.

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	$(\overline{v}+\overline{\beta},\overline{v})$	$(\underline{v}+\beta,\underline{v})$	$(\underline{v}+\underline{\beta},\underline{v})$
r	$(\underline{v},\underline{v}+\underline{\beta})$	$(\underline{v},\underline{v}+\beta)$	$(\overline{v},\overline{v}+\overline{\beta})$

In the above table, $\overline{v}$ 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, $\underline{v} < \overline{v}$ 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 $\overline{\beta}$ 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 $\overline{\beta} = \underline{\beta} = 0$ and $\beta = \overline{v} - \underline{v}$ (with $\overline{v} = 1$ and $\underline{v} = 1 - 2\lambda$), while in Damiano, Li and Suen (2012) we have $\overline{\beta} = \underline{\beta} = \beta > 0$.

By allowing the private value components $\overline{\beta}$, $\underline{\beta}$ 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 R to be common interest states, we need $\underline{\beta} < \overline{v} - \underline{v}$. This assumption also ensures that in delay mechanisms the informed player always persists 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 $\overline{v} + \overline{\beta}$ 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, $(\overline{v} + \underline{v} + \overline{\beta})/2$, is different from the coin-flip payoff of for the uninformed player, which is $\underline{v} + \beta/2$ in the conflict state and $(\overline{v} + \underline{v} + \underline{\beta})/2$ 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 localized 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 localized 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 $\beta/2$. 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.[13]

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 round of $t(i)$.

[13]Lemma 4 also hods. 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 localized variation argument are independent of the payoff structure for the informed.

5.3. General delay mechanisms

Consider a more general mechanism design problem where, in addition to the sequence of delays $\{\delta_t\}_{t=1}^T$ after each regular disagreement in round t, a continuation payoff $\mathcal{U}_t$ after a reverse disagreement in round t has to be chosen as well, where $\mathcal{U}_t$ ranges from $1-\lambda$, which is what we have assumed so far, to $1-\lambda-\Delta$, depending on the delay imposed before flipping a coin. This way of modeling general delay mechanisms after any disagreement is without loss of generality, as long as we restrict to symmetric Bayesian equilibria in finite delay mechanisms in which the informed player persists regardless of history.

The Linkage Lemma (Lemma 1) remains valid in this more general design problem, because $\mathcal{U}_t$ does not appear in the expected payoff of an uninformed player that always persists on the equilibrium path. Since it is based on Lemma 1, Proposition 1 also holds. The argument here amends the original proof by first establishing that $\mathcal{U}_T = 1-\lambda$ if $x_T > 0$. To see this, suppose not. If $x_T = 1$ and the uninformed strictly prefers persisting to conceding, we can keep raising $\mathcal{U}_T$ until either $\mathcal{U}_T = 1-\lambda$, or the uninformed becomes indifferent. If $x_T < 1$, the indifference condition is

$$1-\gamma_T\left(1-\mathcal{U}_T\right)-\gamma_T x_T\left(2\lambda-\left(1-\mathcal{U}_T\right)\right)=1-\left(\gamma_T x_T+1-\gamma_T\right)\left(\delta_T+\lambda\right).$$

Solving for x_T, we have

$$x_T=\frac{\gamma_T\left(1-\mathcal{U}_T\right)-\left(1-\gamma_T\right)\left(\delta_T+\lambda\right)}{\gamma_T\left(1-\mathcal{U}_T-\lambda+\delta_T\right)}.$$

For fixed γ_T, an increase in $\mathcal{U}_T$ decreases x_T and thus increases $U_T(\gamma_T)$. Since the effect of an increase in $U_T(\gamma_T)$ on the last active round $t(r)$ before T can be neutralized for the uninformed by adding auxiliary rounds between $t(r)$ and T, it follows immediately from Lemma 1 that the informed are better off with any change to the payoff of the uninformed, a contradiction establishing that an optimal general mechanism has $\mathcal{U}_T = 1-\lambda$. The rest of the proof of Proposition 1 goes through without change.

Given that Proposition 1 holds and thus $x_T = 0$, we can improve our optimal delay mechanism when it has just one active round before the deadline round T, by lowering the continuation payoff $\mathcal{U}_T$ after a reverse disagreement in the deadline round. To be precise, suppose that $r_*(\gamma_1) = 1$ and we are in case (b) of Main Result. Our original construction requires the first round to be active with no slack, implying that $\gamma_T < \gamma_*$. Now, consider reducing $\mathcal{U}_T$ marginally to $\tilde{\mathcal{U}}_T$ and simultaneously reducing σ_1 to $\tilde{\sigma}_1$, such that

$$-\tilde{\sigma}_1 + 1 - \gamma_T(1 - \tilde{\mathcal{U}}_T) = -\sigma_1 + U_T(\gamma_T).$$

Then, the uninformed still strictly prefers conceding to persisting in the deadline round with the same γ_T, and thus it remains an equilibrium in the first round with the same x_1 and $U_1(\gamma_1)$. Moreover,

$$\tilde{V}_1(\gamma_1) = 1 - x_1\tilde{\sigma}_1 > 1 - x_1\sigma_1 = V_1(\gamma_1),$$

so the informed is better off.

When the optimal delay mechanism given in Main Result has at least two active rounds before the deadline, it generally has one active round—not the first one—with slack, and we can adjust the continuation payoff $\mathcal{U}_t$ after the reverse disagreement in the round with slack to improve on it. More precisely, suppose that $r_*(\gamma_1) \geq 2$, so we are in case (c) of Main Result with $\eta < (\lambda + \Delta)/\lambda$. By Lemma 4 and Lemma 5, the total amount of slack is then $\Delta - \lambda(\eta - 1)$, independent of which active round it appears in except for round 1. In the equilibrium constructed in Main Result, the last active round $t(r_*)$ (round $2r_* - 1$ in case (c) of Main Result) is the one with slack, where $\delta_{t(r_*)} = \lambda(\eta - 1) < \Delta$, belief $\gamma_{t(r_*)}$ satisfies $(1 - \gamma_*)/(1 - \gamma_{t(r_*)}) = \eta$ and $x_{t(r_*)}$ is given by equation (11), with the indifference condition:

$$\begin{aligned} U_{t(r_*)}(\gamma_{t(r_*)}) &= \gamma_{t(r_*)}\left(x_{t(r_*)}(1-2\lambda) + (1 - x_{t(r_*)})\mathcal{U}_{t(r_*)}\right) + 1 - \gamma_{t(r_*)} \\ &= \left(\gamma_{t(r_*)}x_{t(r_*)} + 1 - \gamma_{t(r_*)}\right)\left(-\sigma_{t(r_*)} + U_{t(r+1)}(\gamma_{t(r+1)})\right) + \gamma_{t(r_*)}(1 - x_{t(r_*)}). \end{aligned}$$

Consider reducing $\mathcal{U}_{t(r_*)}$ marginally from $1 - \lambda$ to $\tilde{\mathcal{U}}_{t(r_*)}$ and simultaneously increasing $\sigma_{t(r_*)}$ to $\tilde{\sigma}_{t(r_*)}$, such that the above indifference condition holds at the same $x_{t(r_*)}$ and $U_{t(r+1)}(\gamma_{t(r+1)})$. This is feasible because there is slack in round $t(r_*)$, with $\delta_{t(r_*)} < \Delta$. These modifications of $\mathcal{U}_{t(r_*)}$ and $\sigma_{t(r_*)}$ reduce the continuation payoff for the uninformed in round $t(r_*)$ from $U_{t(r_*)}(\gamma_{t(r_*)})$ to $\tilde{U}_{t(r_*)}(\gamma_{t(r_*)})$, but since $t(r_*)$ is not the first active round, the effects can be neutralized by reducing the effective delay $\sigma_{t(r_*-1)}$ in round $t(r_* - 1)$ by

$$\tilde{U}_{t(r_*)}(\gamma_{t(r_*)}) - U_{t(r_*)}(\gamma_{t(r_*)}) = (\gamma_{t(r_*)}x_{t(r_*)} + 1 - \gamma_{t(r_*)})(\tilde{\sigma}_{t(r_*)} - \sigma_{t(r_*)}).$$

The change in the total delay is positive because $x_{t(r_*)} < 1$, and thus, by Lemma 1, the payoff to the uninformed is kept unchanged while that to the informed is increased.

In spite of the improvements we can make upon the optimal delay mechanism characterized by Main Result through adjusting the continuation payoff $\mathcal{U}_t$ after a reverse disagreement, the design problem of the optimal general delay mechanism remains qualitatively similar. An important reason for this that the Maximum Concession Lemma (Lemma 2) continues to characterize the lowest feasible x_t and γ_{t+1} for any $t < T$ and γ_t sufficiently high. To see this, note that to obtain the lower bounds, in addition to maximizing the current delay by choosing $\delta_t = \Delta$, it is now necessary to maximize the current reverse disagreement payoff by choosing $\mathcal{U}_t = 1 - \lambda$, as well as minimize the payoff that the uninformed gets in round $t+1$ by conceding after a regular disagreement

$$U_{t+1}(\gamma_{t+1}) = \gamma_{t+1}\left(x_{t+1}(1-2\lambda) + (1-x_{t+1})\mathcal{U}_{t+1}\right) + 1 - \gamma_{t+1}.$$

By Assumption (1), since $\mathcal{U}_{t+1}$ is less than $1-2\lambda$, the lowest payoff the uninformed can guarantee after a regular disagreement remains $1 - 2\gamma_{t+1}\lambda$. Because Lemma 2 continues to hold, the important feature of stop-and-start of our optimal delay mechanism is unchanged. In the current proof of the result (Lemma 3), we show that for fixed starting belief $\gamma_{t(i)}$ and ending belief $\gamma_{t(i+2)}$, the total delay implied by the three indifference conditions in rounds $t(i)$, $t(i+1)$ and $t(i+2)$ is convex in the belief $\gamma_{t(i+1)}$. This remains true in the general design problem with choices of continuation payoffs after reverse disagreements, because for fixed choices of $\mathcal{U}_{t(i)}$ and $\mathcal{U}_{t(i+1)}$, the total delay required to update the belief from $\gamma_{t(i)}$ to $\gamma_{t(i+2)}$ in any equilibrium can be shown to be convex in $\gamma_{t(i)+1}$.

A complete characterization of the optimal general delay mechanism is complicated by the fact that Proposition 2 no longer holds. That is, it may not be optimal to end the game at the highest belief that induces concession with probability by the uninformed. This is an implication of our earlier observation that if there is only a single active round before the deadline T under the optimal delay mechanism given in case (b) of Main Result, the mechanism can be improved upon by pushing down the deadline belief γ_T through some delay after the final reverse disagreement (i.e., by lowering $\mathcal{U}_T$). At the critical initial belief $\gamma_1 = g^{-1}(\gamma_*)$ when the optimal delay mechanism given by Main Result switches from case (b) to case (c), this improvement remains valid, so by continuity when γ_1 is just above the critical belief, it is not optimal to have a deadline belief equal to γ_* if the continuation payoff after the last reverse disagreement can be adjusted. We will not attempt to characterize the optimal general delay mechanism, as it adds little insight beyond what we have discussed above.

Appendix

Proof of Proposition 1. Suppose by way of contradiction that $x_T > 0$ in an optimal mechanism. This implies $\gamma_T \geq 1/2$. Let t' be the last round before T for which conceding is weakly preferred to persisting. If there is slack in the last active 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 for the mechanism to be optimal. We first obtain some bounds on the total delay from round t' onward.

Claim 1. Let $t' < T$ be the last round before T in which the uninformed player weakly prefers conceding to persisting. Then, $\delta_{t'} + \sum_{t=t'+1}^{T-1} \delta_t \geq U_T(\gamma_T) - (1 - 2\gamma_{t'}\lambda)$. Further, $\sum_{t=t'+1}^{T-1} \delta_t \leq U_T(\gamma_T) - (1 - 2\gamma_T\lambda)$, with strict inequality if $t' - 1 < T$.

To establish the first inequality of Claim 1, note that in 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 $\gamma_{t'} > \gamma_T \geq 1/2$, 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 $x_{t'} = 1$ gives the first inequality of Claim 1. The second inequality is obtained from the condition that uninformed player strictly prefers persisting to conceding in round $t' + 1$ if $t' + 1 < T$. If $t' = T - 1$ the inequality follows from the uninformed weakly prefers persisting to conceding in round T, with $\sum_{t=t'+1}^{T-1} \delta_t = 0$.

Next, we make the following modification to the original mechanism:

1. Introduce two "extra rounds" s' and s between t' and $t' + 1$ (round s' before s). Set $\delta_{s'} = 0$, and $\delta_s = U_T(\gamma_T) - (1 - 2\gamma_T\lambda) - \sum_{t=t'+1}^{T-1} \delta_t \equiv \zeta$. 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 $\zeta < \delta_{t'} \leq \Delta$ and $\zeta \geq 0$. Given this construction, the effective delay $\sigma_{t'}$ remains unchanged in round t'. Moreover, the uninformed is just indifferent between persisting and conceding with $x_s = 1$ in round s. In 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 in 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 in 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 $x_{t'+1} = \ldots = x_T = 1$, the uninformed is indifferent between persisting and conceding in 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:

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

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

In case (ii), the uninformed player is indifferent between persisting and conceding in round T under the original mechanism. In this case, lowering γ_T would change the value of $U_T(\gamma_T)$. The indifference condition in 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 $U_T(\gamma_T) = 1 - \gamma_T\lambda - \gamma_T x_T \lambda$, and γ_T depends on x_s through Bayes' rule. Differentiate with respect to x_s to obtain:

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

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

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

because the indifference condition in round T ensures that $\mathrm{d}x_T/\mathrm{d}\gamma_T > 0$. Thus we again have $\mathrm{d}\delta_s/\mathrm{d}x_s < 0$ in this case.

For both cases (i) and (ii), when x_s falls, other things being equal, this would increased $U_s(\gamma_s)$. 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 $\mathrm{d}\delta_{s'}/\mathrm{d}x_s = -\gamma_s\lambda < 0$. 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

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

Moreover, since $x_T > 0$ in step 1, we have $\tilde{x}_T > 0$ in step 2 by choosing ϵ to be small. Therefore, it is a best response for the uninformed player to persist throughout in the modified mechanism. Finally, by construction, $U_1(\gamma_1)$ remains unchanged by our localized 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 $\gamma_T < \gamma_*$. Let $t(r)$ be the last active round. There are two cases: (i) $\gamma_{t(r)} \leq \gamma_*$; or (ii) $\gamma_{t(r)} > \gamma_*$. By assumption, there exists another active 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 in round $t(r-1)$ fixed. This can be 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 given by $\tilde{\gamma}_{t(r)+1} = \ldots = \tilde{\gamma}_T = \gamma_{t(r)}$. 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. Therefore step 1 of the construction is feasible.

From the indifference condition in 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 persists through rounds $t(r-1), \dots, T-1$ and concedes in round T (because $\tilde{\gamma}_T = \gamma_{t(r)} < \gamma_*$) in the equilibrium of the modified mechanism, the payoff to the informed in 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 in 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

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

2. Change $\sigma_{t(r-1)}$ in such a way to keep the continuation payoff for the uninformed in 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 in 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. Therefore, 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:

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

where $h(\gamma) \equiv \gamma(\gamma - \gamma_{t(r)})/(1-\gamma)$. 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 when $\tilde{\gamma}_T$ is chosen to be equal to γ_*,

$$\tilde{V}_{t(r-1)}(\gamma_{t(r-1)}) > V_{t(r-1)}(\gamma_{t(r-1)}).$$

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. ∎

Proof of Lemma 3. From the indifference condition in 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 beliefs (and set $t(i+2) = T$ when $i = r-1$)

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

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

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

For fixed $\gamma_{t(i)}$ and $\gamma_{t(i+2)}$, we can use Bayes' rule repeatedly to transform to 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 keep $\gamma_{t(i)}$ 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

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

This is a convex function of $\gamma_{t(i+1)}$, so it attains the maximum at a boundary of the feasible set. If there is 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. In the two mechanisms both $\gamma_{t(i)}$ and the posterior belief after active round $t(i+1)$, $\gamma_{t(i+2)}$, are the same. 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 in round $t(i+1)$) and the point $\gamma_{t(i+1)} = g(\gamma_{t(i)})$ (i.e., no slack in 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 in round $t(2)$ to write the payoff to the uninformed player, the indifference condition in 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,

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

Lowering x_1 thus requires raising σ_1. Since there is slack in round 1, raising σ_1 is feasible.

Similarly, the indifference condition in 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,

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

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

$$\frac{\mathrm{d}U_1(\gamma_1)}{\mathrm{d}x_1} = -\gamma_1 \lambda < 0.$$

Thus, lowering x_1 (while holding $x_1 x_{t(2)}$ 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 best response to the equilibrium strategy of the uninformed. Therefore equation (8) holds. The effect of a change in x_1 (while holding $x_1 x_{t(2)}$ constant) on the payoff of the informed can be calculated by:

$$\begin{aligned}\frac{\mathrm{d}V_1(\gamma_1)}{\mathrm{d}x_1} &= \frac{1}{\gamma_1}\frac{\mathrm{d}U_1(\gamma_1)}{\mathrm{d}x_1} + \frac{1-\gamma_1}{\gamma_1}\left(\frac{\mathrm{d}\sigma_1}{\mathrm{d}x_1} + \frac{\mathrm{d}\sigma_{t(2)}}{\mathrm{d}x_1}\right)\\ &= \frac{1-\gamma_1}{\gamma_1}\frac{\mathrm{d}\sigma_1}{\mathrm{d}x_1} + \left(-1 + \frac{1-\gamma_1}{1-\gamma_1+\gamma_1 x_1 x_{t(2)}}\right)\lambda < 0.\end{aligned}$$

Hence, both the informed and the uninformed can be made better off. ∎

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

$$U_1 = 1 - \gamma_1\lambda - \gamma x_1\lambda$$

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

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

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

Next, fix any $\gamma_1 \in (\gamma_*, g^{-1}(\gamma_*))$. Note that $g^{-1}(\gamma_*) \in \big(g^{-1}(1/2), (\lambda+\Delta)/(2\lambda)\big)$. First, we compare the ex ante payoff $W_1(\gamma_1)$ under the mechanism given in case (b) of Main Result with W_1 under the one-round delay mechanism with maximum delay derived above. Under the mechanism given in case (b) of Main Result, the expected payoff is given by:

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

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 $\underline{\gamma}\in\left(\gamma_*,g^{-1}(1/2)\right)$ such that the one-round delay mechanism with maximum delay dominates the mechanism given in case (b) of Main Result 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

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

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

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

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 6. Fix an infinite mechanism $\{\delta_t\}_{t=1}^{\infty}$ and suppose that it is not effectively finite. This is equivalent to assuming that there exists an equilibrium in which the uninformed player persists with strictly positive probability in every round (i.e., $x_t > 0$ for all t).

First, we claim that there is no equilibrium in which the probability that the informed player persists in round t, y_t, equals 0 for some t. To see this, note that in any such equilibrium $x_t > 0$. From the optimality of the uninformed player persisting in round t we have that

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

where U_{t+1} is the uninformed continuation equilibrium payoff. Because the informed can always persist and then mimic the equilibrium behavior of the uninformed, from the optimality of the informed conceding we have

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

Combining the two inequalities, we have

$$x_t(1 + \delta_t - \lambda - U_{t+1}) \geq 2\lambda,$$

which is not possible because by assumption $\delta_t \leq \Delta < \lambda$ and U_{t+1} is bounded from below by $1 - 2\lambda$.

Second, we claim that in any equilibrium in which both x_t and y_t are strictly positive for all t, the equilibrium payoff V_t to the informed is at least as large as the equilibrium payoff U_t to the uninformed for all t. To see this, note that by assumption both U_t and V_t are equal to the expected payoff from the strategy of persisting in round t and in each succeeding round. For the uninformed this gives

$$U_t = \gamma_t \left(1 - \mathbb{E}\left[\sum_{s=t}^{N-1} \delta_s \middle| \{x_s\}_{s\geq t}\right]\right) + (1-\gamma_t)\left(1 - 2\lambda - \mathbb{E}\left[\sum_{s=t}^{N-1} \delta_s \middle| \{y_s\}_{s\geq t}\right]\right),$$

where N is the terminal round when the game ends, which has an unbounded support, and the expectation is taken with respect to the distribution of N when faced with an opponent with a continuation strategy $\{x_s\}_{s\geq t}$ and $\{y_s\}_{s\geq t}$ respectively. Note that both expectations exist by the assumption that $(\{x_s\}, \{y_s\})_{s\geq t}$ is an equilibrium strategy profile. From the strategy of the informed player to persist in each round starting in round t, we have

$$V_t = 1 - \mathbb{E}\left[\sum_{s=t}^{N-1} \delta_s \middle| \{x_s\}_{s \geq t}\right].$$

By conceding at time t the uninformed can always obtain a payoff at least as large as $1 - 2\lambda$, so optimality of the equilibrium strategy requires that $U_t \geq 1 - 2\lambda$. Since the second term in the expression for U_t above is no larger than $1 - 2\lambda$ while the first term is a fraction of V_t, we have $V_t \geq U_t$.

Third, we claim that $y_t = 1$ for all t. To prove the claim we show that whenever the uninformed weakly prefers persisting to conceding, the informed strictly prefers persisting. Let γ_t be the belief of the uninformed player. For the uninformed player to weakly prefers persisting to conceding, we must have

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

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 the previous claim we know that $U_{t+1} \leq V_{t+1}$, hence

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

implying that the informed player strictly prefers persisting to conceding in round t regardless of x_t.

Finally, 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 and always persisting is optimal for the uninformed at any t. If the limit of γ_t is strictly positive, from Bayes' rule we have $\lim_{n\to\infty} \prod_{t=\tau}^{\tau+n} x_t > 0$, which implies that $\lim_{n\to\infty} \prod_{t=\tau'}^{\tau+n} x_t$ can be made arbitrarily close to 1 by taking $\tau' > \tau$ and sufficiently large. However, in round τ', always persisting results in no alternative being implemented with probability close to 1 and yields a payoff to the informed strictly lower that the payoff from implementing any alternative, contradicting the equilibrium condition. If γ_t converges to zero instead, then for t large enough the expected payoff to the uninformed from conceding is close to the first best payoff of 1 while the strategy of persisting from t onward leads to no alternative being implemented with probability close to 1, again contradicting the equilibrium condition. ■

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 6 can only have equilibria where the game ends with probability 1 in some round N. Since the the game ends with probability 1 in round N, in any such equilibrium the uninformed player concedes, with $x_N = 0$, and because the game does not end with probability 1 before N, $x_t > 0$ for each $t < N$. Given this, in such equilibrium the informed player persists, with $y_N = 1$. It follows that in this case the payoffs 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 know from the proof of Lemma 6 that if the uninformed player weakly prefers persisting to conceding in round t then the informed player strictly prefers persisting to conceding. Since $x_t > 0$ and $y_t = 1$, 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, F. (2000). "Bargaining and Reputation." *Econometrica* 68, 85–117.

Adler, R., Rosen, B. and Silverstein, E. (1998). "Emotions in Negotiations: How to Manage Anger and Fear." *Negotiation Journal*, 14(2), 161–179.

Bester, H., and Strausz, R. (2001). "Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case." *Econometrica* 69(4), 1077–1098.

Cramton, P. (1992). "Strategic Delay in Bargaining with Two-Sided Uncertainty." *Review of Economic Studies* 59(1), 205–225.

Damiano, E., Li, H., and Suen, W. (2012). "Optimal Deadline for Agreements." *Theoretical Economics*, 7, 353–393.

Deneckere, R., and Liang, M. (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, T., and Pesendorfer, W. (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, B. (1984). "On the Theory of Delegation," in *Bayesian Models of Economic Theory*, edited by Boyer and Kihlstrom. North-Holland, New York.

Kolotilin, A., Li, H., and Li, W. (2011). "Optimal Limited Authority for Principal." forthcoming *Journal of Economic Theory*.

Li, H., and Suen, W. (2009). "Decision-making in Committees." *Canadian Journal of Economics* 42(2), 359–392.

Li, H., Rosen, S., and Suen, W. (2001). "Conflicts and Common Interests in Committees." *American Economic Review* 91(5), 1478–1497.

Moore, J., and Repullo, R. (1990). "Nash Implementation: A Full Characterization." *Econometrica* 58, 1083–1099.

Myerson, R., and Satterthwaite, M. (1983). “Efficient Mechanisms for Bilateral Trading.” *Journal of Economic Theory* 29(2), 265–281.

Skreta, V. (2006). “Sequentially Optimal Mechanisms.” *Review of Economic Studies* 73(4), 1085–1111.