Optimal Delay in Committees

Ettore Damiano   Li, Hao   Wing Suen

July 10, 2018—Mechanism Design Workshop

Introduction

- We study collective decision problems (with no transfers) in which disagreements can be either preference-driven or information-driven
- Examples: legislative bargaining, trade negotiations, adoption of industry standards, recruitment committee, workplace practices
- Information aggregation is a key aspect of the model.

Limited Commitment

- Large enough delay (punishment) induces people to give up preference-driven disagreement to achieve first-best, without actually incurring the delay cost, but:
  - a “mistake” made by one player can produce a very bad outcome for all
  - requires commitment power because imposing lengthy delay is costly ex post
- We consider a dynamic mechanism design problem in which:
  - there is an upper bound to the length of delay in each round
  - players commit to a sequence of delays subject to this bound

Research Questions

- Does dynamic mechanism dominate static mechanism when there is an upper bound on delay?
- Does punishment work better when it is front-loaded, back-loaded, or constant through time?
- Is it optimal to have binding deadlines?
- Does the optimal mechanism always produce the efficient decision?

Model

- Two players (*A* and *B*); each of whom has a “favorite” alternative (*a* and *b*, resp.)
- Each player can be high type (*H*) or low type (*L*); the types of the two players are *not* independent
- $\gamma_1 =$ low type’s belief that opponent is *low* type
- $\mu_1 =$ high type’s belief that opponent is *low* type
- Assumption 1: $\gamma_1 < \mu_1$ (negative correlation)

Payoffs

- each player prefers opponent's favorite alternative when he is low type and opponent is high type
  - otherwise prefers his own favorite
- when opponent is low type; payoff gain from choosing own favorite (relative to choosing opponent's favorite) is larger for high type than for low type
- Example: payoff to player I from alternative j is $\theta_j + \mathbf{1}(i = j)\pi$
- Assumption 1 and the payoff assumptions ensure that high type expects to gains more (than a low type does) from an increase in probability that the opponent low type concedes

Impossibility of Information Aggregation

- First-best is to choose a player's favorite if he is high type and opponent is low type; otherwise flip a coin
- If $\gamma_1 \leq \gamma_*$, first-best can be implemented via a voting game (high type always votes for his favorite; low type always concedes)
- If $\gamma_1 > \gamma_*$, no mechanism without transfers can achieve first-best
  - think of γ_1 as the degree of conflict within the group

One Round Delay Mechanism

- if both players choose their favorite alternatives, impose a delay cost δ_1 before the decision is made by flipping a coin
- $x_1 =$ probability low type votes for own favorite
- second-best mechanism: choose lowest δ_1 such that low type concedes ($x_1 = 0$):
  - achieve first-best decision
  - incur some delay cost when two high types meet
- if there is an upper bound Δ on the delay cost, then second-best is not achievable when $\gamma_1 > \gamma^*$

How Does Repeated Voting Help?

- One round: choose $\delta_1 = \Delta$, induce low type to choose own favorite with probability $x_1 < 1$.
- Two rounds:
  - second round: choose some $\delta_2 \leq \Delta$ such that $x_2 < 1$ and continuation payoff for low type is the same as coin flip
  - this is feasible because information revealed in first round reduces conflict
  - first round: choose $\delta_1 = \Delta$ to induce the same x_1
  - low type is indifferent between one-round mechanism and two-round mechanism but high type prefers the latter

Repeated Voting

- Each player votes for *a* or *b* simultaneously at each round *t*.
  - If the votes agree, that decision is implemented immediately and the game ends.
  - If both players concede, then flip a coin to decide immediately.
  - If both persist (vote his own favorite), then each player incurs a delay cost of $\delta_t \leq \Delta$ and votes again in the next round.
- The game can in principal go on indefinitely.
- If the game is finite with *T* rounds, then flip a coin at the very end.

Key results

- Any optimal delay mechanism is finite with a binding deadline T.
- The terminal belief γ_T on entering the last round T is less than or equal to γ^*.
  - efficient decision is always achieved
  - if Δ is not too large and γ_1 is not too close to γ^*, then terminal belief is exactly equal to γ^*
- Stop-and-start: equilibrium play alternates between some concession ($x_t < 1$) and no concession ($x_{t+1} = 1$)
  - optimal delay sequence alternates between $\delta_t = \Delta$ and $\delta_{t+1} < \Delta$.

Remarks about the Design Problem

- Screening Lemma says that $x_t > 0$ implies $y_t = 1$
  - can focus on equilibria in which high type always votes for own favorite
  - equilibrium play depends only on U_t but welfare analysis depends also on V_t.
- The problem is difficult to study because beliefs are solved forwards while payoffs are solved backwards, and the length of the horizon is not fixed:
  - introduce *localized variations method*

Finite Deadline

- An *active round* is one in which $x_t < 1$.
- Proposition 1: Any optimal delay mechanism has a finite number of active rounds. Moreover, $x_t > 0$ for all t before the deadline T.
- Idea of proof:
  - Belief goes down in each active round. If it converges to a positive limit, $\lim_{n\to\infty} \prod_{t=\tau}^{\tau+n} x_t$ is arbitrarily close to 1 for τ large. But then persisting is bad for high type in round τ.
  - If $x_N = 0$ for some $N < T$, the high types are playing a pure war of attrition after round N. We can show that truncating the game after round N and replacing it with a coin toss is better.

Efficient Deadline Concession

- Proposition 2: Any optimal delay mechanism with at least two rounds has efficient deadline concession (i.e., $x_T = 0$ and $y_T = 1$ if $T \geq 2$).
  - means $\gamma_T \leq \gamma^*$
  - proof uses a localized variations method

Localized Variations

- Suppose $\gamma_T > \gamma^*$. Consider a way to marginally drive down γ_T.
- Let (n) be the last active round prior to T. Insert another round round s after (n) but before T with appropriately chosen $\delta_s > 0$ to induces $x_s < 1$
- But a lower x_s means a higher continuation value for the low type after round (n), which would change the entire sequence of play.
- Neutralize this effect by inserting yet another round s' between (n) and s and choose $\delta_{s'} > 0$ equal to the increase in continuation payoff above
- This variation leaves the sequence of play the same up to round (n) and therefore has no effect on U_1, but it induces more concession from the low type, which improves V_1.

Maximal Concession

- Persisting is a worse option is $-\delta_t + U_{t+1}$ is low
- δ_t is bounded above by Δ
- U_{t+1} is bounded below by the payoff from immediately conceding in round $t + 1$
  - the latter payoff is lowest when $x_{t+1} = 1$
- there is *maximal concession* by the low type in round t when $\delta_t = \Delta$ and $x_{t+1} = 1$

Front Loading

- Proposition 3: Any optimal delay mechanism with at least two rounds induces the maximal concession in the first round.
- If concession is not maximal, we employ the following localized variation:
  - maximize concessions (i.e., lower x_1) by increasing the delay penalty (by inserting extra rounds after round 1 if necessary)
  - raise $x_{(2)}$ by lowering delay penalty in the next active round (2) in such a way to keep $x_1 x_{(2)}$ (and therefore $\gamma_{(3)}$) unchanged
  - Because the continuation play starting from round (3) remains the same, we can use a direct computation of these two changes to show that the gain from a lower x_1 is larger than the loss from a higher $x_{(2)}$

Efficient Deadline Belief

- Proposition 4: In any optimal mechanism with more than two active rounds, $\gamma_T = \gamma_*$.
- If $\gamma_T < \gamma_*$, we consider the following localized variation:
  - reduce delay cost $\delta_{(n)}$ prior to round T to drive the belief up to γ_*.
  - a higher $x_{(n)}$ lowers the payoff to the uninformed; we neutralize this by reducing the delay $\delta_{(n-1)}$ to keep $U_{(n-1)}$ unchanged
  - we show that this variation raises $V_{(n-1)}$

Slack

- An active round (i) has *no slack* if $x_{(i)}$ is equal to the maximal concession; it has *slack* if $x_{(i)}$ is less than the maximal concession.
- A mechanism with slack in two successive active rounds is not optimal.

Corner Solution

- Suppose there is slack in both round (i) and round $(i+1)$. The localized variation involves:
  - change $\delta_{(i)}$ to change $x_{(i)}$ marginally (up or down)
  - change $\delta_{(i+1)}$ to change $x_{(i+1)}$ in such a way to keep $x_{(i)}x_{(i+1)}$ constant. This guarantees that $\gamma_{(i+2)}$ and hence the subsequent equilibrium play is unaffected.
  - change $\delta_{(i-1)}$ in such a way to keep the continuation value at round $(i-1)$ constant. This guarantees that the equilibrium play prior to and including round $(i-1)$ is unaffected.
- Equivalent to choosing $\gamma_{(i+1)}$ to maximize total delay, while holding $\gamma_{(i)}$ and $\gamma_{(i+2)}$ constant.
- This maximization problem is convex in $\gamma_{t(i+1)}$.
- Corner solution means no slack in either round (i) or round $(i+1)$.
- Which corner to choose is payoff-equivalent.

Stop-and-Start

- Proposition 5: In any optimal mechanism with at least two active rounds, there can be at most one active round with slack.
- Proof: If there are active two rounds (i) and (j) with slack, we can use the payoff equivalent result to reshuffle these two rounds to make them adjacent. But then it cannot be optimal.
- No slack at round (i) requires $x = 1$ in the next round. Hence it is optimal to have maximal concession followed by no concession.

Optimal Mechanism

- Initial belief is γ_1 and terminal belief is γ_*.
- Belief evolves according to $\gamma_{(i+1)} = g(\gamma_{(i)})$ in each active round where $g(\cdot)$ is given by Bayes' rule under the maximal concession $x_(i)$.
- The concession in the last active round $x_{(n)}$ is chosen in such a way that $\gamma_{(n)}$ updates to γ_*.
- This pins down the entire evolution of beliefs.
- Use definition of no slack to figure out the implied sequence $\{\delta\}_{t=1}^{T}$.

Optimal Mechanism

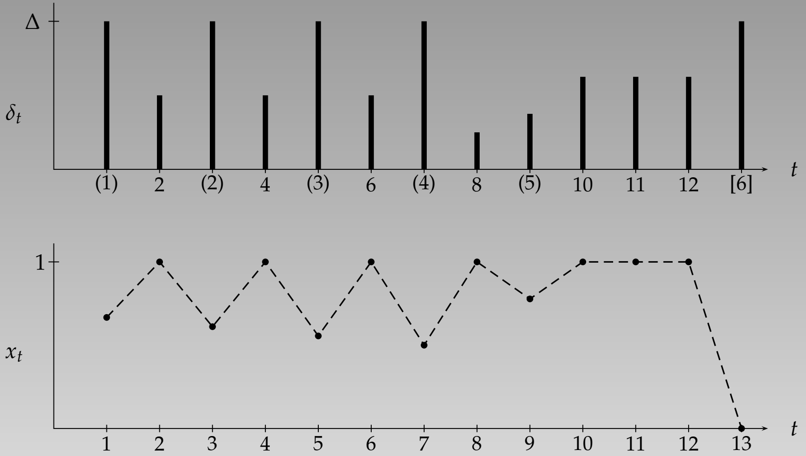

Optimal Mechanism

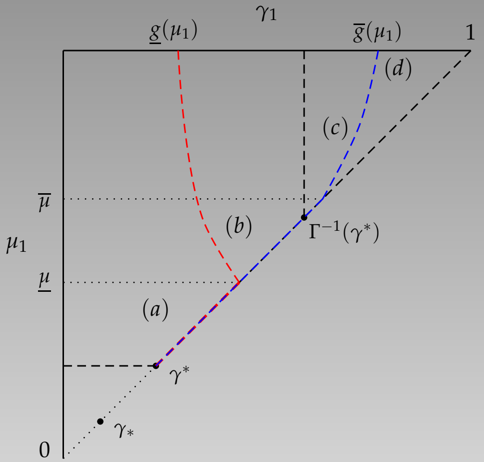

Optimal Mechanism

(a) If γ_1 is very close to γ_*, then one-round mechanism with $\delta_1 = \Delta$ is optimal.

(b) For larger γ_1, there is no slack in round 1 and the game immediately enters the deadline phase. Terminal belief is $\gamma_T \in (\gamma_*, g(\gamma_1))$.

(c) For still larger γ_1, there are both stop-and-start phase and deadline phase. Terminal belief is γ_*.

(d) For γ_1 close to 1, optimal mechanism is to flip a coin.

Discussion

- Continuous-delay limit as Δ goes to 0 is the same as in companion paper, but the optimal mechanism with non-constant delay does strictly better in any discrete time mechanism.
- Can be extended to more general payoff structures. The key component is that the informed type benefits more from concession by the uninformed type than the uninformed type does.
- May do even better if there is delay in implementing agreed decision, but then this also requires commitment power.

Thank you!