Deciding Before You Are Persuaded

Torun Dewan* Kfir Eliaz†

Abstract

We study how a decision maker can use a committed deadline to discipline a biased sender's information acquisition. Evidence arrives through public tests, and both investigation and waiting are costless. With sufficient time, the sender induces favourable decisions by eliminating unfavourable possibilities. A shorter deadline can instead make direct evidence the sender's only route to persuasion. We characterize the optimal deadline and show when it achieves perfect accuracy. For general discovery technologies, the gain compares undiscovered guilt with wrongly convicted innocence. Under continuous posterior paths and feasible stopping, no deadline improves accuracy.

Keywords: Bayesian persuasion; information acquisition; commitment; decision deadlines.
JEL classification: D82; D83; C72.

1 Introduction

It takes time to collect and present information intended to persuade a decision maker. A prosecutor checks alibis and interviews witnesses; campaign managers and the media investigate political candidates. Although it might seem prudent to let such investigations run their course, we show that a decision maker can benefit from limiting their duration, even when collecting evidence is costless. Given enough time, a Bayesian-rational decision maker may be manipulated into choosing the persuader's preferred action more often than she would if the underlying truth were known. Committing in advance to a decision deadline provides a defence against such gradual manipulation. An appropriately chosen deadline can steer the persuader towards an investigation path that is less biased and hence more beneficial to the decision maker.

We study a receiver who commits to a decision date before a biased sender designs an investigation. The sender chooses which public tests to run but cannot conceal their results. By changing which investigation the sender finds worthwhile, a deadline can induce more accurate decisions than allowing additional time.

*London School of Economics and Political Science, Department of Government.
†Department of Economics, Tel Aviv University, and King's College London.

Our example is a prosecutor who wishes to convict and a judge who wants the correct verdict. The judge convicts when guilt is at least as likely as innocence. Each period, the prosecutor can publicly verify whether a particular state is true. Some states establish guilt and the others innocence. The judge fixes a deadline, the prosecutor commits to an investigation for that horizon, and the judge acts on the resulting evidence. Both players are rational and Bayesian; testing and waiting are costless. The equal weighting of the two verdict errors gives a simple decision threshold rather than a literal description of a criminal standard of proof.

With a long deadline, the prosecutor can obtain convictions by eliminating innocent explanations. He tests enough innocent states that, if all tests are negative, the remaining possibilities are equally divided between guilt and innocence. The judge then convicts on the unresolved evidence. The investigation attains the sender's one-shot persuasion value, while leaving the judge's accuracy at the level she could obtain by acquitting on the prior. A sufficiently short deadline prevents the construction of this conviction pool. The prosecutor can then persuade only by identifying a genuinely guilty state, so he directs his tests towards guilt.

The first result characterizes the optimal deadline in a finite-state model. A deadline strictly improves accuracy when the conviction pool takes more than one period to construct. It achieves perfect accuracy when there is enough time to test every guilty state before elimination becomes feasible, which occurs precisely when the prior probability of guilt is below one third. Otherwise, the best short investigation may leave some guilt undiscovered. The receiver chooses time for useful discovery while preventing the sender's preferred elimination strategy.

The second result extends this trade-off to general discovery technologies. An early deadline helps if and only if the guilty probability mass left undiscovered by the best short investigation is smaller than the innocent mass convicted under a long deadline. The third result identifies a contrasting case. If posterior beliefs evolve continuously and the sender can stop information production at a threshold and hold beliefs fixed, every sender-optimal policy gives the receiver the same accuracy as acting on her prior. Gradual learning alone is therefore insufficient for a deadline to help.

The gain arises through the sender's choice of investigation. For a fixed experiment, additional costless evidence weakly improves a Bayesian receiver's expected payoff. A deadline fixed after the sender has chosen that experiment cannot generate our gain. The receiver must commit before the choice she seeks to influence. She commits to a date, while retaining the optimal verdict given the evidence available on that date.

Adjudication and collective decisions provide possible applications. A procedural deadline announced before an interested party investigates, or a common voting date anticipated before campaigners collect evidence, can constrain information production.[1] The application depends on the timing rule actually binding: a date that is routinely extended need not provide the commitment assumed here. An individual voter's early ballot is a different case, since it need not change an investigation directed at a larger electorate. These examples concern the choice of information, rather than a general argument for earlier decisions.[2]

2 Related literature

Our starting point is Kamenica and Gentzkow (2011). Their committed sender chooses an experiment that pools guilty states with enough innocent ones to reach the conviction threshold. We restrict implementation to public tests that take time and give the receiver a prior choice of deadline. The single-state tests follow the learning-by-identification technology of Frug (2018), with public rather than private inspections and outcomes. This separates the sender's control of the questions asked from control of the answers disclosed.

The closest timing comparisons help isolate the receiver's instrument. Brocas and Carrillo (2007) study a leader who controls how much public evidence is collected before a follower acts. Liao (2021) allows the receiver to choose when to stop costly sampling from a sender-designed experiment. Here the receiver instead commits to a deadline before the investigation is chosen, and learning is costless. The resulting gain comes from changing information production rather than saving on investigation or waiting costs.

Dynamic persuasion studies the provision of information over time (Au, 2015; Ely, 2017; Honryo, 2018; Renault et al., 2017). Within this literature, Senkov (2022) studies an interim deadline chosen by the sender, while Orlov et al. (2020) study persuasion to delay an irreversible decision. Che et al. (2023) analyse costly dynamic persuasion without

[1] Under the 2021 ICC Arbitration Rules, the tribunal establishes a procedural timetable under Article 24(2). Article 31 normally provides a six-month time limit for the final award, while authorizing the ICC Court to set a different limit or grant an extension. At the United States International Trade Commission the administrative law judge sets a target completion date early in a Section 337 investigation. Neither is irrevocable, since timetables can be extended. In Thailand registering to vote early gives up the ordinary polling-day ballot, and in Denmark an advance ballot removes the voter's later polling-station override, though a replacement remains possible. So while none provides the precise unconditional commitment we model here, they provide relevant context to our arguments about timing. (Sources: the 2021 ICC Arbitration Rules, Art. 31; 19 CFR 210.51; the Thai Government Public Relations Department notices of December 2025–February 2026; and the Danish Folketing Elections Act, Consolidated Act No. 4 of 2024, sections 53–67.)

[2] At the time of writing, controversy over Vivek Ramaswamy's advertisement featuring Amy Acton's estranged mother coincided with substantial early voting in Ohio (Associated Press, 21 October 2026). The episode illustrates how campaign messages arrive after some voters have decided. Our mechanism concerns the further question of how anticipated decision dates affect the investigations campaigns undertake.

commitment by either player. We retain sender commitment and isolate the effect of allowing the receiver to commit first to a decision date. Receiver commitment also appears in Glazer and Rubinstein (2012), through conditions for persuasion in a model with boundedly rational agents. Our receiver's instrument restricts the duration of a public investigation and leaves the final action posterior-optimal.

3 The model

3.1 States, preferences and tests

There are two players, a sender S and a receiver R. The state ω is uniformly distributed on $\Omega = \{1, \ldots, N\}$. A publicly known subset G contains the guilty states, with $|G| = g$ and $1 \leq g < N/2$. Its complement contains the innocent states. The prior probability of guilt is $p = g/N < 1/2$. Neither player knows the realized state; knowing which labels belong to G does not reveal guilt.

The receiver chooses conviction C or acquittal A. Her payoff is one for a correct verdict and zero otherwise; the sender's payoff is one for conviction and zero for acquittal. At posterior guilt probability q, the receiver therefore chooses

$$a(q) = \begin{cases} C, & q \geq 1/2, \\ A, & q < 1/2. \end{cases} \tag{1}$$

Conviction at equality is an explicit tie-breaking convention. The sender benefits from either a correct or an incorrect conviction.

For each state i, the available test is $X_i(\omega) = \mathbf{1}\{\omega = i\}$. A positive result identifies the state; a negative result excludes it. At most one test runs per period, and its identity and outcome are public. The sender cannot test a union of states, suppress an outcome or garble the record. He may stop testing, idle, and condition future tests on public outcomes and state-independent public randomization. Repeating a known test adds no information. The restriction concerns the questions that can be asked, not the number of binary answers an unrestricted experiment could produce.

3.2 Timing and solution

First, the receiver commits to a finite integer deadline $T \geq 0$. Second, before learning anything about ω, the sender commits to an entire contingent testing plan for that horizon. Third, the plan runs for at most T periods. Finally, the receiver updates by Bayes' rule and

acts according to (1). If a test establishes guilt or innocence earlier, an immediate verdict is payoff-equivalent. There are no investigation costs, waiting costs or discounting.

A sender strategy specifies a test or idling at each public history before the deadline. A receiver strategy specifies a deadline and the posterior-optimal final verdict. We solve by backward induction: find the sender's optimal plans for each deadline and then the receiver's best deadline. The results below show that all sender-best responses at a given deadline give the receiver the same accuracy, so her choice requires no selection among equally persuasive investigations. The commitment concerns when she decides, not a promise to take an action that would be suboptimal on her evidence.

4 Deadlines and investigation

4.1 Elimination and the sender's choice

A positive test settles the verdict. On an all-negative history, let r guilty and s innocent states have been ruled out. The untested states remain equally likely, giving

$$q(r,s) = \frac{g-r}{N-r-s}, \qquad N-r-s>0. \tag{2}$$

Histories eliminating every state have probability zero. Write $d = N - 2g > 0$ for the initial excess of innocent over guilty states. The receiver convicts on an all-negative history if and only if

$$s - r \geq d. \tag{3}$$

Thus conviction without a guilty discovery requires excluding at least d more innocent than guilty states. Testing a guilty state can establish guilt, but a negative answer moves the unresolved pool away from conviction. Testing an innocent state has the opposite effect on that pool.

A deterministic plan is summarized by its distinct tests on the all-negative path, with $0 \leq r \leq g$, $0 \leq s \leq N-g$, and $r+s \leq T$. Its conviction probability is

$$C(r,s) = \frac{r}{N} + \frac{N-r-s}{N}\mathbf{1}\{s-r \geq d\}. \tag{4}$$

The first term counts tested guilty states. The second counts the unresolved pool when it convicts. A positive innocent test leads to acquittal. Before stating the result, consider the two ways an investigation can persuade. A guilty-state test may establish guilt immediately. An innocent-state test cannot do so, but a negative answer raises the posterior of guilt in the

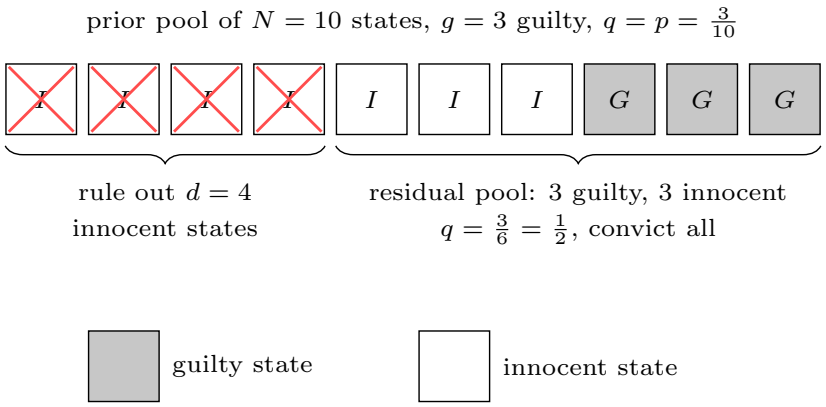


Figure 1: The conviction pool for $N = 10$ and $g = 3$. Four negative innocent-state tests leave three guilty and three innocent possibilities. The posterior reaches $1/2$, and the receiver convicts on this unresolved history. Half of those convicted are innocent.

unresolved pool. Reaching the conviction threshold by elimination requires at least d periods. The deadline determines whether that second route is available.

Write $V_S(T)$ for the sender's maximal conviction probability and $V_R(T)$ for receiver accuracy under a sender-best response.

Proposition 1 (Investigation and the optimal deadline)**.** *Let $d = N - 2g$ and $m(T) = \min\{T, g\}$.*

1. *For every finite integer deadline $T \geq 0$,*

$$V_S(T) = \begin{cases} m(T)/N, & T < d, \\ 2g/N, & T \geq d, \end{cases} \qquad V_R(T) = \begin{cases} 1 - g/N + m(T)/N, & T < d, \\ 1 - g/N, & T \geq d. \end{cases}$$

   An optimal plan tests $m(T)$ guilty states when $T < d$, and d innocent states followed by stopping when $T \geq d$. Every sender-optimal plan at a given deadline gives the stated receiver accuracy.

2. *A shortest optimal deadline is $T^* = \min\{g, d - 1\}$. If $d \geq 2$, every optimal deadline is shorter than d and strictly improves on every deadline $T \geq d$. If $d = 1$, every deadline gives the same accuracy.*

3. *An optimal calendar deadline gives correct verdicts in every state if and only if $3g < N$, equivalently $p < 1/3$. A shortest such deadline is $T^* = g$.*

For $T < d$, conviction requires a genuine guilt discovery. There are no false convictions, and accuracy is the probability of innocence plus the guilty mass identified. For $T \geq d$, the sender can leave g guilty and g innocent states unresolved and obtain conviction on all of them. This achieves the one-shot persuasion value $2p$. A longer horizon allows the sender to idle after reaching it; it does not require further revelation. Accuracy is $1 - p$, the same as acquitting on the prior. Sender indifference over optimal plans does not change these totals.

The receiver therefore gains from allowing useful investigation while keeping elimination infeasible. When $g < d$, every integer deadline from g through $d-1$ is optimal: the prosecutor can test all guilty states but cannot construct the conviction pool. When $g \geq d \geq 2$, the sole optimum is $d - 1$. The gain relative to a long deadline is

$$\Delta(g) = \frac{\min\{g, N - 2g - 1\}}{N}.$$

With few guilty states there is little guilt to discover. With many, the sender can reach the threshold quickly by eliminating innocence. For fixed N, the largest gains occur at priors near one third; this comparison changes the prior and the elimination threshold together, rather than isolating an investigation-speed effect. The exact integer maximum is recorded in Appendix A.1 (supporting proofs).

Perfect accuracy requires time to test every guilty state before elimination can reach the threshold. At $p = 1/3$, the two tasks take equally long. A calendar deadline can no longer permit the first while preventing the second.

Example 1. Let $N = 10$ and $g = 3$. A deadline of at least four periods lets the sender exclude four innocent states. The six unresolved states contain three guilty and three innocent possibilities, so the judge convicts all six (Figure 1). Conviction probability is 0.6 and accuracy 0.7. With three periods, the sender instead tests the three guilty states: conviction probability falls to 0.3 and accuracy rises to one.

Information and commitment. At posterior q, the receiver's indirect payoff is $\max\{q, 1-q\}$, a convex function. For a fixed plan, public beliefs form a martingale, so additional evidence weakly improves expected accuracy. In Example 1, truncating the innocent-state investigation after three tests induces acquittal everywhere and leaves accuracy at 0.7. Announcing that deadline before the investigation is chosen instead induces guilty-state tests and raises accuracy to one. The equilibrium comparison changes the experiment, rather than finding a benefit from discarding free evidence generated by a fixed experiment.

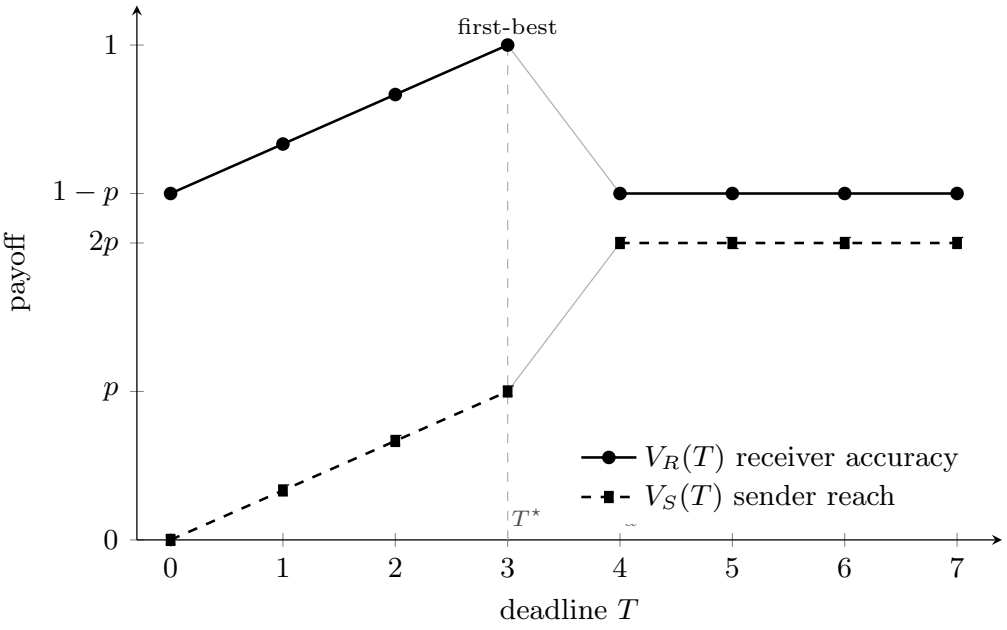


Figure 2: Receiver accuracy and sender conviction probability for $N = 10$, $g = 3$. Accuracy rises to one at $T^* = 3$, then falls to $1 - p$ when elimination becomes feasible at $d = 4$. The sender's value jumps to $2p$ at that threshold.

Calendar deadlines and contingent stopping. The instrument in Proposition 1 is a date fixed before investigation begins. A rule that conditions on the proposed test can do more. Suppose each test is publicly announced before its result, and the receiver can commit to decide immediately before the first proposed innocent-state test or after g periods, whichever comes first. Add a public announcement of each proposed test before its result. The receiver commits to decide immediately before the first proposed innocent-state test, or after g periods, whichever is earlier.

On any unresolved history before an innocent test, only guilty states have been excluded, so the posterior remains below $1/2$ and the receiver acquits. Consequently no innocent state can be convicted. Testing all g guilty states attains conviction probability p, the largest possible without false convictions. Every best response must therefore convict all guilt and no innocence, giving perfect accuracy. A calendar date requires no intervention between a proposed test and its result; the richer rule uses that additional opportunity.

5 General discovery technologies

The same commitment order applies to two channels. A guilt test may conclusively establish guilt and never falsely does so; an innocence test may conclusively establish innocence and

never falsely does so. Each test has only that conclusive outcome or a single inconclusive outcome. One channel is tested each period, or the sender idles.

Let A_k be the conditional probability that k guilt tests fail to discover guilt, given guilt, and B_ℓ the corresponding probability for ℓ innocence tests, given innocence. Assume

$$A_0 = B_0 = 1, \qquad 0 \leq A_{k+1} \leq A_k, \qquad 0 \leq B_{\ell+1} \leq B_\ell.$$

Conditional on guilt, no-discovery probability depends only on the number of guilt tests, irrespective of the placement or number of innocence tests; the analogous restriction holds conditional on innocence. Idling changes neither probability. This count-based technology permits increasing or diminishing returns, flat stretches and certain finite detection, but excludes additional inconclusive signals carrying information.

The unresolved guilty and innocent masses are $x_k = pA_k$ and $y_\ell = (1-p)B_\ell$. Their posterior is $x_k/(x_k + y_\ell)$ when the denominator is positive. Define

$$K = \min\{\ell \geq 1 : (1-p)B_\ell \leq p\}, \quad Y = (1-p)B_K, \quad X = pA_{K-1},$$

and assume K is finite. Innocence testing alone first reaches the conviction threshold after K periods. It leaves innocent mass Y in the convicted pool. Guilt testing for one period less leaves guilty mass X undiscovered.

Proposition 2 (General technology)**.** *For $T < K$, sender conviction probability is $p(1 - A_T)$ and receiver accuracy is $1 - pA_T$. For $T \geq K$, they are $p + Y$ and $1 - Y$. The accuracy formulas hold under every sender-best response. The receiver's best committed accuracy is*

$$V_R^* = 1 - \min\{X, Y\}.$$

An early deadline strictly improves on a long deadline if and only if $X < Y$.

The criterion is an error comparison. At the best short deadline, $K - 1$, every innocent defendant is acquitted, but guilty mass X remains undiscovered and is also acquitted. All errors are false acquittals. With a long deadline, all guilty defendants are convicted, together with innocent mass Y; all errors are false convictions. The receiver values correct decisions in either state, so the shorter deadline helps precisely when the errors it creates are smaller than the errors it avoids. When $X = Y$ the receiver is indifferent. When $X > Y$, the loss of useful discovery outweighs the protection against false convictions.

This formulation separates the speed of conclusive guilt discovery from the time needed to reach the threshold through innocence testing. A technology can permit both channels

and yet offer no benefit from an early deadline. What matters is their relative effectiveness at the relevant horizon, not simply the availability of direct evidence.

The baseline has $A_k = \max\{1 - k/g, 0\}$ and $B_\ell = \max\{1 - \ell/(N-g), 0\}$. Then $K = d$ and $Y = p$, recovering the earlier results. With constant discovery probabilities $\alpha, \beta \in (0,1)$, instead,

$$A_k = (1-\alpha)^k, \quad B_\ell = (1-\beta)^\ell, \quad K = \left\lceil \frac{\log(p/(1-p))}{\log(1-\beta)} \right\rceil,$$

and the gain is $\max\{(1-p)(1-\beta)^K - p(1-\alpha)^{K-1}, 0\}$. Better guilt detection weakly increases the gain for fixed β. Better innocence detection weakly reduces it while K is unchanged, but changes in K can generate discontinuities and nonmonotonicity.

5.1 Continuous posterior paths

Continuous learning supplies a useful contrast to conclusive discovery. The sender's ability to stop exactly at the receiver's threshold removes the accuracy gain in this case.

Proposition 3 (Continuous paths)**.** *Suppose every feasible policy makes the public posterior a bounded, continuous martingale starting at $p < 1/2$, and every such policy can be stopped at any stopping time and its posterior held fixed. At any finite deadline for which a sender-optimal policy exists, every sender-optimal policy gives receiver accuracy $1 - p$. No deadline improves on that benchmark.*

The sender can freeze a continuous path at its first arrival at one half. This preserves all eventual convictions and prevents later acquittals on paths that have reached that threshold. Optimality implies that freezing cannot increase conviction probability, so it preserves the original decision almost surely. A verdict based on a posterior of exactly one half adds no accuracy relative to acquittal. By contrast, a conclusive guilt discovery jumps directly to one. Such jumps leave scope for a deadline to induce useful evidence, with the precise gain governed by $X < Y$.

6 Conclusion

A receiver can improve decision accuracy by choosing a deadline before a biased sender designs an investigation. In the finite-state model, a sufficiently short horizon makes conclusive evidence of guilt the only route to conviction. A longer horizon permits persuasion through elimination. The general discovery model compares the guilt left undiscovered by an early decision with the innocence wrongly convicted after a longer investigation. Continuous

posterior paths, with feasible stopping and freezing, show why gradual learning by itself does not make deadlines valuable.

The institutional instrument is a committed calendar date. Its effectiveness depends on constraining the sender's investigation before that investigation is chosen. The receiver continues to act optimally on the evidence she observes; the gain comes from inducing a different supply of evidence. The model therefore identifies a role for decision timing alongside the rules governing which investigations and disclosures are feasible.

References

Au, Pak Hung. 2015. "Dynamic Information Disclosure." *RAND Journal of Economics* 46 (4): 791–823.

Brocas, Isabelle, and Juan D. Carrillo. 2007. "Influence Through Ignorance." *RAND Journal of Economics* 38 (4): 931–947.

Che, Yeon-Koo, Kyungmin Kim, and Konrad Mierendorff. 2023. "Keeping the Listener Engaged: A Dynamic Model of Bayesian Persuasion." *Journal of Political Economy* 131 (7): 1797–1844.

Ely, Jeffrey C. 2017. "Beeps." *American Economic Review* 107 (1): 31–53.

Frug, Alexander. 2018. "Strategic Gradual Learning and Information Transmission." *Journal of Economic Theory* 177: 594–615.

Glazer, Jacob, and Ariel Rubinstein. 2012. "A Model of Persuasion with Boundedly Rational Agents." *Journal of Political Economy* 120 (6): 1057–1082.

Honryo, Takakazu. 2018. "Dynamic Persuasion." *Journal of Economic Theory* 178: 36–58.

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

Liao, Xiaoye. 2021. "Bayesian Persuasion with Optimal Learning." *Journal of Mathematical Economics* 97: 102534.

Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov. 2020. "Persuading the Principal to Wait." *Journal of Political Economy* 128 (7): 2543–2578.

Renault, Jérôme, Eilon Solan, and Nicolas Vieille. 2017. "Optimal Dynamic Information Provision." *Games and Economic Behavior* 104: 329–349.

Senkov, Maxim. 2022. "Setting Interim Deadlines to Persuade." Working paper, CERGE-EI. https://arxiv.org/abs/2210.08294.

A Proofs

A.1 Proof of Proposition 1

Elimination threshold. At a positive-probability all-negative history the untested states remain equally likely. Thus $q = (g-r)/(N-r-s) \geq 1/2$ exactly when $s - r \geq N - 2g = d$. A history excluding every state has probability zero and needs no posterior assignment for the payoff calculations.

Sender value in part (i). After a positive test the verdict is fixed. Before it, a deterministic plan has a single all-negative path. Repeated tests and idle periods can be omitted, leaving counts (r, s) with $r + s \leq T$. Its conviction probability is the sum of the tested guilty mass r/N and the unresolved mass $(N - r - s)/N$ if that pool convicts.

For $T < d$, $s - r \leq s + r \leq T < d$, so the unresolved pool acquits. Hence conviction is at most $\min\{T, g\}/N$, achieved by that many guilty tests. For $T \geq d$, testing d innocent states and stopping yields $2g/N$. A plan whose unresolved pool acquits yields at most g/N. If the pool convicts, then $s \geq d + r$, and

$$C(r,s) = \frac{N-s}{N} \leq \frac{N-d-r}{N} = \frac{2g-r}{N} \leq \frac{2g}{N}.$$

The same bound covers an empty unresolved pool by its zero probability. Public randomization selects deterministic contingent plans independently of the state, so averaging cannot improve these bounds.

Receiver accuracy in part (i). Write TP and FP for true- and false-conviction probabilities. At $T < d$, every feasible plan has FP$= 0$, and sender optimality fixes TP$= \min\{T, g\}/N$. Accuracy is $1-p+$TP$-$FP, giving the short formula. At $T \geq d$, optimality gives TP$+$FP$= 2p$. Bayesian conviction implies FP$\leq$TP, while TP$\leq p$. Thus TP$=$FP$= p$ and accuracy is $1 - p$. This also covers all optimal randomized plans.

Optimal deadline and perfect accuracy: parts (ii) and (iii). For $T < d$, accuracy increases with $\min\{T, g\}$; for $T \geq d$, it is $1 - p$. The maximal short gain is $\min\{g, d-1\}/N$, attained first at $T^* = \min\{g, d-1\}$. It is positive iff $d \geq 2$. If $d = 1$, both branches have the

same accuracy. If $g < d$, all $T \in \{g, \ldots, d-1\}$ are optimal; otherwise, for $d \geq 2$, only $d-1$ is. Perfect accuracy requires $g \leq d-1$, equivalently $3g < N$.

The integer maximum of the gain. For the integer maximum, let $m = \min\{g, N-2g-1\}$. Then $g \geq m$ and $N \geq 2g + m + 1 \geq 3m + 1$, so $m \leq \lfloor (N-1)/3 \rfloor$. For $N \geq 4$, setting $g = \lfloor (N-1)/3 \rfloor$ attains this bound and satisfies $1 \leq g < N/2$. At $N = 3$, the only admissible $g = 1$ gives gain zero. The worked example substitutes $N = 10, g = 3$ into these formulas.

A.2 Proof of Proposition 2

A deterministic investigation has one unresolved path. Let its counts be (k, ℓ) with $k + \ell \leq T$. Count dependence implies unresolved masses $x_k = pA_k$, $y_\ell = (1-p)B_\ell$, and discovered guilty mass $p - x_k$. These masses also follow recursively from conditional discovery hazards $1 - A_{k+1}/A_k$ and $1 - B_{\ell+1}/B_\ell$ whenever the denominators are positive. A zero conditional survival probability leaves no corresponding mass to discover.

If the unresolved pool acquits, conviction probability is $p - x_k \leq p$. If it convicts, conviction probability is

$$(p - x_k) + (x_k + y_\ell) = p + y_\ell.$$

The formula remains valid for a zero-mass unresolved pool.

For $T < K$, $y_\ell > p \geq x_k$, so unresolved evidence cannot convict. Conviction is at most $p(1 - A_T)$, attained by T guilt tests. There are no false convictions; therefore every sender-best response gives accuracy $1 - pA_T$.

For $T \geq K$, innocence testing K times and stopping yields $p + Y$. To establish the upper bound for an arbitrary mixed-channel plan, unresolved conviction requires $y_\ell \leq x_k \leq p$, hence $\ell \geq K$. Monotonicity then gives $y_\ell \leq Y$, and conviction is at most $p + Y$. A plan ending in unresolved acquittal gives at most p, so cannot beat it. This step also covers plans that discover some guilt before reaching the unresolved threshold.

If $Y > 0$, every optimal deterministic plan must convict on its unresolved pool, hence convict all guilt and innocent mass exactly Y. Accuracy is $1 - Y$. If $Y = 0$, an optimal plan with unresolved acquittal must discover all guilt and makes no errors; an optimal plan with unresolved conviction likewise has no false convictions. An optimal public lottery assigns probability one to optimal deterministic plans, preserving these accuracy conclusions.

Short-deadline accuracy is nondecreasing in T and reaches $1 - X$ at $K - 1$. Long-deadline accuracy is $1 - Y$. Comparing them proves the optimal accuracy and the strict-gain criterion, including equality and $K = 1$. Substitution gives the linear-survival special case. For geometric survival, the logarithmic expression is the first integer satisfying $(1-p)(1-\beta)^K \leq p$.

Within a fixed-K region, increasing α reduces the undiscovered guilty mass and increasing β reduces the innocent convicted mass; taking the positive part preserves the stated weak signs. No global monotonicity in β is inferred across integer changes in K.

A.3 Proof of Proposition 3

Fix a finite deadline and a sender-optimal policy. On the same underlying state and signal realizations, follow that policy until its posterior first reaches $1/2$, then freeze. Stopping-time closure makes this alternative feasible. Let H be the event of reaching $1/2$ by the deadline and C the event of conviction under the original policy. Continuity and $p < 1/2$ imply $C \subseteq H$. The alternative convicts exactly on H, so optimality implies $\Pr(H \setminus C) = 0$. The two policies therefore give the same verdict almost surely, and the same expected accuracy.

Under the frozen policy, the terminal posterior never exceeds $1/2$. Where it equals $1/2$, conviction and acquittal have identical expected accuracy. Replacing those convictions by acquittals leaves expected accuracy unchanged, and always acquitting yields $1 - p$. The argument applies to every optimal policy. It assumes existence, Bayesian posterior martingales and feasible freezing; it is not a limit argument obtained by making the discrete tests faster.