Voluntary Disclosure of Evidence in Plea Bargaining

Pablo Cuellar*

November 2026

Abstract

I study how voluntary disclosure of evidence affects outcomes in plea bargaining. A prosecutor negotiates a sentence in continuous time with a defendant who privately knows whether he is guilty or innocent. While negotiating, the prosecutor privately gathers hard evidence about the defendant's type; if no agreement is reached by a deadline, a trial assigns payoffs based on the evidence. Voluntary disclosure induces endogenous second-order belief uncertainty. Under a consistency requirement on off-path beliefs, I show that even a purely sentence-motivated prosecutor discloses exculpatory evidence when the prior probability of guilt is high, and that voluntary disclosure then generates costly trials. Mandatory disclosure eliminates trials, raises the expected sentence of guilty defendants without making innocent defendants worse off, and makes the prosecutor better off.

Keywords: Bargaining, plea bargaining, second-order belief uncertainty, disclosure, Brady rule.
JEL classifications: C78, D82, D83, K40.

*Email: pcuel@bu.edu. I am indebted to Juan Ortner, Bart Lipman, and Dilip Mookherjee for invaluable guidance and support. I thank Gonzalo Cisternas, Albert Ma, Chiara Margaria, Michael Manove, Andrew Newman, Jawwad Noor, Beixi Zhou, and participants in the Microeconomic Theory Workshop at Boston University for helpful comments and suggestions. All errors are my own.

1 Introduction

1.1 Motivation and results

In many bargaining situations, one party can gather information and keep the outcome private. A real estate agent who inspects a house for sale might conceal information that would decrease its price. After reviewing a company, an investor might reveal only information that increases the return she demands. In plea bargaining, the context I examine in this paper, a prosecutor who finds exculpatory evidence might conceal it from the defendant. I study the disclosure decision of the newly informed party when disclosure is voluntary, and I ask whether mandatory disclosure is socially desirable.

Plea bargaining is of particular relevance in the U.S. criminal justice system, where more than 90% of criminal cases end in a plea bargain rather than a trial.[1] In plea bargaining, a prosecutor and a defendant negotiate a sentence to avoid trial. During the negotiation, the prosecutor can search for evidence regarding the defendant's culpability. In several U.S. jurisdictions, disclosure of exculpatory evidence is not constitutionally required during plea bargaining: the prosecutor can conceal exculpatory evidence during the negotiation but must disclose it at trial. If the prosecutor wants the judge to assign as high a sentence as possible, will she disclose exculpatory evidence? And even if she sometimes does, is it socially desirable to require disclosure during plea bargaining?

To answer these questions, I study a continuous-time plea-bargaining model between a prosecutor (she) and a defendant (he). The defendant is innocent or guilty and privately knows his type. The prosecutor starts with inconclusive default evidence and, while negotiating, investigates for conclusive hard evidence that is exculpatory if the defendant is innocent and incriminating if he is guilty. Evidence arrives at a Poisson rate, so the prosecutor may not find any by the deadline. At any time, the prosecutor may voluntarily disclose evidence she has found and make an offer, which the defendant accepts or rejects. If no agreement is reached by the deadline, the case goes to trial, modeled as a rule that assigns a sentence based on the evidence: exculpatory evidence sets the defendant free, default evidence leads to a low sentence, and incriminating evidence leads to a high sentence.

Because the prosecutor's evidence is private, voluntary disclosure creates *second-order belief uncertainty*: an innocent defendant knows he is innocent but does not know whether the prosecutor knows it. This uncertainty is what allows a prosecutor holding exculpatory evidence to pool with a prosecutor holding only default evidence. Off-path beliefs therefore play a central role, and I discipline them with a consistency requirement in the spirit of Grossman and Perry (1986): after an unexpected offer, the defendant believes it comes from the set of prosecutor types who would benefit from making it, given his response to that belief.

The first main result shows that, in some cases, the prosecutor discloses exculpatory evidence. When the prior probability of guilt is low, a prosecutor who finds exculpatory evidence conceals it and pools with the prosecutor holding default evidence on a low offer that the innocent defendant

[1] See American Bar Association (2023) and Devers (2011).

accepts. When the prior is intermediate, the prosecutor first screens the defendant with an offer that only a guilty defendant accepts (with some probability), lowering her belief enough to make pooling sustainable. When the prior is high, she discloses exculpatory evidence as soon as she finds it. In that region the prosecutor with default evidence prefers to make the high offer that only a guilty defendant accepts, and a prosecutor holding exculpatory evidence cannot mimic that offer without going to trial, which is costly. The disclosure regime is self-enforcing under the consistency requirement: an early, lower offer by a prosecutor with default evidence would be interpreted as coming from the pool of all prosecutors who could have concealed exculpatory evidence, and against that pool the best accepted offer is too low to be attractive.

The second main result compares voluntary and mandatory disclosure. First, mandatory disclosure is efficient: the parties always settle before trial, whereas under voluntary disclosure an innocent defendant goes to trial with positive probability when the prior of guilt is high. Second, mandatory disclosure is fairer in a precise sense: it raises the expected sentence of the guilty defendant and leaves that of the innocent defendant unchanged. Third, the prosecutor herself is better off under mandatory disclosure. Hiding exculpatory evidence has a downside: when the prosecutor holds only default evidence, she cannot extract the full surplus from the defendant. If she could commit ex ante to disclose everything she finds, she would.

In all regimes the game exhibits a deadline effect: agreements cluster at the deadline, as in Spier (1992). Under voluntary disclosure with a high prior of guilt, some innocent defendants do not settle at all and go to trial; under mandatory disclosure everyone settles.

Appendix A shows that these results extend to a prosecutor who cares about convicting the innocent, to a trial that sentences according to the posterior probability of guilt, and to inconclusive incriminating evidence. It also considers the case in which investigating is the prosecutor's choice. There, voluntary disclosure weakens the prosecutor's incentive to investigate, and mandatory disclosure makes innocent defendants strictly better off for intermediate priors.

Outline. Section 1.2 provides background on plea bargaining and the Brady rule, and Section 1.3 discusses the related literature. Section 2 introduces the model and the solution concept. Section 3 analyzes a one-offer benchmark. Section 4 characterizes equilibrium in continuous time. Section 5 describes the mandatory-disclosure case, and Section 6 compares the two regimes. Section 7 studies the timing of agreements, and Section 8 concludes. Appendix A presents extensions, and Appendix B contains the proofs.

1.2 Background: plea bargaining and the Brady rule

In U.S. criminal law, plea bargaining is the pretrial process in which the prosecutor and the defendant negotiate an agreement in which the defendant pleads guilty in exchange for a lower sentence.[2] A plea bargain allows the parties to avoid a trial and its associated cost and uncertainty. If they do

[2]In the U.S. system, the judge must approve the plea bargain. I assume the judge always approves an agreement reached by the prosecutor and the defendant.

not reach an agreement, the case goes to trial. The prosecutor's role is to represent society in the criminal case brought against the defendant.

At trial, the prosecutor must disclose exculpatory evidence. Trials are governed by the Brady rule, named for *Brady v. Maryland* (1963), which requires prosecutors to disclose materially exculpatory evidence in their possession to the defendant.[3] Whether the rule extends to plea bargaining is unsettled. In *United States v. Ruiz* (2002), the Supreme Court held that the Constitution does not require prosecutors to disclose material *impeachment* evidence before a defendant pleads guilty, but it did not decide whether the same holds for *exculpatory* evidence.[4] A circuit split has developed on that question. The Seventh and Tenth Circuits have suggested that material exculpatory evidence must be disclosed before a guilty plea, while the First, Fourth, and Fifth Circuits have held that Brady is a trial right that does not apply during plea bargaining (Casey, 2020; Falk, 2025). The Ninth Circuit has not taken a direct position, although its case law has been read as favoring pre-plea disclosure (Falk, 2025). The Supreme Court has repeatedly declined to resolve the split.[5] Professional ethics rules, such as ABA Model Rule 3.8(d), call for timely disclosure of evidence that tends to negate guilt, but they create no enforceable right for defendants and are rarely enforced against prosecutors (Falk, 2025). In practice, therefore, whether a prosecutor must reveal exculpatory evidence before a plea depends on the jurisdiction.

Whether to apply the Brady rule to plea bargaining is a policy question that has attracted attention from scholars and the media.[6] Arguments in favor emphasize that a guilty plea must be knowing and voluntary, which failure to disclose exculpatory evidence precludes; extending the rule would therefore reduce convictions of innocent defendants. Arguments against hold that extending the rule would raise costs and reduce efficiency. This paper shows that, in a bargaining model with endogenous second-order uncertainty, extending the Brady rule to plea bargaining *increases* efficiency, raises the sentences of guilty defendants, and benefits the prosecutor herself.

1.3 Related literature

The prosecutor investigates while negotiating, and voluntary disclosure generates second-order uncertainty for the defendant. The paper therefore relates to the literatures on pretrial negotiation, bargaining with information arrival, higher-order uncertainty in bargaining, and verifiable disclosure.

Pretrial negotiation and plea bargaining. Spier (1992) and Fuchs and Skrzypacz (2013) study pretrial bargaining with incomplete information and a deadline, and show that many agreements occur at the deadline. My model has a similar deadline effect, but the focus is on disclosure. Garoupa

[3] See *Brady v. Maryland*, 373 U.S. 83 (1963).

[4] See *United States v. Ruiz*, 536 U.S. 622 (2002).

[5] The Court denied certiorari in *Alvarez v. City of Brownsville*, 904 F.3d 382 (5th Cir. 2018), cert. denied, 139 S. Ct. 2690 (2019), and in *Mansfield v. Williamson County*, 30 F.4th 276 (5th Cir. 2022), cert. denied, 143 S. Ct. 486 (2022). See Petegorsky (2012), Grossman (2016), Casey (2020), and Falk (2025).

[6] See Casey (2020), Daughety and Reinganum (2020), and Sanders (2019). See also the New York Times editorial "Beyond the Brady Rule," May 19, 2013.

and Rizzolli (2011) study a model in which the prosecutor might decide not to investigate before trial and conclude that the Brady rule at trial may hurt innocent defendants; their mechanism operates through investigation incentives at trial, whereas mine operates through second-order uncertainty during negotiation, and I find that extending disclosure to the negotiation stage leaves innocent defendants no worse off. Daughety and Reinganum (2018) study a prosecutor with career concerns who may violate the Brady rule at trial. Closest to this paper, Daughety and Reinganum (2020) develop a dynamic model of a criminal case, from arrest through plea bargaining and possibly trial, in which prosecutors and defendants may discover exculpatory evidence and prosecutors choose whether to disclose it. They compare disclosure regimes that differ in timing and show that, when disclosure is required before plea bargaining, a prosecutor who has discovered exculpatory evidence discloses it and drops the case. My paper differs in three respects. First, evidence arrives *during* the negotiation, and the prosecutor can make offers at any time, so concealment creates second-order uncertainty that evolves with the negotiation; this is what allows a prosecutor with exculpatory evidence to pool with one holding only default evidence. Second, I characterize when a prosecutor who cares only about the sentence discloses exculpatory evidence voluntarily, and show that she does so when the prior of guilt is high, because concealment would require mimicking an offer that an innocent defendant rejects. Third, I show that the prosecutor herself prefers mandatory disclosure, because it solves a commitment problem.

In the plea-bargaining literature, Landes (1971) examines how the probability of winning at trial affects pretrial negotiations. Grossman and Katz (1983) and Reinganum (1988) study the welfare effects of plea bargaining depending on the probability of conviction. Baker and Mezzetti (2001) examine a model in which the prosecutor chooses the costly precision of a signal about the defendant's type. Bjerk (2007) presents a model in which new information can be revealed at trial. Vasserman and Yildiz (2019) study negotiating parties who are optimistic about the trial outcome and anticipate public information before trial. More recent theoretical work studies plea bargaining with richer trial stages and prosecutorial objectives. Little and Simpson (2025) show that plea bargaining may sort perversely, with innocent defendants more likely than guilty ones to plead guilty, when defendants differ in risk attitudes or prosecutors value trial wins; their prosecutor also maximizes sentences, but the evidence is common knowledge before trial. Kim (2025) obtains a separating equilibrium in which the defendant privately knows both his guilt and the strength of his own evidence and the jury's conviction decision is rational. Guha (2024) integrates plea bargaining with costly juror effort, and Christmann (2023) studies plea bargaining and investigation effort under inquisitorial procedure. On the empirical and experimental side, Lundberg (2024) asks whether prosecutors induce innocent defendants to plead guilty, and Ralston et al. (2023) study experimentally how the way conviction rates are computed affects prosecutors' plea offers. None of these papers studies the disclosure of privately acquired evidence during negotiation. Ispano and Vida (2024) study interrogations in which a law enforcement official may disclose previously acquired evidence; their focus is the optimal interrogation policy, while mine is bargaining over the sentence.

Bargaining with information arrival. Duraj (2020) studies a buyer who chooses how accurately to learn her valuation and can disclose it. Esö and Wallace (2019) consider bargaining in which the value of the good is privately revealed and can be disclosed, and show that the possibility of learning can delay agreement. Esö and Wallace (2014) show that the proposer is better off with verifiable evidence in a one-period model. Hwang and Li (2017) study a buyer whose outside option arrives stochastically and may be disclosed. In these papers, the party with new information conceals detrimental evidence and discloses beneficial evidence. In my model, the informed party discloses detrimental information, and I characterize when doing so is optimal. Daley and Green (2020), Fuchs and Skrzypacz (2010), Hwang (2018), Lomys (2017), Ortner (2017), and Ortner (2023) study variations of the Coase conjecture with arrival of new information but do not consider disclosure.

Higher-order uncertainty in bargaining. Feinberg and Skrzypacz (2005) study bargaining in which one party privately knows his type and the other holds a private belief about it; the resulting second-order uncertainty is exogenous and generates delay. In my model, second-order uncertainty is endogenous and one party can eliminate it by disclosing. Friedenberg (2019) studies delay arising from strategic uncertainty; there is no strategic uncertainty in my model.

Disclosure and off-path beliefs. Dye (1985) studies a disclosure model with an evidence structure similar to mine, and Acharya et al. (2011) study a dynamic extension. In those papers detrimental information is not disclosed, or only after negative public news; in my model the prosecutor voluntarily reveals detrimental information. More recently, Kolb and Zhou (2026) study a sender who privately obtains hard evidence at a random date and chooses when to disclose it, which is the evidence structure of my model; in their setting the receiver does not bargain with the sender. Shishkin (2022) studies a sender whose test produces evidence only with some probability and who can pretend not to have obtained it, so that the receiver faces uncertainty about whether evidence exists, as the innocent defendant does here. Moreover, in my model the receiver knows the state (his type) and the sender investigates it. Methodologically, my consistency requirement on off-path beliefs builds on Grossman and Perry (1986), Cho and Kreps (1987), and Mailath et al. (1993). Because the game is in continuous time, I follow Simon and Stinchcombe (1989) and Bergin and MacLeod (1993) in restricting strategies so that outcomes are well defined.

2 Model

There are two players, a prosecutor (she) and a defendant (he). The defendant's type is $\alpha \in \{I, G\}$ (innocent or guilty), which he knows privately. The prosecutor's prior that the defendant is guilty is $\theta \in (0, 1)$. The prosecutor's only objective is to obtain the highest possible sentence, regardless of the defendant's guilt, while the defendant wants the lowest possible sentence.[7]

[7] This is an extreme case chosen to show that even a purely sentence-motivated prosecutor discloses exculpatory evidence. Many prosecutors appear to be motivated by high sentences rather than justice: see Medwed (2004),

Time and evidence. The game is divided into two phases: *plea bargaining* and *trial*. Plea bargaining negotiation takes place in continuous time $t \in [0,1]$; the deadline is normalized to $T = 1$. If a plea agreement is not reached by the deadline, the plea bargaining phase concludes and the game goes to trial.

The prosecutor starts with default evidence $y = d$. Conclusive evidence arrives at a random time τ that is exponentially distributed with rate $\lambda > 0$, independently of the defendant's type. If $\tau \leq 1$, at time τ the prosecutor privately learns $y = e$ (exculpatory) if $\alpha = I$ and $y = h$ (incriminating) if $\alpha = G$; new evidence arrives at most once. I say that the prosecutor is *y-type* when she holds evidence $y \in \{d, e, h\}$. Define

$$q \equiv e^{-\lambda}, \qquad r(t) \equiv e^{-\lambda(1-t)},$$

so q is the probability that no evidence arrives over the whole negotiation and $r(t)$ is the probability that none arrives in $(t, 1]$; note $r(0) = q$ and $r(1) = 1$. Because arrival is independent of α, a prosecutor who has not found evidence does not update her belief about the defendant's type.

Disclosure. Evidence is hard: the prosecutor can disclose new evidence she holds and cannot fabricate evidence. Disclosure is public and can be made at any time after arrival. Under *voluntary disclosure* the prosecutor chooses whether and when to disclose; disclosure is mandatory at trial. Section 5 studies *mandatory disclosure*, under which new evidence is disclosed upon arrival.

Offers. At any time $t \in [0,1]$ the prosecutor may make an offer $x \in \mathbb{R}$, which the defendant accepts or rejects immediately. If he accepts, the game ends with payoffs $u^P = x$ and $u^D = -x$. If an offer made at $t = 1$ is rejected, or no offer is accepted by $t = 1$, the case goes to trial. There is no discounting.

Trial. The trial assigns the sentence $s(y)$ according to the evidence held at $t = 1$:

$$s = \begin{cases} 0 & \text{if } y = e, \\ d & \text{if } y = d, \\ h & \text{if } y = h, \end{cases}$$

where $0 < d < h$. Trial costs the prosecutor $c > 0$, so trial payoffs are $u^P = s(y) - c$ and $u^D = -s(y)$.[8] I call the absolute value of a defendant's expected payoff his *expected loss*.

2.1 Strategies and beliefs

A public history at time t records the evidence disclosed and the offers made and rejected before t. The prosecutor's private history adds her evidence. A pure strategy for the prosecutor maps her

Keenan et al. (2011), Garrett (2017), Pfaff (2017), and Alschuler (2015). Appendix A.1 allows the prosecutor to care about convicting the innocent.

[8] Adding a trial cost for the defendant does not change the results, since the prosecutor makes the offers and would incorporate that cost into them.

private history to a disclosure decision and an offer decision at each t; a (behavioral) strategy for the defendant maps his type, the public history, and the current offer into a probability of acceptance. To make continuous-time outcomes well defined, I restrict attention to strategies with *inertia*: on any history, the set of times at which the prosecutor makes an offer that is accepted with positive probability is finite (Simon and Stinchcombe, 1989; Bergin and MacLeod, 1993). Offers that both defendant types reject with probability one are payoff-irrelevant and I call them *non-serious*.

The guilty defendant knows that only a prosecutor holding d or h can face him; the innocent defendant knows that only a prosecutor holding d or e can face him. Let $\beta(t)$ denote the innocent defendant's belief that the prosecutor is d-type at a public history with no disclosure.

2.2 Solution concept

The solution concept is perfect Bayesian equilibrium (PBE) with the prosecutor using pure strategies, refined in four ways.

Consistent off-path beliefs. Consider a history at time t with no disclosure, and suppose the prosecutor makes an offer x that she does not make there on the equilibrium path. For $y \in \{d, e\}$, let $\ell_y(t)$ denote the innocent defendant's expected loss if he rejects x and the prosecutor of type y follows her equilibrium strategy thereafter. For a nonempty set $K \subseteq \{d, e\}$ of prosecutor types, define the posterior that would obtain if exactly the types in K could be at this history:

$$\beta_{\{d\}}(t) = 1, \qquad \beta_{\{e\}}(t) = 0, \qquad \beta_{\{d,e\}}(t) = e^{-\lambda t}.$$

The last expression counts as part of the pool every e-type whose evidence arrived before t and who *could* have concealed it: conditional on innocence, the probability that no evidence arrived by t is $e^{-\lambda t}$. Given β_K, the innocent defendant accepts x if $x < \beta_K \ell_d(t) + (1 - \beta_K)\ell_e(t)$ and rejects if the inequality is reversed; indifference is resolved as described below. The guilty defendant's response is unaffected by β_K.

Definition 1 (Consistency)**.** A set K is *self-consistent* for the offer x if K is exactly the set of types $y \in \{d, e\}$ that strictly prefer making x, given the defendants' responses when the innocent defendant holds belief β_K, to following the equilibrium. A PBE is *consistent* if no off-path offer at a history without disclosure admits a nonempty self-consistent set.

If some off-path offer admitted a self-consistent set K, then beliefs equal to β_K would make every type in K strictly gain from the deviation; consistency rules out equilibria that are vulnerable to such credible deviations, as in Grossman and Perry (1986). The definition is applied at every history, including the deadline. Standard refinements based on equilibrium dominance, such as the intuitive criterion of Cho and Kreps (1987), do not pin down beliefs here: at the offers that matter, both the d-type and the e-type would gain if the offer were accepted, so neither type can be excluded. Consistency resolves this by requiring the belief to be self-confirming: the types to

whom the defendant attributes an offer must be exactly those who would want to make it given his response.

Tie-breaking at the default sentence. At the deadline, an innocent defendant who believes the prosecutor is d-type is indifferent between accepting $x = d$ and going to trial. I let him accept with probability μ, and maintain throughout:

Assumption 1. $\mu \in [0, \bar{\mu}]$ with $\bar{\mu} \equiv c/(c+d)$, and $(1-\mu)c > d(1-q)$.

At every other indifference, the defendant accepts. The upper bound $\bar{\mu}$ is necessary for exculpatory evidence ever to be disclosed: if the innocent defendant accepted d with probability above $\bar{\mu}$, a prosecutor with exculpatory evidence would conceal it and offer d. The second part of Assumption 1 says that trials are not too cheap; with $\mu = 0$ it reduces to $c > d(1-q)$.[9] I keep μ general because the welfare comparisons depend on it quantitatively.

Stationary disclosure rules. I restrict attention to equilibria in which either every e-type discloses upon arrival, or no e-type discloses before the deadline. This rules out equilibria in which the disclosure decision depends on the arrival time of the evidence. Such equilibria are fragile: an e-type who discloses obtains 0, while an e-type who conceals pools with the d-type and obtains the d-type's offers, so whenever concealment yields a positive payoff to some e-types, every e-type has an incentive to conceal.

Selection. I break payoff ties as follows. (i) A prosecutor who will offer $x = 0$ to an innocent defendant discloses exculpatory evidence (rather than making the offer without disclosing). (ii) An h-type discloses and settles upon arrival. (iii) Absent discounting, payoffs do not depend on when a given expected transfer takes place, so the timing of agreements is not pinned down by payoffs alone. I select equilibria in which, absent new evidence, the prosecutor makes serious offers only at $t = 1$, and at $t = 0$ when a serious offer at that date is needed to sustain the equilibrium. (iv) Among the consistent equilibria with stationary disclosure rules, I select the one that maximizes the prosecutor's payoff at $t = 0$. At $t = 0$ no evidence can have arrived, so the prosecutor has no private information; criterion (iv) is the proposer's choice of regime, not a signaling device. I do not characterize the full set of consistent equilibria.

3 A One-Offer Benchmark

I first consider the simplest version of the game: evidence arrives over $[0, 1]$ as above, but the prosecutor can make a single offer, at $t = 1$. The probability that she holds new evidence at that point is $1 - q$. Define

$$\tilde{\theta}(\mu) \equiv 1 - \frac{d(1-q)}{(1-\mu)c}.$$

[9] If it fails, the prosecutor discloses exculpatory evidence for every prior.

By Assumption 1, $\tilde{\theta}(\mu) \in (0,1)$; with $\mu = 0$ it equals $1 - d(1-q)/c$.

Proposition 1 (Disclosure with a single offer)**.** *In the one-offer benchmark, the consistent equilibrium outcome is as follows. A prosecutor with incriminating evidence discloses it and offers h, which the guilty defendant accepts.*

(i) If $\theta \leq \tilde{\theta}(\mu)$, a prosecutor with exculpatory evidence conceals it. Both the d-type and the e-type offer $x = dq$, which both defendant types accept.

(ii) If $\theta > \tilde{\theta}(\mu)$, a prosecutor with exculpatory evidence discloses it and offers $x = 0$, which the innocent defendant accepts. The d-type offers $x = d$; the guilty defendant accepts, and the innocent defendant accepts with probability μ and otherwise goes to trial.

The argument proceeds in six steps. The first two describe what each defendant type believes about the prosecutor's evidence; the remaining four derive the prosecutor's offers and disclosure decision from those beliefs.

Step 1: Incriminating evidence is always disclosed. A prosecutor who finds $y = h$ discloses it and offers $x = h$. After disclosure, the guilty defendant's loss at trial is h, so rejecting $x = h$ does not help him, and he accepts. No offer above h is ever accepted by a guilty defendant, so disclosure gives the h-type the highest payoff she can obtain.

This has an immediate implication for the guilty defendant's beliefs. Because every h-type discloses, a guilty defendant who sees no disclosure knows that the prosecutor did not find incriminating evidence. A guilty defendant can face only a d-type or an h-type, so

$$P^G(d\text{-type} \mid \text{no disclosure}) = 1.$$

His expected loss at trial is therefore d, and he accepts any offer $x \leq d$.

Step 2: Concealment creates second-order uncertainty for the innocent defendant. Suppose that a prosecutor who finds exculpatory evidence conceals it. The innocent defendant knows he is innocent, but he does not know whether the prosecutor knows it: the investigation either found exculpatory evidence (probability $1 - q$) or found nothing (probability q), and both outcomes look the same to him. His beliefs about the prosecutor's evidence after no disclosure are therefore

$$P^I(d\text{-type} \mid \text{no disclosure}) = q, \qquad P^I(e\text{-type} \mid \text{no disclosure}) = 1 - q.$$

At trial, all evidence must be disclosed, so the innocent defendant receives sentence d if the prosecutor holds default evidence and 0 if she holds exculpatory evidence. His expected loss at trial is

$$q \cdot d + (1-q) \cdot 0 = dq,$$

and he accepts any offer $x \leq dq$. Table 1 summarizes the two defendant types' beliefs and trial losses when exculpatory evidence is concealed.

Defendant	$P(d\text{-type} \mid \text{no disclosure})$	Expected loss at trial	Accepts x if
Guilty	1	d	$x \leq d$
Innocent	q	dq	$x \leq dq$

Table 1: Defendants' beliefs and trial losses when exculpatory evidence is concealed.

Second-order uncertainty is what makes concealment valuable. An innocent defendant who knew that the prosecutor had found exculpatory evidence would accept nothing above 0. Because he cannot rule out that she holds only default evidence, he is willing to accept $dq > 0$.

Step 3: The d-type's offer. Given these acceptance thresholds, a d-type prosecutor has two sensible offers:

- *Offer* $x = dq$. Both defendant types accept, and the prosecutor's payoff is dq.
- *Offer* $x = d$. The guilty defendant accepts. The innocent defendant's expected trial loss is only $dq < d$, so he rejects, and the prosecutor goes to trial with an innocent defendant while holding default evidence, obtaining $d - c$.[10] Her expected payoff is $\theta d + (1-\theta)(d-c)$.

Any offer strictly between dq and d is accepted only by the guilty defendant and is dominated by $x = d$; any offer below dq is dominated by $x = dq$. Comparing the two,

$$dq \;\geq\; \theta d + (1-\theta)(d-c) \quad \Longleftrightarrow \quad \theta \leq 1 - \frac{d(1-q)}{c} = \tilde{\theta}(0).$$

The trade-off is intuitive. Offering d extracts an extra $d - dq = d(1-q)$ from a guilty defendant, but costs $c - d(1-q)$ in expectation from an innocent defendant, since the prosecutor pays the trial cost instead of receiving dq and obtains d at trial. When guilt is unlikely, the second effect dominates and the prosecutor offers the low, pooling sentence dq.

Step 4: Low prior ($\theta \leq \tilde{\theta}(\mu)$): the e-type conceals. A prosecutor who finds exculpatory evidence knows the defendant is innocent. If she discloses, the innocent defendant accepts nothing above 0. If she conceals, she can make the same offer as the d-type, $x = dq$. Because the d-type and the e-type make the same offer, the innocent defendant learns nothing from it: his belief stays at q, his trial loss stays at dq, and he accepts. Concealment gives the e-type $dq > 0$, so she conceals.

The beliefs in Step 2 are then confirmed. With the e-type concealing, a fraction $1 - q$ of the prosecutors who face an innocent defendant without disclosing hold exculpatory evidence, exactly as the innocent defendant believed. If the prosecutor instead offered some $x > dq$ without disclosing, both the d-type and the e-type would like the offer to be accepted, so the innocent defendant attributes it to the whole pool, keeps belief q, and rejects it.

[10] If the innocent defendant knows the prosecutor holds default evidence, he is indifferent at $x = d$ and accepts with probability μ; this is the case in the disclosure regime of Step 5, where the prosecutor's payoff becomes $\theta d + (1-\theta)[\mu d + (1-\mu)(d-c)]$. To keep the exposition simple, I write the trade-off here for $\mu = 0$; the general case only replaces c by $(1-\mu)c$.

Step 5: High prior ($\theta > \tilde{\theta}(\mu)$): the e-type cannot imitate the d-type. Now the d-type prefers to offer $x = d$ (Step 3). To conceal exculpatory evidence, the e-type would have to make an offer that does not reveal her evidence *and* is accepted by the innocent defendant. She cannot do both.

- *Imitate the d-type and offer $x = d$.* The innocent defendant accepts with probability at most μ, and otherwise the case goes to trial, where the exculpatory evidence must be disclosed. The e-type obtains $\mu d - (1-\mu)c \leq 0$ because $\mu \leq \bar{\mu}$. Imitation does not pay.

- *Conceal but offer something lower, $x \in (0, d)$.* The innocent defendant asks who could be making this offer. If he thought it came from the whole pool of d- and e-types, his belief would be q and he would accept only up to dq. But a d-type does not want to offer dq when $\theta > \tilde{\theta}(\mu)$; she prefers d. So an offer below d can only come from an e-type. The innocent defendant infers that the prosecutor holds exculpatory evidence,

$$P^I(d\text{-type} \mid \text{no disclosure},\ x < d) = 0,$$

  so his expected loss at trial is 0, and he rejects every $x > 0$. The e-type then obtains $-c$.

Since every way of concealing yields at most 0, the e-type discloses and offers $x = 0$, which the innocent defendant accepts.

The innocent defendant's beliefs change accordingly. In this regime the e-type always discloses, so an innocent defendant who sees no disclosure and the offer $x = d$ knows the prosecutor holds default evidence:

$$P^I(d\text{-type} \mid \text{no disclosure},\ x \geq d) = 1.$$

His trial loss is then exactly d; he is indifferent, accepts with probability μ, and otherwise goes to trial.

Step 6: Why the innocent defendant cannot always accept $x = d$. One might ask why the innocent defendant does not simply accept $x = d$ with probability one, since he is indifferent. If he did, the e-type would conceal her evidence and offer d, obtaining $d > 0$. But then the pool of prosecutors offering d would again contain e-types, the innocent defendant's belief would fall back to q, his trial loss to $dq < d$, and he would strictly prefer to reject. Hence acceptance of d must be limited to $\mu \leq \bar{\mu} = c/(c+d)$, the largest probability at which the e-type still prefers disclosure to imitation.

3.1 Inefficiency and commitment

When $\theta > \tilde{\theta}(\mu)$, the parties go to trial with probability $(1-\theta)q(1-\mu)$, which is socially wasteful. Moreover, the prosecutor would like to commit to disclosing everything she finds: with voluntary disclosure, either her default-evidence offer to guilty defendants is depressed to dq, or she incurs trial costs with innocent defendants. Section 6 makes this precise.

4 Continuous-Time Negotiation

I now allow offers at any time in $[0,1]$. Two features distinguish this case from the benchmark. First, the prosecutor can learn about the defendant's type through rejected offers. Second, she can settle before the deadline. I show that the disclosure result survives, that early settlement is not profitable under consistent beliefs, and that the possibility of screening enlarges the range of priors in which exculpatory evidence is concealed.

Three regimes are candidates for equilibrium. In all of them incriminating evidence is disclosed upon arrival and the guilty defendant accepts h.

Concealment (C). No e-type discloses. Before the deadline the prosecutor makes only non-serious offers; at $t = 1$ the d-type and the e-type offer dq, which both defendant types accept.

Screening (S). At $t = 0$ the prosecutor offers $x_0 \equiv (1-q)h + q\,dq$. The guilty defendant accepts with probability

$$s(\theta) \equiv \frac{\theta - \tilde{\theta}(\mu)}{\theta\big(1-\tilde{\theta}(\mu)\big)},$$

the innocent defendant rejects, and after a rejection the prosecutor's belief is $\tilde{\theta}(\mu)$. Play then continues as in regime C.

Disclosure (D). Every e-type discloses upon arrival and offers 0, which the innocent defendant accepts. Before the deadline the d-type makes only non-serious offers; at $t = 1$ she offers d, which the guilty defendant accepts and the innocent defendant accepts with probability μ.

Lemma 1 (Continuation values)**.** *Consider a d-type prosecutor at time t with belief θ, and write $r = r(t)$.*

(i) In regime D, the innocent and guilty defendants' expected losses are rd and $(1-r)h + rd$, and the prosecutor's continuation value is

$$V^D(t;\theta) = \theta\big[(1-r)h + rd\big] + (1-\theta)\,r\big(d - (1-\mu)c\big).$$

(ii) In regime C, the innocent and guilty defendants' expected losses are dq and $(1-r)h + r\,dq$, and the prosecutor's continuation value is

$$V^C(t;\theta) = \theta\big[(1-r)h + r\,dq\big] + (1-\theta)\,dq.$$

(iii) The ex ante value of regime S equals $V^C(0;\theta)$.

The innocent defendant's loss in regime C does not depend on t: whatever happens before the deadline, he ends up paying dq. In regime S, the guilty defendant is indifferent between paying x_0 and waiting, so the prosecutor's ex ante value is the same as in regime C; screening matters because

it makes regime C sustainable when $\theta > \tilde{\theta}(\mu)$.

Define

$$\bar{\theta}(\mu) \equiv \frac{(1-\mu)c}{(1-\mu)c + d(1-q)}.$$

Lemma 2 (Cutoffs)**.** *$0 < \tilde{\theta}(\mu) < \bar{\theta}(\mu) < 1$, and $V^C(0;\theta) \geq V^D(0;\theta)$ if and only if $\theta \leq \bar{\theta}(\mu)$.*

The two cutoffs answer different questions. The cutoff $\tilde{\theta}(\mu)$ is the deadline comparison: a d-type who reaches $t = 1$ without new evidence chooses between pooling at dq, which both defendant types accept, and offering d, which only the guilty type accepts for sure. Offering d gains $d(1-q)$ per defendant but costs the trial with innocent defendants, so she prefers dq if and only if $\theta \leq \tilde{\theta}(\mu)$. The cutoff $\bar{\theta}(\mu)$ is the ex ante comparison of the two regimes, which can be read off the difference

$$V^C(0;\theta) - V^D(0;\theta) = -\theta\, qd(1-q) + (1-\theta)\, q(1-\mu)c.$$

Concealment costs the prosecutor $d(1-q)$ on guilty defendants when no evidence arrives, and saves her the trial cost on innocent defendants when no evidence arrives. Ex ante, the innocent defendant pays dq in both regimes, because under concealment the e-types collect dq from him. At the deadline, by contrast, the d-type no longer takes into account what the e-types collect, so she undervalues concealment. This is why $\bar{\theta}(\mu) > \tilde{\theta}(\mu)$: for priors between the two cutoffs, the prosecutor prefers concealment ex ante but would abandon it at the deadline.

Lemma 3 (Concealment)**.** *Regime C is a consistent equilibrium if and only if $\theta \leq \tilde{\theta}(\mu)$.*

Concealment is easy to sustain before the deadline: whatever the innocent defendant believes, he ends up paying dq at $t = 1$, so he never accepts more than dq earlier and the prosecutor has nothing to gain from an early offer. The binding constraint is at the deadline. If $\theta > \tilde{\theta}(\mu)$, a d-type prefers to offer d rather than dq, and this deviation is credible: only a d-type would make it, because an e-type who offered d would be rejected and go to trial. The innocent defendant therefore attributes the offer to a d-type, and concealment breaks down. If $\theta \leq \tilde{\theta}(\mu)$, the d-type has no reason to deviate, and the pooling offer dq is stable.

Lemma 4 (Screening)**.** *For every $\theta > \tilde{\theta}(\mu)$, regime S is a consistent equilibrium.*

Screening resolves the problem identified in Lemma 3. When $\theta > \tilde{\theta}(\mu)$, concealment fails only because the d-type's belief at the deadline is too high. At $t = 0$, before any evidence can arrive, the prosecutor makes an offer x_0 that only the guilty defendant is willing to accept, and he accepts it with just enough probability that a rejection lowers her belief to $\tilde{\theta}(\mu)$. From then on, concealment is sustainable. The guilty defendant is indifferent because x_0 equals exactly what he expects to pay if he waits, and the screening offer costs the prosecutor nothing: her ex ante payoff is the same $V^C(0;\theta)$ she would obtain if concealment were sustainable without screening. Screening is available for every $\theta > \tilde{\theta}(\mu)$, but the prosecutor uses it only when concealment is the regime she prefers, that is, when $\theta \leq \bar{\theta}(\mu)$.

Lemma 5 (Disclosure)**.** *Regime D is a consistent equilibrium if and only if $V^D(t;\theta) \geq dq$ for all $t \in [0,1]$. In particular, it is a consistent equilibrium for every $\theta > \bar{\theta}(\mu)$.*

The key step in Lemma 5 concerns early settlement. In regime D, a d-type at time t might offer the innocent defendant's continuation loss $r(t)d$, which approaches d near the deadline and would let her avoid the trial cost. If the innocent defendant believed this offer came from a d-type, he would accept it. But then every e-type who found evidence before t would also want to make it, so the belief $\beta_{\{d\}} = 1$ is not self-consistent. Against the pool $\{d, e\}$, the innocent defendant's acceptable offer is

$$\beta_{\{d,e\}}(t)\,\ell_d(t) = e^{-\lambda t} \cdot r(t)\,d = dq,$$

the same at every t. A d-type prefers waiting to settling at dq exactly when $V^D(t;\theta) \geq dq$. Because V^D is affine in r, it suffices to check $t = 0$ and $t = 1$, and both conditions hold whenever $\theta > \bar{\theta}(\mu)$.

4.1 Equilibrium characterization

Proposition 2 (Equilibrium in continuous time)**.** *Under Assumption 1, the selected consistent equilibrium is:*

(i) regime C if $\theta \leq \tilde{\theta}(\mu)$;

(ii) regime S if $\tilde{\theta}(\mu) < \theta \leq \bar{\theta}(\mu)$;

(iii) regime D if $\theta > \bar{\theta}(\mu)$.

In particular, a prosecutor with exculpatory evidence discloses it upon arrival if and only if $\theta > \bar{\theta}(\mu)$, and a prosecutor with incriminating evidence always discloses it upon arrival.

Figure 1 summarizes the characterization.

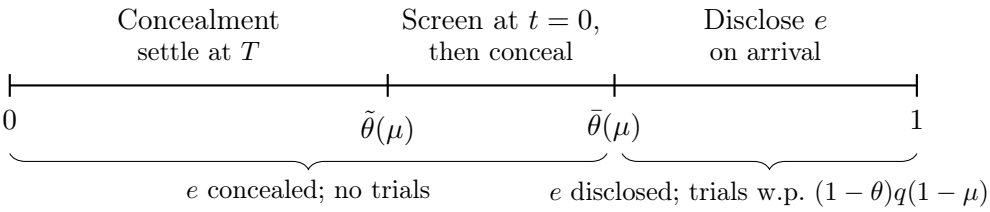


Figure 1: Equilibrium regimes in continuous time as a function of the prior θ.

The trade-off behind $\bar{\theta}(\mu)$ is the following. Concealing exculpatory evidence lets the prosecutor obtain dq from an innocent defendant even when she knows he is innocent, but it forces her to accept dq from a guilty defendant when she holds only default evidence. Disclosing lets her obtain d from a guilty defendant with default evidence, but she obtains nothing from an innocent defendant when she finds exculpatory evidence and pays the trial cost when she does not settle with him. When guilt is likely enough, the second option dominates.

Remark 1 (The role of beliefs). Consistency is what makes regime D robust to early settlement. Suppose instead that, after an unexpected offer at a history without disclosure, the innocent defendant kept his on-path belief that the prosecutor is d-type. Then a d-type at time t could settle at $r(t)d$, gaining

$$(1-\theta)\, r(t)(1-\mu)c - \theta\big(1-r(t)\big)h$$

relative to $V^D(t;\theta)$. This gain is positive for t close enough to 1 for every $\theta < 1$, so regime D would not be an equilibrium. Such "passive" beliefs, however, are not self-consistent: if early offers were accepted at that price, prosecutors with exculpatory evidence would conceal and make them too.

4.2 Comparison with the one-offer benchmark

Corollary 1 (Value of negotiation flexibility)**.** *Compared with the one-offer benchmark, in continuous time: exculpatory evidence is disclosed for a smaller set of priors ($\theta > \bar{\theta}(\mu)$ instead of $\theta > \tilde{\theta}(\mu)$); the probability of trial is weakly lower for every θ; the prosecutor's ex ante payoff is weakly higher and strictly higher for $\theta \in (\tilde{\theta}(\mu), \bar{\theta}(\mu))$; the guilty defendant's expected loss is weakly lower and strictly lower for $\theta \in (\tilde{\theta}(\mu), \bar{\theta}(\mu)]$; and the innocent defendant's expected loss is dq in both cases.*

The ability to screen at $t = 0$ lets the prosecutor transfer part of the guilty defendant's surplus to herself while preserving the second-order uncertainty that allows her to extract dq from innocent defendants.

5 Mandatory Disclosure

The Brady rule requires prosecutors to disclose materially exculpatory evidence at trial; as discussed in Section 1.2, whether it extends to plea bargaining is unsettled.

In this section, suppose that the Brady rule applies to plea bargaining: any new evidence is disclosed upon arrival. The prosecutor no longer induces second-order uncertainty, because the defendant always knows the prosecutor's evidence.

Lemma 6 (Mandatory disclosure)**.** *Under mandatory disclosure, in equilibrium: if e arrives, the prosecutor offers 0; if h arrives, she offers h; both are accepted immediately. If no evidence arrives, she makes only non-serious offers before $t = 1$ and offers d at $t = 1$, which both defendant types accept.*[11]

Without voluntary disclosure there is no second-order uncertainty: at every moment the defendant knows exactly what evidence the prosecutor holds, so every offer is evaluated against the sentence the trial would actually impose. When exculpatory or incriminating evidence arrives, the trial outcome is certain (0 or h), and the parties settle on it immediately. When no evidence has arrived, the prosecutor has no reason to settle early. At time t the innocent defendant would accept at most $r(t)d$, which is below d because exculpatory evidence may still arrive; by waiting, the

[11] Here the innocent defendant must accept d with probability one: otherwise the prosecutor would offer $d - \varepsilon$ and would have no best response.

prosecutor lets that possibility resolve itself and, against a guilty defendant, keeps the chance of finding incriminating evidence. At the deadline the default evidence is final, both defendant types face a trial sentence of d, and both accept d. In every case the prosecutor obtains the trial sentence without paying for a trial.

Corollary 2 (Efficiency under mandatory disclosure)**.** *Under mandatory disclosure the parties never go to trial. The prosecutor's ex ante payoff is*

$$V^M(\theta) \equiv \theta\big[(1-q)h + qd\big] + (1-\theta)\,qd.$$

The prosecutor's payoff under mandatory disclosure is simply the expected trial sentence. With probability $1-q$ evidence arrives and she obtains h from a guilty defendant and 0 from an innocent one; with probability q no evidence arrives and she obtains d from either type. Because agreement always replicates the trial outcome, the trial cost is never incurred, and the prosecutor extracts the full surplus from both defendant types. This payoff is the benchmark against which Section 6 measures the losses from voluntary disclosure: under voluntary disclosure, the prosecutor either receives only dq instead of d from a guilty defendant when she holds default evidence, or pays the trial cost when an innocent defendant rejects d.

6 Voluntary versus Mandatory Disclosure

I compare the two regimes in terms of efficiency, fairness, and the prosecutor's payoff. Throughout, "voluntary disclosure" refers to the equilibrium in Proposition 2; the same comparisons hold in the one-offer benchmark with $\bar{\theta}(\mu)$ replaced by $\tilde{\theta}(\mu)$.

Efficiency. Under voluntary disclosure the probability of trial is

$$P(\text{trial}) = \begin{cases} 0 & \text{if } \theta \le \bar{\theta}(\mu), \\ (1-\theta)\,q\,(1-\mu) & \text{if } \theta > \bar{\theta}(\mu), \end{cases}$$

while under mandatory disclosure it is zero. Since sentences are transfers, the social loss from voluntary disclosure equals the expected trial cost $(1-\theta)q(1-\mu)c$ when $\theta > \bar{\theta}(\mu)$. Its magnitude depends on μ, the probability that an innocent defendant accepts the default sentence when he is indifferent. Under voluntary disclosure this probability cannot exceed $\bar{\mu} < 1$, because otherwise a prosecutor with exculpatory evidence would conceal it and offer d; under mandatory disclosure no such constraint exists, and the innocent defendant accepts d with probability one (footnote 11). The trial cost is therefore positive for every admissible μ.

Fairness. I call a system fair if it assigns sentence 0 to the innocent defendant and h to the guilty defendant.

Proposition 3 (Defendants). *(i) The innocent defendant's expected loss is dq under both voluntary and mandatory disclosure.*

(ii) The guilty defendant's expected loss is $(1-q)h + qd$ under mandatory disclosure. Under voluntary disclosure it is $(1-q)h + q\,dq$ if $\theta \leq \bar{\theta}(\mu)$ and $(1-q)h + qd$ if $\theta > \bar{\theta}(\mu)$.

Mandatory disclosure therefore moves outcomes toward the fair benchmark by raising the expected sentence of guilty defendants, without affecting innocent defendants. The innocent defendant's loss is the same because, under voluntary disclosure, he either pays dq when evidence is concealed or faces the default sentence d with probability q when it is disclosed.

Commitment. Voluntary disclosure gives the prosecutor an option that mandatory disclosure takes away: after finding exculpatory evidence, she can decide whether to reveal it. One might expect a player who cares only about the sentence to value this option, since she can always choose to disclose and can conceal when concealment pays. Suppose the prosecutor could choose the disclosure regime before the negotiation starts and commit to it. Which regime would she choose? The answer is that she prefers to give up the option. The defendant anticipates that a prosecutor who can conceal exculpatory evidence will do so whenever it pays, and this anticipation lowers what he is willing to accept when she holds only default evidence.

Proposition 4 (The prosecutor prefers mandatory disclosure). *For every $\theta \in (0,1)$, the prosecutor's ex ante payoff is strictly higher under mandatory disclosure. The difference is $\theta qd(1-q)$ if $\theta \leq \bar{\theta}(\mu)$ and $(1-\theta)q(1-\mu)c$ if $\theta > \bar{\theta}(\mu)$.*

When the prosecutor conceals exculpatory evidence, she obtains the same amount from an innocent defendant as under mandatory disclosure, but only dq instead of d from a guilty defendant when she holds default evidence. When she discloses, she obtains the same amount from a guilty defendant, but pays the trial cost when an innocent defendant rejects d. In the disclosure region the prosecutor's loss equals the social loss from trials, since the defendants' losses are unchanged. Because she cannot commit not to conceal exculpatory evidence, the defendant anticipates concealment and the prosecutor cannot extract the full surplus. Figure 2 illustrates.

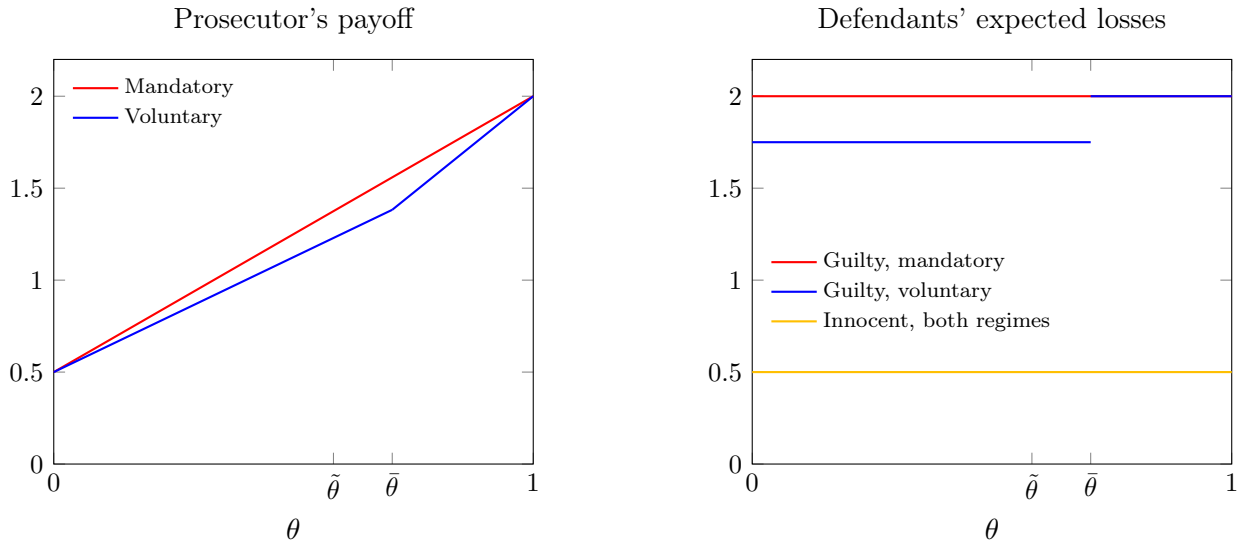


Figure 2: Payoffs and losses under voluntary and mandatory disclosure.
Parameters: $h = 3$, $d = 1$, $c = 1.2$, $q = 1/2$, $\mu = 0$, so $\tilde{\theta} \approx 0.58$ and $\bar{\theta} \approx 0.71$.

7 The Timing of Agreements

Deadline effects in pretrial negotiation have been studied by Spier (1992), Ma and Manove (1993), Fuchs and Skrzypacz (2013), and Vasserman and Yildiz (2019), and observed experimentally by Roth et al. (1988) and Güth et al. (2005).[12] Let τ^A denote the time at which the game ends by agreement; if no agreement is reached by $t = 1$, the case goes to trial.

Proposition 5 (Deadline effects)**.** *Under voluntary disclosure, the distribution of τ^A is as follows.*

(i) If $\theta \leq \tilde{\theta}(\mu)$: for the guilty defendant, $P(\tau^A \leq t) = 1 - e^{-\lambda t}$ for $t < 1$ and $P(\tau^A \leq 1) = 1$; for the innocent defendant, $P(\tau^A \leq t) = 0$ for $t < 1$ and $P(\tau^A \leq 1) = 1$.

(ii) If $\tilde{\theta}(\mu) < \theta \leq \bar{\theta}(\mu)$: for the guilty defendant, $P(\tau^A = 0) = s(\theta)$, $P(\tau^A \leq t) = s(\theta) + (1 - s(\theta))(1 - e^{-\lambda t})$ for $t \in (0, 1)$, and $P(\tau^A \leq 1) = 1$; for the innocent defendant, as in (i).

(iii) If $\theta > \bar{\theta}(\mu)$: for the guilty defendant, as in (i); for the innocent defendant, $P(\tau^A \leq t) = 1 - e^{-\lambda t}$ for $t < 1$ and $P(\tau^A \leq 1) = 1 - q(1 - \mu)$, so he goes to trial with probability $q(1 - \mu)$.

Under mandatory disclosure, for both defendant types, $P(\tau^A \leq t) = 1 - e^{-\lambda t}$ for $t < 1$ and $P(\tau^A \leq 1) = 1$.

In every case there is a mass point of agreements at the deadline. The reason is that, absent discounting, a prosecutor who holds only default evidence has no reason to settle early: the possibility that conclusive evidence arrives before the deadline keeps waiting valuable, while an early agreement would require her to give up that option. Agreements before the deadline therefore occur

[12] Other papers that find deadline effects include Cramton and Tracy (1992) and Fershtman and Seidmann (1993).

only when evidence arrives, and every case without new evidence is resolved at T. This differs from the deadline effect in Spier (1992), which arises from screening a privately informed defendant; here it arises from the arrival of evidence during the negotiation.

Under voluntary disclosure, the timing of agreements differs across defendant types, and the differences reflect how exculpatory evidence is handled. In the concealment regime, an innocent defendant never settles before the deadline, even when the prosecutor has found exculpatory evidence. A prosecutor who settled early after finding such evidence would reveal it, because a prosecutor with default evidence would not settle early; to keep the evidence hidden, she must wait and make the same offer at T as the prosecutor with default evidence. A guilty defendant, by contrast, settles as soon as incriminating evidence arrives, since that evidence is always disclosed. In the screening regime, the same pattern holds after an initial mass of agreements at $t = 0$, when some guilty defendants accept the screening offer. In the disclosure regime, both types settle as soon as evidence arrives, so before the deadline the two distributions coincide. At the deadline, however, they diverge: the guilty defendant accepts d, while the innocent defendant accepts only with probability μ and otherwise goes to trial.

Under mandatory disclosure, the distribution is the same for both types and every case settles: evidence ends the negotiation when it arrives, whatever it shows, and absent evidence both types accept the default sentence at the deadline. Comparing the two regimes, when exculpatory evidence is concealed, an innocent defendant for whom it has been found is not released when the evidence arrives, as he would be under mandatory disclosure; instead, he pleads guilty at the deadline and receives the sentence dq. This suggests an empirical implication: where disclosure before a plea is not required, guilty pleas by defendants who are later exonerated should be concentrated close to the trial date.

Figure 3 illustrates Proposition 5.

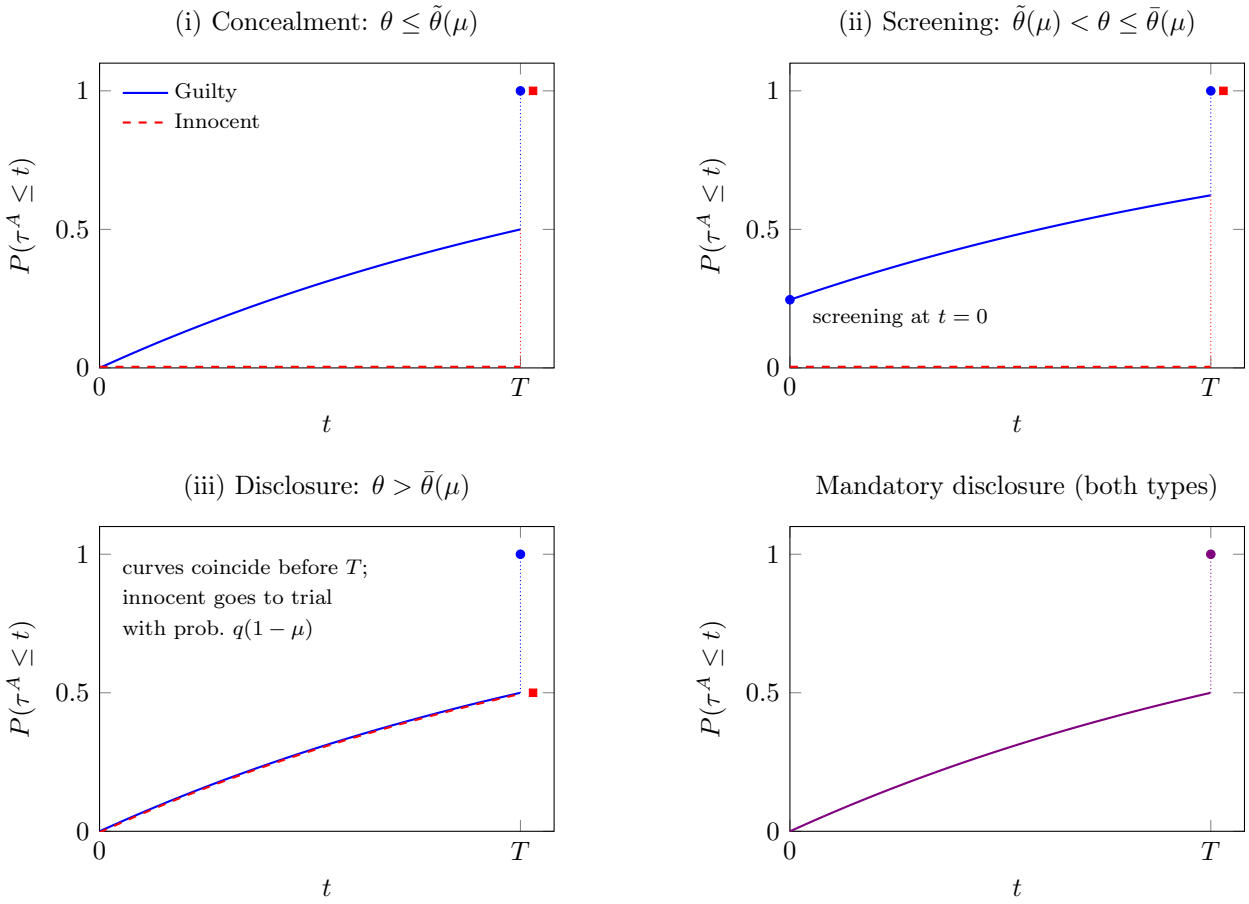


Figure 3: Probability that the game ends by agreement by time t (Proposition 5). Parameters as in Figure 2: $q = 1/2$ ($\lambda = \ln 2$), $\mu = 0$, so $\tilde{\theta} \approx 0.58$ and $\bar{\theta} \approx 0.71$; panel (ii) uses $\theta = 0.65$, for which $s(\theta) \approx 0.25$.

8 Concluding Remarks

With voluntary disclosure of evidence in plea bargaining, the prosecutor conceals exculpatory evidence when the prior probability of guilt is low or intermediate, screening guilty defendants first in the intermediate case, and discloses it when the prior is high. A purely sentence-motivated prosecutor may therefore still disclose exculpatory evidence. The result rests on a consistency requirement on beliefs: an unexpected early offer is interpreted as coming from the pool of all prosecutors who could have concealed exculpatory evidence. Without it, a prosecutor with default evidence would settle just before the deadline and the disclosure regime would unravel.

Mandatory disclosure during plea bargaining is socially desirable for two reasons: it is efficient, because the parties always settle, and it better aligns sentences with culpability, because it raises the expected sentence of guilty defendants without affecting innocent defendants. Finally, the prosecutor herself prefers mandatory disclosure, because it lets her commit to a disclosure policy that she cannot commit to under voluntary disclosure.

The extensions in Appendix A suggest that these conclusions are robust, and that they become stronger when investigation is the prosecutor's choice. Allowing exculpatory evidence to be inconclusive, so that guilty defendants also face second-order uncertainty, is left for future work.

References

Acharya, V. V., DeMarzo, P., and Kremer, I. (2011). Endogenous information flows and the clustering of announcements. *American Economic Review*, 101(7):2955–2979.

Alschuler, A. W. (2015). A nearly perfect system for convicting the innocent. *Albany Law Review*, 79:919.

American Bar Association (2023). Plea bargain task force report.

Baker, S. and Mezzetti, C. (2001). Prosecutorial resources, plea bargaining, and the decision to go to trial. *Journal of Law, Economics, and Organization*, 17(1):149–167.

Bergin, J. and MacLeod, W. B. (1993). Continuous time repeated games. *International Economic Review*, 34(1):21–37.

Bjerk, D. (2007). Guilt shall not escape or innocence suffer? The limits of plea bargaining when defendant guilt is uncertain. *American Law and Economics Review*, 9(2):305–329.

Casey, C. (2020). Lost opportunity: Supreme Court declines to resolve circuit split on Brady obligations during plea-bargaining. *Boston College Law Review*, 61(9):E-Supp.

Cho, I.-K. and Kreps, D. M. (1987). Signaling games and stable equilibria. *Quarterly Journal of Economics*, 102(2):179–221.

Christmann, R. (2023). Plea bargaining and investigation effort: Inquisitorial criminal procedure as a three-player game. *European Journal of Law and Economics*, 56(3):497–532.

Cramton, P. C. and Tracy, J. S. (1992). Strikes and holdouts in wage bargaining: Theory and data. *American Economic Review*, 82(1):100–121.

Daley, B. and Green, B. (2020). Bargaining and news. *American Economic Review*, 110(2):428–474.

Daughety, A. F. and Reinganum, J. F. (2018). Evidence suppression by prosecutors: Violations of the Brady rule. *Journal of Law, Economics, and Organization*, 34(3):475–510.

Daughety, A. F. and Reinganum, J. F. (2020). Reducing unjust convictions: Plea bargaining, trial, and evidence disclosure. *Journal of Law, Economics, and Organization*, 36(2):378–414.

Devers, L. (2011). Plea and charge bargaining. Research summary, Bureau of Justice Assistance.

Duraj, J. (2020). Bargaining with endogenous learning. Working paper, SSRN 3615758.

Dye, R. A. (1985). Disclosure of nonproprietary information. *Journal of Accounting Research*, 23(1):123–145.

Esö, P. and Wallace, C. (2014). Information and evidence in bargaining. *Economic Theory Bulletin*, 2(1):23–32.

Esö, P. and Wallace, C. (2019). Bargaining with hard evidence. *Economic Journal*, 129(621):2039–2063.

Falk, L. (2025). Brady rights in a system of pleas: Analyzing the Ninth Circuit's "apparent position." *Arizona Law Review*, 67:241–264.

Feinberg, Y. and Skrzypacz, A. (2005). Uncertainty about uncertainty and delay in bargaining. *Econometrica*, 73(1):69–91.

Fershtman, C. and Seidmann, D. J. (1993). Deadline effects and inefficient delay in bargaining with endogenous commitment. *Journal of Economic Theory*, 60(2):306–321.

Friedenberg, A. (2019). Bargaining under strategic uncertainty: The role of second-order optimism. *Econometrica*, 87(6):1835–1865.

Fuchs, W. and Skrzypacz, A. (2010). Bargaining with arrival of new traders. *American Economic Review*, 100(3):802–836.

Fuchs, W. and Skrzypacz, A. (2013). Bargaining with deadlines and private information. *American Economic*

Journal: Microeconomics, 5(4):219–243.

Garoupa, N. and Rizzolli, M. (2011). The Brady rule may hurt the innocent. *American Law and Economics Review*, 13(1):168–200.

Garrett, B. L. (2017). Actual innocence and wrongful convictions. *Academy for Justice, A Report on Scholarship and Criminal Justice Reform*.

Grossman, J. M. (2016). Getting Brady right: Why extending Brady v. Maryland's right to plea negotiations better protects a defendant's constitutional rights in the modern legal era. *BYU Law Review*, 2016:1525.

Grossman, G. M. and Katz, M. L. (1983). Plea bargaining and social welfare. *American Economic Review*, 73(4):749–757.

Grossman, S. J. and Perry, M. (1986). Perfect sequential equilibrium. *Journal of Economic Theory*, 39(1):97–119.

Guha, B. (2024). Plea bargaining when juror effort is costly. *Economic Theory*, 78(3):945–977.

Güth, W., Levati, M. V., and Maciejovsky, B. (2005). Deadline effects in sequential bargaining: An experimental study. *International Game Theory Review*, 7(2):117–135.

Hwang, I. (2018). A theory of bargaining deadlock. *Games and Economic Behavior*, 109:501–522.

Hwang, I. and Li, F. (2017). Transparency of outside options in bargaining. *Journal of Economic Theory*, 167:116–147.

Ispano, A. and Vida, P. (2024). Designing interrogations. *Review of Economic Studies*, 91(6):3504–3531.

Keenan, D., Cooper, D. J., Lebowitz, D., and Lerer, T. (2011). The myth of prosecutorial accountability after Connick v. Thompson: Why existing professional responsibility measures cannot protect against prosecutorial misconduct. *Yale Law Journal Forum*, 121:203.

Kim, J.-Y. (2025). Credible plea bargaining with the jury's rational conviction decision. *European Journal of Law and Economics*, 59(3):491–521.

Kolb, A. and Zhou, B. (2026). Dynamic disclosure with(out) timestamps. Working paper, SSRN 7499918.

Landes, W. M. (1971). An economic analysis of the courts. *Journal of Law and Economics*, 14(1):61–107.

Little, A. T. and Simpson, H. K. (2025). Guilt and guilty pleas. *American Political Science Review*, 119(2):1003–1017.

Lomys, N. (2017). Learning while trading: Experimentation and Coasean dynamics. Working paper.

Lundberg, A. (2024). Do prosecutors induce the innocent to plead guilty? *Economic Inquiry*, 62(2):650–674.

Ma, C.-t. A. and Manove, M. (1993). Bargaining with deadlines and imperfect player control. *Econometrica*, 61(6):1313–1339.

Mailath, G. J., Okuno-Fujiwara, M., and Postlewaite, A. (1993). Belief-based refinements in signalling games. *Journal of Economic Theory*, 60(2):241–276.

Medwed, D. S. (2004). The zeal deal: Prosecutorial resistance to post-conviction claims of innocence. *Boston University Law Review*, 84:125.

Ortner, J. (2017). Durable goods monopoly with stochastic costs. *Theoretical Economics*, 12(2):817–861.

Ortner, J. (2023). Bargaining with evolving private information. *Theoretical Economics*, 18(3):885–916.

Petegorsky, M. N. (2012). Plea bargaining in the dark: The duty to disclose exculpatory Brady evidence during plea bargaining. *Fordham Law Review*, 81:3599.

Pfaff, J. (2017). Prosecutorial guidelines. *Academy for Justice: A Report on Scholarship and Criminal Justice Reform*, 3.

Ralston, J., Aimone, J., Rentschler, L., and North, C. (2023). Prosecutor plea bargaining and conviction rate structure: Evidence from an experiment. *Public Choice*, 196:299–329.

Reinganum, J. F. (1988). Plea bargaining and prosecutorial discretion. *American Economic Review*, 78(4):713–728.

Roth, A. E., Murnighan, J. K., and Schoumaker, F. (1988). The deadline effect in bargaining: Some experimental evidence. *American Economic Review*, 78(4):806–823.

Sanders, B. (2019). Exculpatory evidence pre-plea without extending Brady. *University of Chicago Law Review*, 86(8):2243–2300.

Shishkin, D. (2022). Evidence acquisition and voluntary disclosure. Working paper.

Simon, L. K. and Stinchcombe, M. B. (1989). Extensive form games in continuous time: Pure strategies. *Econometrica*, 57(5):1171–1214.

Spier, K. E. (1992). The dynamics of pretrial negotiation. *Review of Economic Studies*, 59(1):93–108.

Vasserman, S. and Yildiz, M. (2019). Pretrial negotiations under optimism. *RAND Journal of Economics*, 50(2):359–390.

A Extensions

This appendix extends the continuous-time analysis in four directions: a prosecutor who cares about convicting the innocent, a trial that sentences according to the posterior probability of guilt, an investigation that is the prosecutor's choice, and incriminating evidence that is not conclusive. Throughout, $\tilde{\theta} = \tilde{\theta}(\mu)$ and $\bar{\theta} = \bar{\theta}(\mu)$ are as in the main text, and the solution concept is the one in Section 2.

A.1 A prosecutor who cares about convicting the innocent

Suppose the prosecutor incurs a cost $k \geq 0$ whenever an innocent defendant receives a positive sentence, whether by agreement or at trial. Define

$$A_k \equiv q(1-\mu)c - (1-q)k, \qquad \bar{\theta}(\mu, k) \equiv \max\left\{0,\ \frac{A_k}{A_k + qd(1-q)}\right\}.$$

Note that $\bar{\theta}(\mu, 0) = \bar{\theta}(\mu)$ and that $\bar{\theta}(\mu, k)$ is decreasing in k.

Proposition A.1 (Concern for the innocent)**.** *(i) If $k < dq$, the selected consistent equilibrium is regime C if $\theta \leq \tilde{\theta}$, regime S if $\tilde{\theta} < \theta \leq \bar{\theta}(\mu, k)$, and regime D if $\theta > \max\{\tilde{\theta}, \bar{\theta}(\mu, k)\}$. In particular, exculpatory evidence is disclosed if and only if $\theta > \max\{\tilde{\theta}, \bar{\theta}(\mu, k)\}$.*

(ii) If $k \geq d$, a prosecutor with exculpatory evidence discloses it upon arrival for every θ.

The deadline cutoff $\tilde{\theta}$ does not depend on k: at the deadline, an innocent defendant facing a d-type receives a positive sentence whether he accepts dq or d or goes to trial, so the cost k is incurred in both cases and cancels. The ex ante cutoff does depend on k. Under disclosure, an innocent defendant goes free whenever exculpatory evidence arrives, whereas under concealment he always receives the positive sentence dq. A prosecutor who cares about the innocent therefore values disclosure more, the screening region shrinks, and when k is large enough to make $\bar{\theta}(\mu, k) \leq \tilde{\theta}$ it disappears. When the concern is strong enough ($k \geq d$), concealment never pays. The main result, that a prosecutor discloses exculpatory evidence when the prior of guilt is high, thus holds a fortiori.

Proof. (i) Under k, the continuation values of Lemma 1 become $V_k^C(t;\theta) = V^C(t;\theta) - (1-\theta)k$ and $V_k^D(t;\theta) = V^D(t;\theta) - (1-\theta)r(t)k$, and the ex ante value of regime S is $V_k^C(0;\theta)$.

Regime C. An e-type obtains $dq - k > 0$ from concealment and 0 from disclosure, so she conceals. Before the deadline, an offer accepted by both defendant types is at most dq and yields $dq - (1-\theta)k \leq V_k^C(t;\theta)$. At the deadline, the comparison between $x = dq$ and $x = d$ is

$$dq - (1-\theta)k \ \text{ versus } \ d - (1-\theta)(1-\mu)c - (1-\theta)k,$$

so k cancels and the argument of Lemma 3 applies: regime C is consistent if and only if $\theta \leq \tilde{\theta}$.

Regime S. The proof of Lemma 4 applies verbatim for $\theta > \tilde{\theta}$, with e-types concealing because $dq - k > 0$.

Regime D. An e-type who mimics the d-type obtains $\mu(d-k)-(1-\mu)c \le \mu d-(1-\mu)c \le 0$. As in Lemma 5, the only candidate credible deviation is the pooling offer $x = dq$ with $K = \{d, e\}$: the e-type gains because $dq - k > 0$, and the d-type gains if and only if $dq - (1-\theta)k > V_k^D(t;\theta)$. (With $K = \{d\}$, any accepted offer $x > k$ makes the e-type gain, and any accepted offer $x \le k < dq$ yields the d-type less than the pooling offer.) The function

$$F(r) \equiv V_k^D(t;\theta) - \big[dq - (1-\theta)k\big]$$

is affine in $r = r(t) \in [q, 1]$. At $r = 1$, $F(1) = V_d - dq$, which is nonnegative if and only if $\theta \ge \tilde{\theta}$. At $r = q$, $F(q) = V_k^D(0;\theta) - dq + (1-\theta)k$, which exceeds $V_k^C(0;\theta) - dq + (1-\theta)k = V^C(0;\theta) - dq \ge 0$ whenever $V_k^D(0;\theta) > V_k^C(0;\theta)$. Hence regime D is consistent whenever $\theta > \tilde{\theta}$ and $V_k^D(0;\theta) > V_k^C(0;\theta)$, and it is not consistent if $\theta < \tilde{\theta}$.

Selection. Direct computation gives

$$V_k^C(0;\theta) - V_k^D(0;\theta) = -\theta\, qd(1-q) + (1-\theta)\big[q(1-\mu)c - (1-q)k\big],$$

which is nonnegative if and only if $\theta \le \bar{\theta}(\mu, k)$. If $\theta \le \tilde{\theta}$, only regime C is consistent. If $\theta > \tilde{\theta}$, regimes S and (when $\theta > \bar{\theta}(\mu, k)$) D are consistent, and the prosecutor prefers S if and only if $\theta \le \bar{\theta}(\mu, k)$.

(ii) The innocent defendant can guarantee himself a loss of at most d by rejecting every offer, so he never pays more than d. Any outcome in which an e-type obtains a positive sentence therefore gives her at most $d - k \le 0$, a trial gives her $-c < 0$, and disclosure followed by the offer 0 gives her 0. By selection rule (i), she discloses. □

A.2 A Bayesian trial

Suppose there is no separate default sentence. Absent conclusive evidence, the trial assigns the sentence $\theta^T h$, where θ^T is the posterior probability of guilt given the evidence presented at trial. Plea negotiations are not admissible at trial (Federal Rule of Evidence 410), so the court does not learn from rejected offers. Because the arrival of evidence is independent of the defendant's type, the absence of evidence is uninformative, and $\theta^T = \theta$. The model is therefore the baseline model with d replaced by θh. To ensure that Assumption 1 holds for every θ, assume:

Assumption A.1. $\mu \le c/(c+h)$ and $(1-\mu)c > h(1-q)$.

Define

$$\tilde{\theta}_B \equiv \frac{(1-\mu)c}{(1-\mu)c + h(1-q)}, \qquad \bar{\theta}_B \equiv \frac{-(1-\mu)c + \sqrt{(1-\mu)^2c^2 + 4h(1-q)(1-\mu)c}}{2h(1-q)}.$$

Proposition A.2 (Bayesian trial)**.** *Under Assumption A.1, $\tilde{\theta}_B < \bar{\theta}_B$, and the selected consistent equilibrium is regime C if $\theta \le \tilde{\theta}_B$, regime S if $\tilde{\theta}_B < \theta \le \bar{\theta}_B$, and regime D if $\theta > \bar{\theta}_B$, with $d = \theta h$ throughout. In particular, exculpatory evidence is disclosed if and only if $\theta > \bar{\theta}_B$.*

The qualitative result is unchanged. The only difference is that the default sentence now rises with the prior, which makes the high offer more attractive to the d-type and the cutoffs depend on θ in a nonlinear way. The inadmissibility of plea discussions matters: if the court could learn from the negotiation, screening would lower the trial sentence and the analysis would require a fixed point.

Proof. With $d = \theta h$, Assumption A.1 implies $\mu \leq c/(c+d)$ and $(1-\mu)c > d(1-q)$, so Proposition 2 applies for each θ. The condition $\theta \leq \tilde{\theta}(\mu)$ becomes $\theta \leq 1 - \theta h(1-q)/((1-\mu)c)$, which is equivalent to $\theta \leq \tilde{\theta}_B$. The condition $\theta \leq \bar{\theta}(\mu)$ becomes $\theta[(1-\mu)c + \theta h(1-q)] \leq (1-\mu)c$, that is, $Q(\theta) \equiv h(1-q)\theta^2 + (1-\mu)c\,\theta - (1-\mu)c \leq 0$. Since Q is convex with $Q(0) < 0$, this holds if and only if $\theta \leq \bar{\theta}_B$, the positive root. Finally, $Q(\tilde{\theta}_B) = \tilde{\theta}_B[h(1-q)\tilde{\theta}_B + (1-\mu)c] - (1-\mu)c < \tilde{\theta}_B[h(1-q) + (1-\mu)c] - (1-\mu)c = 0$, so $\tilde{\theta}_B < \bar{\theta}_B$. □

A.3 Investigation as a decision

In the main text, investigation is automatic. Suppose instead that at each t the prosecutor chooses whether to investigate, and that evidence arrives at rate λ only while she investigates. Under *public* investigation the defendant observes this choice; under *private* investigation he does not. The consistency requirement of Definition 1 applies to every off-path action of the prosecutor at a history without disclosure, including a publicly observed decision to stop investigating. Define

$$\theta_P \equiv \frac{d}{h - dq},$$

and maintain:

Assumption A.2. $h \geq dq + d/\tilde{\theta}(\mu)$, equivalently $\theta_P \leq \tilde{\theta}(\mu)$.

Assumption A.2 says that incriminating evidence is valuable enough for investigation to be attractive throughout the concealment and screening regions. It implies $d/h < \theta_P \leq \tilde{\theta}(\mu) < \bar{\theta}(\mu)$. Without it, under public investigation, a prosecutor in the concealment regime could gain by publicly stopping her investigation shortly after it begins, and the characterization below would change for intermediate priors.

Proposition A.3 (Investigation under voluntary disclosure)**.** *Suppose Assumption A.2 holds. Under public investigation, and under private investigation when the prosecutor believes that a defendant who rejects an offer that both types accept on the equilibrium path is innocent, the selected consistent equilibrium is:*

(i) if $\theta \leq \theta_P$, the prosecutor offers d at $t = 0$, both defendant types accept, and she never investigates;

(ii) if $\theta > \theta_P$, the prosecutor investigates until agreement, and the outcome is as in Proposition 2.

Proposition A.4 (Investigation under mandatory disclosure)**.** *Under mandatory disclosure, with public or private investigation, the prosecutor offers d at $t = 0$, both defendant types accept, and she never investigates if $\theta \leq d/h$. If $\theta > d/h$, she investigates until agreement and the outcome is as in Lemma 6.*

At $t = 0$ no investigation has taken place, so there is no second-order uncertainty yet: if a rejection leads to no investigation, both defendant types expect to pay d and accept it. The prosecutor therefore compares settling immediately at d with the value of investigating. Investigation pays when guilt is likely enough, because it may produce incriminating evidence. Under voluntary disclosure, however, investigation is worth less, because it creates the second-order uncertainty that forces the prosecutor holding default evidence to settle with a guilty defendant at dq rather than d. As a result, the prosecutor gives up investigating for a larger set of priors ($\theta \leq \theta_P$) than under mandatory disclosure ($\theta \leq d/h$). Under public investigation, this outcome does not depend on the prosecutor's belief about a defendant who rejects the immediate offer. Under private investigation, it requires the belief that such a defendant is innocent; under passive beliefs the cutoff falls to $(d - (1-\mu)c)/(h - (1-\mu)c)$ when $d > (1-\mu)c$.

Corollary A.1 (Welfare with an investigation decision)**.** *For $\theta \leq d/h$, outcomes under voluntary and mandatory disclosure coincide. For $\theta \in (d/h, \theta_P]$:*

(i) the innocent defendant pays d under voluntary disclosure and dq under mandatory disclosure;

(ii) the guilty defendant pays d under voluntary disclosure and $(1-q)h + qd$ under mandatory disclosure;

(iii) the prosecutor obtains d under voluntary disclosure and $V^M(\theta) > d$ under mandatory disclosure.

For $\theta > \theta_P$, Propositions 3 and 4 apply.

When investigation is a choice, mandatory disclosure strictly benefits innocent defendants for intermediate priors: it restores the prosecutor's incentive to investigate, and investigation frees innocent defendants with probability $1 - q$. The prosecutor still prefers mandatory disclosure, weakly for $\theta \leq d/h$ and strictly otherwise.

Proof of Proposition A.3. *Step 1: incentives to investigate in regimes C, S, and D.* Fix the defendants' strategies. Under private investigation, the prosecutor's choice affects her payoff only through the probability p that evidence arrives before agreement. In regime C (and in the continuation of regime S), her payoff is $\theta[ph + (1-p)dq] + (1-\theta)dq$, which is increasing in p. In regime D, it is

$$\theta\big[ph + (1-p)d\big] + (1-\theta)(1-p)\big(d - (1-\mu)c\big),$$

whose slope in p is $\theta(h-d) - (1-\theta)(d-(1-\mu)c)$. This is positive whenever $\theta > \theta_I \equiv (d-(1-\mu)c)/(h-(1-\mu)c)$. Since $\theta_I < d/h < \theta_P < \bar{\theta} < \theta$ in regime D (the first inequality because $h > d$), the prosecutor investigates at every t.

Under public investigation, a decision to stop investigating is observed. In regime D, the innocent defendant believes the prosecutor holds d on path, so stopping changes nothing except p, and the argument above applies. In regime C, consider a public decision to stop at time t, and let $u = e^{-\lambda t} \in [q, 1]$. If the innocent defendant attributed it to the pool $\{d, e\}$, his belief would be u, he would accept up to du, and a d-type would gain if and only if $du > V^C(t;\theta)$. Write $V^C(t;\theta) - du = G(u)$ with

$$G(u) = \theta h - \theta \frac{q}{u}(h - dq) + (1-\theta)dq - du.$$

G is concave, $G(q) = 0$, and $G(1) = V^C(0;\theta) - d \geq 0$ when $\theta \geq \theta_P$ (see Step 3). In regime C we have $\theta > \theta_P$, and in the continuation of regime S the belief is $\tilde{\theta} \geq \theta_P$ by Assumption A.2. Hence $G \geq 0$ on $[q, 1]$, the d-type does not gain, and $\{d, e\}$ is not self-consistent. If $K = \{d\}$, the innocent defendant would accept d and the e-type would gain, so $\{d\}$ is not self-consistent; $K = \{e\}$ yields no gain. Assigning belief 0 after a public stop, the innocent defendant rejects every positive offer, so stopping yields the d-type at most $\theta d + (1-\theta)(d-c) \leq dq \leq V^C(t;\theta)$ in regime C (where $\theta \leq \tilde{\theta}$), and it yields the e-type at most $0 < dq$. The same argument applies to the continuation of regime S, where the belief is $\tilde{\theta}$.

Step 2: immediate settlement at d. Consider the profile in which the prosecutor offers d at $t = 0$ and both defendant types accept. After a rejection, specify the continuation in which the prosecutor never investigates and offers d at the next offer date. Under public investigation with $\theta \leq \theta_P$, this continuation is optimal for the prosecutor whatever her belief about the defendant, since by Step 3 the value of investigating does not exceed d; both defendant types then expect to pay d and are willing to accept d at $t = 0$. Under private investigation, after a rejection the prosecutor believes the defendant is innocent. Investigating can then only produce exculpatory evidence; the defendant, expecting no investigation, believes the prosecutor holds d and accepts d, so an e-type that conceals obtains d as well. The prosecutor is therefore indifferent about investigating, the continuation is an equilibrium, and both defendant types are willing to accept d at $t = 0$. At $t = 0$ no evidence can have arrived, so the innocent defendant's belief that the prosecutor holds d is 1 after any offer, his continuation loss is d, and no offer above d is accepted. No off-path offer admits a self-consistent set.

Step 3: selection. We have

$$V^C(0;\theta) - d = (1-q)\big[\theta(h - dq) - d\big],$$

which is nonnegative if and only if $\theta \geq \theta_P$. For $\theta \leq \theta_P \leq \bar{\theta}$, Lemma 2 gives $V^D(0;\theta) \leq V^C(0;\theta) \leq d$, so immediate settlement is selected. For $\theta > \theta_P$, $\max\{V^C(0;\theta), V^D(0;\theta)\} > d$, and by Step 1 the regimes of Proposition 2 remain consistent equilibria with continuous investigation, so the outcome of Proposition 2 is selected. □

Proof of Proposition A.4. Under mandatory disclosure, the prosecutor's payoff from an investigation policy with arrival probability p is

$$\theta\big[ph + (1-p)d\big] + (1-\theta)(1-p)d = d + p(\theta h - d),$$

because every case settles at the trial sentence. If $\theta \leq d/h$, this is maximized at $p = 0$: the prosecutor does not investigate, both defendant types know her evidence is d, and they settle at d. Since no evidence will arrive, the timing of this agreement is payoff-irrelevant, and I report agreement at $t = 0$. If $\theta > d/h$, it is maximized by investigating at every t, and the outcome is that of Lemma 6. Since evidence is disclosed upon arrival, the defendant's beliefs do not depend on whether investigation is observed. Immediate settlement can also be supported for $\theta > d/h$ under private investigation, but it gives the prosecutor $d < V^M(\theta)$ and is not selected. □

Proof of Corollary A.1. For $\theta \leq d/h$, both regimes settle at d at $t = 0$. For $\theta \in (d/h, \theta_P]$, voluntary disclosure settles at d (Proposition A.3), while mandatory disclosure leads to investigation and the payoffs of Lemma 6 and Proposition 3; moreover $V^M(\theta) - d = \theta(1-q)h + qd - d = (1-q)(\theta h - d) > 0$. For $\theta > \theta_P$, both regimes involve investigation and the comparisons of Section 6 apply. □

A.4 Inconclusive incriminating evidence

In the main text, evidence is conclusive. Here I let incriminating evidence be inconclusive, while exculpatory evidence remains conclusive (as with a DNA exclusion), and I allow the probability of finding evidence to depend on the defendant's type. I work with the one-offer benchmark of Section 3. Evidence is found by $t = 1$ with probability $1 - q_G$ if the defendant is guilty and $1 - q_I$ if he is innocent, where $q_G, q_I \in (0, 1)$. A guilty defendant's evidence is always h; an innocent defendant's evidence is h with probability $\pi \in [0, 1)$ and e with probability $1 - \pi$. The trial assigns the sentences $s(e) = 0$, $s(d) = d$, and $s(h) = h$ as before, so an innocent defendant can be convicted on incriminating evidence.

Because the probability of finding evidence differs across types, the absence of evidence is now informative. A prosecutor who finds nothing updates her belief to

$$\theta^d \equiv \frac{\theta q_G}{\theta q_G + (1-\theta)q_I}.$$

If exculpatory evidence is concealed, the innocent defendant's belief that the prosecutor holds default evidence is

$$\rho \equiv \frac{q_I}{q_I + (1-q_I)(1-\pi)}.$$

Define

$$\tilde{\theta}_{NC} \equiv 1 - \frac{d(1-\rho)}{(1-\mu)c}, \qquad \theta^*_{NC} \equiv \frac{q_I \tilde{\theta}_{NC}}{q_I \tilde{\theta}_{NC} + q_G(1 - \tilde{\theta}_{NC})},$$

and assume $\mu \leq c/(c+d)$ and $(1-\mu)c > d(1-\rho)$, so that $\tilde{\theta}_{NC} \in (0,1)$. Note that $\theta^d \leq \tilde{\theta}_{NC}$ if and only if $\theta \leq \theta^*_{NC}$.

Proposition A.5 (Inconclusive incriminating evidence)**.** *In the one-offer benchmark with inconclusive incriminating evidence, the consistent equilibrium outcome is as follows. A prosecutor with incriminating evidence discloses it and offers h, which both defendant types accept.*

(i) If $\theta \leq \theta^*_{NC}$, exculpatory evidence is concealed; the d-type and the e-type offer $d\rho$, which both defendant types accept.

(ii) If $\theta > \theta^*_{NC}$, exculpatory evidence is disclosed and the e-type offers 0; the d-type offers d, which the guilty defendant accepts and the innocent defendant accepts with probability μ.

Corollary A.2 (Welfare with inconclusive incriminating evidence)**.** *(i) The innocent defendant's expected loss is $(1-q_I)\pi h + q_I d$ under voluntary disclosure (in both regimes) and under mandatory disclosure.*

(ii) The guilty defendant's expected loss is $(1-q_G)h + q_G d$ under mandatory disclosure. Under voluntary disclosure it is $(1-q_G)h + q_G d\rho$ if $\theta \leq \theta^*_{NC}$ and $(1-q_G)h + q_G d$ otherwise.

(iii) Under voluntary disclosure, the case goes to trial with probability $(1-\theta)q_I(1-\mu)$ if $\theta > \theta^*_{NC}$ and never otherwise; under mandatory disclosure it never goes to trial.

(iv) The prosecutor's payoff is strictly higher under mandatory disclosure, by $\theta q_G d(1-\rho)$ if $\theta \leq \theta^*_{NC}$ and by $(1-\theta)q_I(1-\mu)c$ otherwise.

The structure of the baseline survives. The innocent defendant's second-order uncertainty now reflects both the probability of finding evidence and the probability that the evidence found is exculpatory, and the prosecutor's decision depends on her posterior after finding nothing. Inconclusive incriminating evidence exposes innocent defendants to the high sentence with probability $(1-q_I)\pi$, but that risk is the same under both disclosure regimes, so the welfare comparisons of Section 6 carry over unchanged. If exculpatory evidence were also inconclusive, guilty defendants would face second-order uncertainty as well; I leave that case, and the continuous-time version of this extension, for future work.

Proof of Proposition A.5. An h-type who discloses and offers h obtains h from either defendant type, whose trial loss is then h; this is the highest payoff available, so by selection rule (ii) she discloses. A guilty defendant never faces an e-type, so after no disclosure he believes the prosecutor holds d and accepts x if and only if $x \leq d$. If e-types conceal, Bayes' rule gives the innocent defendant the belief ρ, so his trial loss is $d\rho$, and $\beta_{\{d,e\}} = \rho$. A d-type holds belief θ^d; an e-type knows the defendant is innocent. The proof of Proposition 1 now applies with q replaced by ρ in the innocent defendant's beliefs and with θ replaced by θ^d in the d-type's comparison. In particular, the d-type prefers the pooling offer $d\rho$ to the offer d if and only if

$$d\rho \;\geq\; d - (1-\theta^d)(1-\mu)c \iff \theta^d \leq \tilde{\theta}_{NC},$$

and the credible deviations that rule out each regime are the same as in that proof. Finally, $\theta^d \leq \tilde{\theta}_{NC}$ if and only if $\theta q_G(1-\tilde{\theta}_{NC}) \leq (1-\theta)q_I\tilde{\theta}_{NC}$, that is, $\theta \leq \theta^*_{NC}$. □

Proof of Corollary A.2. (i) Under concealment, the innocent defendant pays h with probability $(1-q_I)\pi$ and $d\rho$ otherwise; since $[q_I+(1-q_I)(1-\pi)]\rho = q_I$, his expected loss is $(1-q_I)\pi h+q_I d$. Under disclosure and under mandatory disclosure he pays h with probability $(1-q_I)\pi$, 0 with probability $(1-q_I)(1-\pi)$, and d with probability q_I. (ii) The guilty defendant pays h with probability $1-q_G$ and, otherwise, $d\rho$ under concealment and d under disclosure or mandatory disclosure. (iii) Under disclosure, an innocent defendant without new evidence rejects d with probability $1-\mu$. (iv) Subtract the prosecutor's payoffs, which equal the defendants' expected losses minus the expected trial cost. □

B Proofs

Throughout, write $\tilde{\theta} = \tilde{\theta}(\mu)$, $\bar{\theta} = \bar{\theta}(\mu)$, and $V_d \equiv d-(1-\theta)(1-\mu)c$. Two facts are used repeatedly:

$$V_d > dq \iff \theta > \tilde{\theta}, \qquad \mu d-(1-\mu)c \leq 0 \iff \mu \leq \bar{\mu}. \tag{1}$$

The first follows from $d-(1-\theta)(1-\mu)c > dq \iff (1-\theta)(1-\mu)c < d(1-q) \iff \theta > \tilde{\theta}$.

Preliminary: incriminating evidence. Since the guilty defendant's loss at trial is at most h, he rejects any offer above h; hence h is an upper bound on the prosecutor's payoff from a guilty defendant. An h-type who discloses and offers h obtains exactly h (the guilty defendant's trial loss after disclosure is h). By selection rule (ii), she does so upon arrival. Consequently, at any history without disclosure, the guilty defendant believes the prosecutor is d-type, and his responses do not depend on the innocent defendant's beliefs.

Proof of Proposition 1

With a single offer at $t=1$, a prosecutor without disclosure is a d-type with probability one from the guilty defendant's point of view, so the guilty defendant accepts x if and only if $x \leq d$. The innocent defendant's trial loss is d against a d-type and 0 against an e-type, so $\ell_d = d$ and $\ell_e = 0$, and $\beta_{\{d,e\}} = q$ (all e-types in the pool).

Concealment. Let the d- and e-types both offer dq without disclosure. On path the innocent defendant's belief is q, his trial loss is dq, and he accepts. Consider deviations.

- $x < dq$: accepted by both defendant types; lower payoff for both prosecutor types.
- $x > d$: rejected by the guilty defendant. A d-type gets at most $d-c < dq$ (by Assumption 1, $c > d(1-q)$), and an e-type gets at most $0 < dq$.

- $x \in (dq, d]$: consider candidate sets K. If $K = \{d, e\}$, the belief is q, the innocent defendant rejects ($dq < x$), and the e-type obtains $-c < dq$, so $e \notin$ gainers and K is not self-consistent. If $K = \{e\}$, the innocent defendant rejects and no type gains. If $K = \{d\}$ and $x < d$, the innocent defendant accepts, so the e-type would gain ($x > dq$) and K is not self-consistent. If $K = \{d\}$ and $x = d$, the innocent defendant accepts with probability μ; the e-type obtains $\mu d - (1-\mu)c \leq 0 < dq$ and does not gain, while the d-type obtains $\theta d + (1-\theta)[\mu d + (1-\mu)(d-c)] = V_d$, which exceeds dq if and only if $\theta > \tilde{\theta}$ by (1). Hence $\{d\}$ is self-consistent for $x = d$ if and only if $\theta > \tilde{\theta}$.

Thus for $\theta > \tilde{\theta}$ concealment is not consistent. For $\theta \leq \tilde{\theta}$, no off-path offer admits a self-consistent set; assign belief q after every $x > dq$, so the innocent defendant rejects. A d-type's best deviation then yields $\theta d + (1-\theta)(d-c) \leq V_d \leq dq$, and an e-type's yields $-c$. Disclosure gives an e-type $0 < dq$. So concealment is a consistent equilibrium if and only if $\theta \leq \tilde{\theta}$.

Disclosure. Let the e-type disclose and offer 0, and the d-type offer d. On path, after no disclosure the innocent defendant's belief is 1; he is indifferent and accepts with probability μ. The e-type's best concealment strategy that mimics the d-type yields $\mu d - (1-\mu)c \leq 0$ by (1), so she weakly prefers to disclose, and selection rule (i) applies. Consider off-path offers without disclosure.

- $x \in (0, d)$: if $K = \{d, e\}$, the innocent defendant accepts if and only if $x \leq dq$. The d-type gains only if $x > V_d$, which is impossible for $x \leq dq$ when $\theta > \tilde{\theta}$; for $x \in (dq, d)$ the offer is rejected and no type gains. If $K = \{d\}$, the innocent defendant accepts, so the e-type gains and K is not self-consistent. If $K = \{e\}$, the offer is rejected and no type gains. Hence, for $\theta > \tilde{\theta}$, no self-consistent set exists. Assign belief 0, so the innocent defendant rejects any $x > 0$. A d-type's deviation yields $\theta x + (1-\theta)(d-c) < V_d$, and an e-type's yields $-c < 0$.

- $x > d$: rejected by both defendant types; worse for both prosecutor types.

- $x \leq 0$: yields at most $0 < V_d$ to the d-type.

For $\theta < \tilde{\theta}$ the disclosure profile is not consistent: for $x = dq$, the set $K = \{d, e\}$ is self-consistent, because the innocent defendant accepts dq given belief q, the e-type gains ($dq > 0$), and the d-type gains ($dq > V_d$). Combining both regimes yields the proposition; at $\theta = \tilde{\theta}$ the two regimes give the prosecutor the same payoff and I report concealment. □

Proof of Lemma 1

(i) In regime D, if the defendant is innocent, an e-type arrives in $(t, 1]$ with probability $1 - r$ and the innocent defendant then pays 0; otherwise he pays d at $t = 1$, either by accepting or at trial. If he is guilty, an h-type arrives with probability $1 - r$ and he pays h; otherwise he accepts d at $t = 1$. The prosecutor obtains d from an innocent defendant at $t = 1$ with probability μ and $d - c$ with probability $1 - \mu$, which gives V^D.

(ii) In regime C, an innocent defendant pays dq at $t = 1$ regardless of arrival, since the e-type pools with the d-type. A guilty defendant pays h if h arrives in $(t, 1]$ and dq otherwise.

(iii) In regime S, the guilty defendant pays x_0 at $t = 0$ with probability $s(\theta)$; otherwise he continues in regime C, where by (ii) at $t = 0$ his expected loss is $(1 - q)h + q\,dq = x_0$. Either way he pays x_0 in expectation, and the innocent defendant pays dq. Hence the ex ante value is $\theta x_0 + (1 - \theta)dq = V^C(0; \theta)$. □

Proof of Lemma 2

Let $a \equiv d(1 - q)/((1 - \mu)c) \in (0, 1)$ by Assumption 1. Then $\tilde{\theta} = 1 - a$ and $\bar{\theta} = 1/(1 + a)$, so $0 < \tilde{\theta} < \bar{\theta} < 1$ because $(1 - a)(1 + a) = 1 - a^2 < 1$. Next,

$$V^C(0; \theta) - V^D(0; \theta) = \theta\, q\,(dq - d) + (1 - \theta)\big[dq - qd + q(1 - \mu)c\big] = -\theta\, qd(1 - q) + (1 - \theta)\, q(1 - \mu)c,$$

which is nonnegative if and only if $\theta \leq (1 - \mu)c/\big((1 - \mu)c + d(1 - q)\big) = \bar{\theta}$. □

Proof of Lemma 3

Fix regime C with prior θ. The guilty defendant's belief is d-type at histories without disclosure. On path, conditional on innocence and no disclosure by t, the innocent defendant's belief is $e^{-\lambda t}$, which equals q at $t = 1$; his trial loss there is dq, and he accepts dq.

Before the deadline. In regime C, both a d-type and an e-type end with the pooled offer dq, so $\ell_d(t) = \ell_e(t) = dq$ and the innocent defendant accepts an offer before $t = 1$ if and only if it is at most dq, whatever his belief. No set K is self-consistent for any off-path offer before $t = 1$, because the innocent defendant's response does not depend on β_K. An offer accepted by both types yields at most $dq \leq V^C(t; \theta)$, since $(1 - r)h + r\,dq \geq dq$. An offer accepted only by the guilty defendant must be at most his continuation loss $(1 - r)h + r\,dq$; if he accepts with probability one, a rejection reveals innocence and the innocent defendant continues to pay dq, so the prosecutor obtains at most $V^C(t; \theta)$. If he accepts with probability less than one, he must be indifferent, so the offer equals his continuation loss after a rejection. After a rejection the prosecutor's belief falls below $\theta \leq \tilde{\theta}$, play continues in regime C, and the guilty defendant's continuation loss is again $(1 - r)h + r\,dq$. The prosecutor therefore obtains $\theta[(1 - r)h + r\,dq] + (1 - \theta)dq = V^C(t; \theta)$, no more than by following the equilibrium. Hence no deviation before $t = 1$ is profitable.

At the deadline. Conditional on no arrival, the d-type's belief at $t = 1$ is θ. Here $\ell_d = d$ and $\ell_e = 0$, and $\beta_{\{d,e\}}(1) = q$. The analysis of offers at $t = 1$ is identical to the concealment part of the proof of Proposition 1, with the e-type's equilibrium payoff dq. Therefore regime C admits a self-consistent deviation (the offer d with $K = \{d\}$) if and only if $\theta > \tilde{\theta}$, and otherwise is a PBE with the off-path beliefs specified there. □

Proof of Lemma 4

Let $\theta > \tilde{\theta}$. At $t = 0$ no evidence has arrived, so every prosecutor is a d-type and $\beta_K(0) = 1$ for any K that contains d.

Defendants. Following a rejection of x_0, the prosecutor's posterior is

$$\frac{\theta(1-s)}{\theta(1-s)+1-\theta} = \tilde{\theta} \qquad \text{when } s = s(\theta),$$

and play continues in regime C with prior $\tilde{\theta}$, which is a consistent equilibrium by Lemma 3. The guilty defendant's loss in that continuation is $(1-q)h + q\,dq = x_0$, so he is indifferent and mixing with probability $s(\theta)$ is optimal. The innocent defendant's continuation loss is $dq < x_0$ (because $h > dq$), so he rejects.

Prosecutor at $t = 0$. Specify that, after any offer at $t = 0$ other than x_0, the prosecutor makes the screening offer $x_{t'} = (1-r(t'))h + r(t')\,dq$ at some later date $t' > 0$ and play continues in regime C (the inertia restriction makes this well defined). Under this continuation, e-types arriving before t' conceal, the guilty defendant pays h if h arrives before t' and $x_{t'}$ in expectation otherwise, and the innocent defendant pays dq. Hence the prosecutor's ex ante payoff is

$$\theta\big[(1-e^{-\lambda t'})h + e^{-\lambda t'}x_{t'}\big] + (1-\theta)dq = \theta\big[(1-q)h + q\,dq\big] + (1-\theta)dq = V^C(0;\theta),$$

using $e^{-\lambda t'}r(t') = q$. Delaying screening is therefore not profitable, and by selection rule (iii) screening occurs at $t = 0$. An offer at $t = 0$ accepted by both defendant types is at most the innocent defendant's continuation loss $dq \le V^C(0;\theta)$. Since $\beta_K(0) = 1$ and $\ell_d(0) = dq$, the innocent defendant's response to any off-path offer at $t = 0$ is the same for every K, so no self-consistent set exists. After $t = 0$, the play is regime C with prior at most $\tilde{\theta}$, which is consistent by Lemma 3. □

Proof of Lemma 5

Fix regime D. The guilty defendant believes the prosecutor is d-type at histories without disclosure; the innocent defendant does too on path, because every e-type discloses. At a history at time t without disclosure, $\ell_d(t) = r(t)d$ by Lemma 1(i), and $\ell_e(t) = 0$ because an e-type's equilibrium strategy is to disclose.

Exculpatory types. An e-type that conceals can either disclose later (payoff 0), mimic the d-type until $t = 1$ (payoff $\mu d - (1-\mu)c \le 0$ by (1)), or make an off-path offer. So disclosing upon arrival is optimal provided off-path offers are rejected by the innocent defendant, which is established below.

Off-path offers without disclosure. Let $x \in (0,\infty)$ be an off-path offer at time t. If $x > \ell_G(t) \equiv (1-r)h + rd$, the guilty defendant rejects. Consider the candidate sets.

- $K = \{d\}$: then $\beta = 1$ and the innocent defendant accepts $x \le r(t)d$ (at $t = 1$ and $x = d$, with probability μ). For any accepted $x > 0$, an e-type would strictly gain relative to her equilibrium payoff of 0, so K is not self-consistent. For $x > r(t)d$ the offer is rejected by the innocent defendant and the e-type does not gain, so $K = \{d\}$ would require the d-type alone to gain; but a d-type whose offer is rejected by the innocent defendant obtains at most $\theta\ell_G(t)$ from a guilty defendant and, with belief $\beta = 1$, at most $r(t)\big(d - (1-\mu)c\big)$ from an innocent one, for a total of at most $V^D(t;\theta)$, so she does not gain either.

- $K = \{e\}$: then $\beta = 0$, the innocent defendant rejects every $x > 0$, and no type gains.
- $K = \{d, e\}$: then $\beta = e^{-\lambda t}$ and the innocent defendant accepts if and only if

$$x \leq e^{-\lambda t}\, r(t)\, d = dq.$$

  For $x \leq dq$, the e-type gains. The d-type obtains $x \leq dq$ from both defendant types, so she strictly gains if and only if $x > V^D(t;\theta)$. For $x > dq$, the innocent defendant rejects and the e-type does not gain.

Therefore a self-consistent set exists for some off-path offer if and only if $V^D(t;\theta) < dq$ for some t (take $x = dq$ and $K = \{d, e\}$). When $V^D(t;\theta) \geq dq$ for all t, assign belief 0 after every off-path offer without disclosure. Then the innocent defendant rejects every $x > 0$, so an e-type that conceals obtains at most 0, and a d-type that deviates obtains at most $\theta\ell_G(t) + (1-\theta)r(t)(d-c) \leq V^D(t;\theta)$ (if the guilty defendant accepts for sure, a rejection reveals innocence, and the innocent defendant, believing the prosecutor holds e, rejects all positive offers until the trial). Offers $x \leq 0$ yield at most $0 < dq \leq V^D$. Hence regime D is a consistent equilibrium if and only if $V^D(t;\theta) \geq dq$ for all $t \in [0,1]$.

Sufficiency of $\theta > \bar{\theta}$. $V^D(t;\theta)$ is affine in $r(t) \in [q,1]$, so $V^D \geq dq$ on $[0,1]$ if and only if it holds at $r = 1$ and $r = q$. At $r = 1$, $V^D(1;\theta) = V_d > dq$ because $\theta > \bar{\theta} > \tilde{\theta}$ (Lemma 2 and (1)). At $r = q$, Lemma 2 gives $V^D(0;\theta) > V^C(0;\theta)$ for $\theta > \bar{\theta}$, and $V^C(0;\theta) \geq dq$ because $(1-q)h + q\,dq \geq dq$. □

Proof of Proposition 2

By Lemmas 3–5, the consistent regimes are: C for $\theta \leq \tilde{\theta}$; S for $\theta > \tilde{\theta}$; and D at least for $\theta > \bar{\theta}$ (possibly also for some lower θ). By Lemma 1(iii), C and S both give the prosecutor $V^C(0;\theta)$, and D gives $V^D(0;\theta)$. By Lemma 2, $V^C(0;\theta) \geq V^D(0;\theta)$ if and only if $\theta \leq \bar{\theta}$. Selection rule (iv) therefore picks C for $\theta \leq \tilde{\theta}$, S for $\theta \in (\tilde{\theta}, \bar{\theta}]$, and D for $\theta > \bar{\theta}$, where D is available by Lemma 5. The disclosure statements follow from the definitions of the regimes and the preliminary step on incriminating evidence. □

Proof of Remark 1

Under belief $\beta = 1$, the innocent defendant accepts $r(t)d$ (it equals $\ell_d(t)$), and so does the guilty defendant, since $r(t)d < \ell_G(t)$. The d-type's gain is

$$rd - V^D(t;\theta) = rd - \theta\big[(1-r)h + rd\big] - (1-\theta)r\big(d - (1-\mu)c\big) = (1-\theta)r(1-\mu)c - \theta(1-r)h,$$

which converges to $(1-\theta)(1-\mu)c > 0$ as $t \to 1$. □

Proof of Corollary 1

In the benchmark, by Proposition 1 and the computations in Lemma 1, the prosecutor obtains $V^C(0;\theta)$ for $\theta \le \tilde{\theta}$ and $V^D(0;\theta)$ for $\theta > \tilde{\theta}$. In continuous time she obtains $\max\{V^C(0;\theta), V^D(0;\theta)\}$ (Proposition 2 and Lemma 2), which is weakly higher; on $(\tilde{\theta}, \bar{\theta})$ the benchmark gives $V^D(0;\theta) < V^C(0;\theta)$, so the inequality is strict there, and at $\theta = \bar{\theta}$ the two values coincide. The guilty defendant's loss is $(1-q)h + q\,dq$ when evidence is concealed and $(1-q)h + qd$ when it is disclosed, and concealment occurs on a larger set in continuous time. The innocent defendant's loss is dq in every regime. Trials occur only when e is disclosed. □

Proof of Lemma 6 and Corollary 2

After disclosure of e (h), the defendant's trial loss is 0 (h), so the prosecutor's best offer is 0 (h), accepted. Without new evidence at t, the innocent and guilty defendants' continuation losses are $r(t)d$ and $(1-r(t))h + r(t)d$, and the prosecutor's continuation value is $V^M(t) = \theta[(1-r)h + rd] + (1-\theta)rd$. An offer accepted by both defendant types is at most $r(t)d \le V^M(t)$; an offer accepted only by the guilty defendant gives at most $V^M(t)$. The first bound is strict for $t < 1$ and $\theta > 0$, since $V^M(t) - r(t)d = \theta(1-r(t))h$; the second holds with equality only for an offer that the guilty defendant accepts for sure, and such offers are excluded by selection rule (iii). Hence the prosecutor makes no serious offer before $t = 1$, where both defendant types know the evidence is d and accept d. No case goes to trial, and the ex ante payoff is $V^M(0)$. □

Proof of Proposition 3

(i) Under voluntary disclosure, the innocent defendant pays dq in regimes C and S (Lemma 1(ii)). In regime D he pays 0 with probability $1-q$ and d with probability q, which is again dq. Under mandatory disclosure he pays 0 with probability $1-q$ and d with probability q.

(ii) Under mandatory disclosure, the guilty defendant pays h with probability $1-q$ and d otherwise. Under voluntary disclosure, in regimes C and S he pays $(1-q)h + q\,dq$ in expectation (Lemma 1), and in regime D he pays $(1-q)h + qd$. □

Proof of Proposition 4

By Corollary 2, Proposition 2, and Lemma 1,

$$V^M(\theta) - V^C(0;\theta) = \theta\, q\, d(1-q) > 0, \qquad V^M(\theta) - V^D(0;\theta) = (1-\theta)\, q\, (1-\mu)c > 0,$$

the latter because $\mu \le \bar{\mu} < 1$. The first difference applies for $\theta \le \bar{\theta}$ and the second for $\theta > \bar{\theta}$. □

Proof of Proposition 5

Evidence arrives by time t with probability $1 - e^{-\lambda t}$. In regime C, a guilty defendant settles at h upon arrival and otherwise at $t = 1$; an innocent defendant settles at $t = 1$. In regime S, the

guilty defendant additionally settles at $t = 0$ with probability $s(\theta)$; conditional on rejecting, he settles upon arrival or at $t = 1$. In regime D, both defendant types settle upon arrival of evidence; without arrival, the guilty defendant settles at $t = 1$ and the innocent defendant settles at $t = 1$ with probability μ and otherwise goes to trial, so he goes to trial with probability $q(1 - \mu)$. Under mandatory disclosure, both types settle upon arrival and otherwise at $t = 1$ (Lemma 6). □