EVIDENCE ARBITRAGE: THE FABRICATION OF EVIDENCE AND THE VERIFIABILITY OF CONTRACT PERFORMANCE[♣]

CHRIS SANCHIRICO[*]
GEORGE TRIANTIS[+]

THIS DRAFT: MARCH 2004

Abstract

Contract theory identifies verifiability as a critical determinant of the incompleteness of contracts. Although verifiability refers to the cost of proving relevant facts to a court, very little scholarship connects explicitly the evidentiary process to the drafting of substantive contract terms. This paper begins to explore this relationship to provide a more rigorous explanation of contract design. In particular, the paper concerns the very core of verifiability – truth-finding by a court -- and examines the impact of the prospect of evidence fabrication on contracting. It thereby also explores the puzzling tolerance of the adjudicatory system for fabrication and the incentives to fabricate created by thresholds in burdens of proof. The paper suggests that, despite undermining truth finding, evidence fabrication may be harnessed by contracting parties to improve the (evidentiary) cost-efficiency of performance incentives in their relationship.

We divide the future into evidentiary states. In each state, nature produces evidence of breach by the promisor; in the zero state, nature produces no evidence of breach. The probability of each *positive* state is greater if the promisor actually breaches (q) than if she performs (p). In each positive state, the prospect of damages liability has two effects: it induces the promisee to spend resources to present evidence (truthful or fabricated) to secure a damages award and it gives a performance incentive to the promisor (who wishes to reduce the probability of that state). At the time of contracting, the parties view these prospects in expected value terms: (i) the deterrence effect from the *expected* cost of breach and (ii) the *expected* evidentiary cost. The parties contracting objective is to maximize across states the performance incentive bang at lowest cost. They may arbitrage across states to get the largest incentive bang for their evidentiary buck: selecting for states with high incentive bang (high q-p) at low expected evidentiary cost (low q,p).

As a particular instance of arbitrage, we show that when the likelihood ratio (q/p) increases as the evidentiary state becomes large, the parties may wish to induce less evidentiary presentment in low states and more in higher states. This may involve evidence fabrication at higher states once all the truthful evidence is presented. In this way, the parties engage in a joint, conscious contracting strategy to mislead the court in some future evidentiary states of the world.

[♣] The authors are grateful for financial support from the Eastwood D. Herbert Fund (Sanchirico) and the Horace W. Goldsmith Research Chair (Triantis) at the University of Virginia Law Foundation.
[*] Professor of Law, Business and Public Policy, University of Pennsylvania Law School and Wharton School, Business and Public Policy Department; csanchir@law.upenn.edu.
[+] Perre Bowen Professor of Law, University of Virginia School of Law; gtriantis@virginia.edu

I. INTRODUCTION

Contract theory identifies verifiability as a critical determinant of the incompleteness of contracts: specifically, contracts condition obligations only on contingencies that can be verified to a court.[1] Verifiability refers to the feasibility of establishing the truth or reality to a court. Yet, to many contracting parties, verifiability is at best an intermediate goal because litigation creates value by virtue of its ex ante effect on contract incentives (e.g. for performance or specific investment). Anticipating the judicial resolution of future disputes, contracting parties are likely to be interested in the likelihood or cost of judicial truth-finding only to the extent that the court's ability to discern the truth efficiently improves contract incentives and the gains from trade.

Therefore, parties may increase their gains from trade by contracting for terms whose enforcement relies on facts that cannot be verified, but nevertheless promote efficient contract incentives. In fact, actual contracts frequently contain terms that do not seem to be verifiable to a court.[2] More strikingly, courts are often fooled by fabricated, suppressed or otherwise manipulated evidence due to a number of systemic causes. Perjury prosecutions are very rare in civil cases and therefore scarcely deter false testimony. In addition, burden of proof thresholds encourage fabrication by parties with slightly less evidence than that needed to satisfy the burden. Judge Posner recently described the apparent tolerance of the legal system for evidence fabrication in the following passage:

> Even judges have a certain ambivalence about perjury in civil litigation. It is not unusual for one judge to say to another that he or she has just presided at a trial at which several of the witnesses were obviously lying… I have heard expert witnesses referred to as 'paid liars'. These comments are generally not made in a tone of indignation, and they very rarely lead to a referral to the department of justice to inquire into the possibility of an

[1] E.g., B. Salanie, The Economics of Contracts ch. 7 (1997); Alan Schwartz, Relational Contracts in the Courts: An Analysis of Incomplete Agreements and Judicial Strategies, 21 J. Legal Stud. 271 (1992).

[2] The observation and several explanations are explored in George G. Triantis, The Efficiency of Vague Contract Terms, 62 La. L. Rev. (2002).

> obstruction of justice. Part of this reaction is due to a sense that the court system has been designed, or at least has evolved, to be robust in the face of the known inefficacy of the oath and of the threat of prosecution for perjury… and as result, of the frequency of these crimes…[3]

If contracting parties were concerned only about truth finding in future adjudication of their disputes, the judicial system's leniency toward false evidence would intensify their inclination to limit agreements to easily verifiable terms.[4] In this article, however, we suggest that the parties themselves might in fact prefer to permit evidence fabrication as part of a conscious contracting strategy that emphasizes efficient performance incentives over accuracy or fairness in the resolution of disputes.

For the purpose of exposition, we analyze the factual determination of whether a promisor (seller) has performed as promised in her contract.[5] We describe sufficient conditions under which parties may improve incentives for performance by agreeing to contract terms that deliberately create incentives for the future plaintiff (the promisee/ buyer) to fabricate evidence (or, equivalently, to suppress negative evidence). Moreover, we show that fabrication may enhance the efficient provision of performance incentives even though the cost of presenting fabricated evidence is likely to be greater than the cost of truthful evidence.

In our analysis, we distinguish between evidence created by nature (truthful evidence) and evidence fabricated by the plaintiff, and we assign different exogenous costs to the presentation of each at trial. The incentive of the plaintiff to present truthful or fabricated evidence is a function of those presentation costs and her expected increase in litigation payoffs from presentation. Building on a suggestion in Kaplow (1994),[6] we allow the contracting parties to agree ex ante to a *litigation*

---

[3] Richard A. Posner, An Affair of State 147 (1999).

[4] Arguably, contracts that provide for alternative dispute resolution may be opting for a better informed adjudicator and more severe reputational sanctions on evidence fabrication. Our paper raises doubts as to this motivation for preferring nonjudicial referees.

[5] The analysis applies readily to the companion issue of whether a contingency excusing performance materialized.

[6] In discussing the accuracy in adjudicating a standard comprising many elements, Kaplow writes: "One might say that, once an element is deemed to be relevant in a particular manner, there remains the question of how accurately to ascertain that element. But

payoff schedule that associates a damages award to each evidentiary presentation. The promisor subsequently decides whether to perform in light of how performance affects her expected liability for breach which, in turn, is determined by the anticipated evidentiary strategy of the promisee (and future plaintiff). The parties choose the litigation payoff schedule that optimizes the creation of performance incentives: that is, it achieves an optimal balance between the benefits of performance incentives and expected evidence production costs.

We define the "evidentiary state" as the truly existing evidence of breach, and "positive evidentiary states" as those states whose the probability of occurring if the promisor actually breaches (q_i) is larger than their probability of occurring if she performs ($p_i < q_i$). It is important for performance incentives that the litigation payoff schedule does not induce fabrication in the non-positive states because that would effectively sanction the promisor for performance rather than breach. However, a payoff schedule that induces the promisee to fabricate evidence in positive evidentiary states may enhance performance incentives. It would increase the level of liability in these positive states and thereby raise the promisor's sanction for breach.

However, the parties care not just about the performance incentive but also about the cost of creating it. If fabrication is used to improve incentives, it is likely to increase the cost of evidence presentment because fabricating evidence is typically more costly than truthfully presenting it. Yet, we show that inducing some amount of fabrication in positive evidentiary states may still be the most cost effective means of setting performance incentive.

The key insight is that, at the time of contracting, the parties are concerned with *ex ante* breach sanctions and *ex ante* litigation costs. It follows that (1) the effect of increasing the promisor's

the distinction is largely semantic. In principle, one can state a functional relationship directly from evidence to outcomes: findings on elements are a function of evidence, and outcomes are a function of findings on elements; thus outcomes are a composite of these two functions and can be determined from the evidence." Louis Kaplow, The Value of Accuracy in Adjudication: An Economic Analysis, 23 J. Legal Stud. 307, 327 f.49 (1994) (hereafter Accuracy in Adjudication).

liability in any given evidentiary state on the ex ante incentive depends on the difference in the probabilities of that state at the time of the contract given performance and given nonperformance and, similarly, (2) the effect of presenting truthful or fabricated evidence in any given state on ex ante evidence production costs depends on the ex ante likelihood of that state. More precisely, the discount factors that translate these ex post state-by-state values into their contribution to ex ante values are, respectively, (1) the difference between the probability of a given state if the promisor breaches less the probability if she performs (that is, $q_i - p_i$) and (2) the ex ante probability of a given state, given the performance incentive (which we call r_i).

Most significantly in our analysis, we draw attention to the ratio of the first factor over the second —i.e., the probability difference for the state divided by the state's ex ante probability. For reasons that will become clear, we call this the "incentive efficiency ratio" of the state: $\frac{q_i - p_i}{r_i}$. All else the same, the parties can reduce the cost of achieving any given performance incentive by increasing the promisor's sanction in states in which this efficiency ratio is high and reducing it appropriately where the ratio is low. Where this ratio ranges widely across different evidentiary states—and so is much greater in some states than in others--the parties may be willing to increase the promisor's sanction in states with high efficiency ratios even if that requires paying the greater evidentiary costs of fabrication in those states.

Consider the following simple example. A seller contracts to manufacture and deliver 100 widgets to a buyer. The seller agrees to exercise its best efforts to ensure that the widgets are in excellent condition when delivered. The sole evidence of the seller's breach of this promise is the number of broken widgets. There are a hundred positive evidentiary states, each representing a distinct number of broken widgets, and one residual state (no broken widgets). By definition, each evidentiary state is more likely if the seller has in fact breached than if she has performed ($q_i - p_i$). However, the probability difference varies, as does the ex ante probability of each evidentiary state. For the sake of exposition, suppose that the incentive efficiency ratio is greater for higher than for lower states (we suggest later why this may be the case). Then, as a general proposition, the buyer should optimally be discouraged from presenting even truthful evidence in the lower (less incentive efficient) states and encouraged to present both truthful and even fabricated evidence in high (more incentive efficient) states.

Compare in this light, two alternative litigation payoff schedules. The first awards the plaintiff incremental damages for each broken widget that is presented to the court. It deters fabrication by ensuring that the incremental payoff is less than the cost of fabricating evidence. The alternative schedule awards no damages for presenting fewer than, say, 50 broken widgets and a fixed payment where the buyer presents 50 or more broken widgets. Thus, the alternative schedule has a threshold character to it, much like the evidence-reward structure induced by "burdens of proof." The discontinuity in payoffs also encourages fabrication. Suppose that the fixed payment is set so that any buyer with between 40 and 49 broken widgets will find it profitable to fabricate evidence necessary to come up with 50 broken widgets. The effect of this alternative payoff schedule is that the seller of fewer than 40 broken widgets bears no liability. The seller of greater than 40 broken widgets bears a fixed liability. Yet, if the incentive efficiency ratios for the states above 40 are sufficiently higher than the incentive efficiency ratios for the states below 40, the alternative payoff schedule may be preferable even though it encourages fabrication. Later in the paper, we demonstrate this result formally for the more general case where the likelihood ratio $\frac{q_i}{p_i}$ is greater for higher evidentiary states.[7] However, we emphasize that it is the volatility of the ratio that creates the fabrication result, not the direction of its movement across states. One might think of the parties as engaging in evidence arbitrage among future states of the world.

A numerical example in Part II provides a more elaborate illustration of these concepts. Part III sets out a more general and formal model. Using this model, Part IV formally defines the incentive efficiency ratio and establishes our general proposition that contracts encouraging evidence fabrication in at least some future states of the world may achieve *at lower cost* any deterrence level offered by non-fabrication alternatives. Part V offers some implications for the theory of incomplete contracts and an agenda for future research.

Our research is related to several strands of scholarship. First, although most incomplete contracts scholarship assumes that verifiability is binary and exogenous (i.e., some facts are feasibly proven to courts, others are not), recent work has begun to integrate models of evidence production and adjudication with contract design.[8] These papers assume that although litigating parties may omit to produce evidence, they may not fabricate evidence. Second, Sanchirico proposes viewing evidence production as a costly signaling game, the anticipation of which influences primary activity behavior.[9] These models allow parties to fabricate, but do not address the question of whether fabrication would form part of the optimal incentive setting mechanism. Third, the conception of verifiability in contracts has been addressed under the more general concern about accuracy in several important articles by Louis Kaplow and Steven Shavell. We

[7] The proof for the opposite case of decreasing likelihood ratio is available from the authors.

[8] E.g., Jesse Bull and Joel Watson, Evidence Disclosure and Verifiability (Working Paper 2001); Jesse Bull, Costly Evidence Production and the Limits of Verifiability (Working Paper 2001).

[9] Chris Sanchirico, Relying on the Information of Interested--and Potentially Dishonest--Parties, 3 Am. L. & Econ. Rev. 320 (2001), Games, Information and Evidence Production: With Application to English Legal History, 2 Am. L. & Econ. Rev. 342 (2000). A variant of this model appears in Antonio Bernardo, Eric Talley and Ivo Welch, A Theory of Legal Presumptions, 16 J. L., Econ. & Org. 1 (2000). An earlier paper suggests analyzing evidence production as costly signaling, but in fact models the court's interpretation of evidence as exogenous, and also does not consider the feedback effect of evidence production on primary activity incentives: Daniel L. Rubinfeld and David E.M. Sappington, Efficient awards and standards of proof in judicial proceedings, 18 RAND J. Econ. 308 (1987).

share with them the focus on the cost efficiency of incentives provided by adjudication (the tradeoff between performance incentives and litigation costs).[10] However, we differ from their work in that we consider the prospect that parties expend litigation resources to mislead – rather than inform–the court. Thus, we arrive at a conclusion that contracting parties may not only be uninterested ex ante in investing in accuracy, but they may in fact promote the expenditure of resources (evidence fabrication) to positively reduce the accuracy of adjudication.

II. AN EXAMPLE

Suppose that a buyer sues a seller for breach of a contract to deliver a good, and suppose also that the buyer may have no evidence, low evidence, medium evidence, or high evidence of the seller's breach. Each individual positive level of evidence is more likely when the seller has actually breached than when she has not, whereas no evidence is more likely when the seller has not breached.[11] Let the probability of each evidentiary state given performance be p_i, where $i = 0, L, M, H$. The corresponding probabilities following breach are q_i. Thus, $q_0 < p_0$, but in all other states $q_i > p_i$.

The buyer's cost of truthfully presenting each level of evidence at trial is \$0, \$1, \$2, or \$3 respectively. The buyer can also fabricate evidence. For example, when she truly has only low evidence, she can combine this with additional fabricated evidence in order to present medium evidence of breach. The cost of fabricating each additional level of evidence is \$2. For example, if the buyer truly has low evidence and fabricates up to medium evidence, her total cost of evidence is \$1+\$2 = \$3, whereas if she fabricates up to high evidence her total cost is \$1+\$2+\$2 = \$5.

[10] Louis Kaplow and Steven Shavell, Accuracy in the Determination of Liability, 37 J. Law & Econ. 1 (1994): Louis Kaplow and Steven Shavell, Accuracy in the Assessment of Damages, 39 J. Law & Econ. 191 (1996); Kaplow, Accuracy in Adjudication, supra note --; Kathryn E. Spier, Settlement Bargaining and the Design of Damage Awards, 10 J. Law, Econ. & Organ. 84 (1994).

[11] This implies that the probability distribution of evidence levels given breach first order stochastically dominates the probability distribution of evidence levels in the absence of breach.

After the buyer presents evidence at trial, the court may require the seller to pay some level of damages to the buyer. The buyer's decision about how much evidence to present (truthfully or otherwise) depends on her net payoffs from different levels of presentation. For example, if she truly has only low evidence, but her recovery increases by $3 if she presents medium rather than low evidence, she earns $3 by spending only $2 to fabricate from the low to medium level. Thus, as between truthfully presenting low evidence or partially fabricating medium evidence, she will chose the later. In other cases, the buyer may decide to withhold evidence. For example, if her litigation payoffs increase only by $.50 when she presents her actual low evidence rather than no evidence, then she would prefer to present no evidence.

We start with a litigation payoff schedule[12] that rewards the buyer for each higher level of evidence with (slightly less than) $2 additional dollars of transfers from the seller. For example, if the buyer presents low evidence she gets $2, and if she presents high evidence, she gets $6. This schedule induces the buyer to present precisely the amount of evidence that she truly has. To see this, suppose she truly had medium evidence. If she presents medium evidence she receives $4 at a cost to her of $2. If she withholds some evidence, presenting only low evidence, her $2 loss of reward would not justify the $1 of evidence costs savings. On the other hand, if she fabricates evidence, the (slightly less than) $2 of additional reward that she obtains by fabricating to the next level does not justify the additional $2 of evidence costs. The same no-fabrication-no-withholding result applies to the other "types" of buyers (i.e. buyers holding no evidence, low evidence or high evidence).

This litigation payoff schedule enables a sophisticated court to deduce the true state, because it induces the buyer to present a distinct level of evidence in each evidentiary state. The true evidentiary state is an informative signal of whether there was breach because the probability of

[12] We discuss the interpretation of this schedule later in the paper.

each evidentiary state depends on the breach decision. The more the court knows about the true evidentiary state, therefore, the more it knows about whether there was breach. Thus, because the schedule we are considering is maximally informative about the true evidentiary state, it is also maximally informative about whether there was breach. Put another way, breach is maximally verifiable under this schedule.

However, the contract objective of the parties is to set efficient incentives for the seller to perform, not to maximize a court's ability to verify whether there was performance. The no-fabrication payoff schedule described in the previous two paragraphs does indeed create an incentive for the seller to perform. The buyer's award—thus, the seller's liability—is greater in evidentiary states that are more likely following breach. For example, when there is no evidence of breach, the buyer presents no evidence and the seller pays nothing to the buyer. On the other hand, when there is medium evidence the buyer presents medium evidence and the seller must pay $4 to the buyer. Breaching effectively shifts probability weight from the no evidence state to the medium evidence state, inter alia. Therefore, the seller has an incentive to perform so as to decrease the likelihood that she will face $4 rather than $0 of liability.

Yet, this no-fabrication schedule may not be the most efficient means of creating a performance incentive. Consider an alternative schedule under which the seller pays (slightly less than) $5 if the buyer presents high evidence, and nothing for any lower level of evidence. The high level of evidence under this schedule might be thought of as the burden of proof.

The buyer's evidence presentment incentives are different than under the original schedule. If the buyer has no evidence, then the $6 cost to her of fabricating up to high evidence is not worth the $5 award. If the buyer has low evidence, the $5 (= $1 + $2 + $2) cost of producing high evidence is also not worth the (slightly less than) $5 award. However, if the buyer has medium evidence, then the $5 payoff from meeting the threshold of high evidence is worth the cost ($4 = $1 + $1 + $2), even though fabrication is required. Similarly, presenting high evidence is also worthwhile for the buyer when she truly has high evidence.

The court learns less about performance under this alternative payoff schedule than under the previous no-fabrication schedule. The court learns only whether or not the buyer has at least medium evidence. The buyer presents the same amount of evidence—none—whenever she has either no evidence or low evidence. The buyer also presents the same amount of evidence—high—whenever she has either medium or high evidence.

The more important economic question, however, is whether this pooling of evidentiary states improves or worsens the efficiency of contract incentives. At first blush, one might believe the alternative schedule is inferior on two scores. First, it appears to generate less of a performance incentive. Comparing the damages paid by the seller under the no fabrication schedule in each state ($0, $2, $4, $6) with the same for the fabrication-inducing schedule ($0, $0, $5, $5),[13] we see that the fabrication-inducing schedule's damages are $1 greater in the medium state ($5 versus $4), but $1 less in the high state ($5 versus $6) and $2 less in the low state ($0 versus $2). Across states, then, the fabrication-inducing schedule appears to offer $2 less of a performance incentive. See Figure 1A.

Second, the fabrication inducing schedule also seems to be more costly. Compare the evidence produced by the buyer under the no fabrication schedule in each state (none, low, medium, high) to the same for the fabrication inducing schedule (none, none, high, high). Under the fabrication-inducing schedule, the buyer presents one "unit" less in the low state (none versus low) and one unit more in the medium state (high versus medium). Moreover, the fabrication-inducing

[13] Note that these are damages actually paid given then buyer's evidence production decision in each state, which involves fabrication in the medium state.

schedule's additional unit of production in the medium state is fabricated, thus costing an additional $2 dollars, while the no fabrication schedule's additional unit of production in the low state is truthfully presented, thus costing an addition $1. Across states, then, it appears that the fabrication-inducing schedule is $1 more expensive. See Figure 1C.

This analysis is flawed, however, because one cannot simply add across ex post states. Rather, one must discount the sanction and cost of each ex post state by the relevant probabilistic expression to determine that state's contribution to the ex ante performance incentive and ex ante evidence costs.

Consider first the performance incentive. To determine the performance incentive we must weight the damages in each state by that state's *probability difference*, $(q_i - p_i)$. To take an extreme example, if the probability of the low state were the same regardless of performance ($q_L = p_L$), then the level of damages in that state would have no effect on behavior. More to the point, one cannot conclude that the no fabrication schedule produces a greater incentive than the fabrication schedule without considering these probability differences. For example, suppose that the probabilities of low, medium and high evidence states given performance are 20%, 20% and 10%, while the same probabilities given breach are 15.5%, 2% and 1%. Then, the probability differences (in terms of percentage points) are 4.5, 18, and 9, respectively. Therefore, the fact that the fabrication-inducing schedule has greater damages in the medium evidence state will be magnified in importance, relative to the fact that it has lower damages in the low and high states. Indeed, one can check that the performance incentive is the same under both schedules in this example.[14] See Figure 1B.

[14] Multiplying the probability differences by the differences in damages across the two schedules we have 4.5(-2)+18(1)+9(-1) = 0.

Figure 1A: Ex post damages

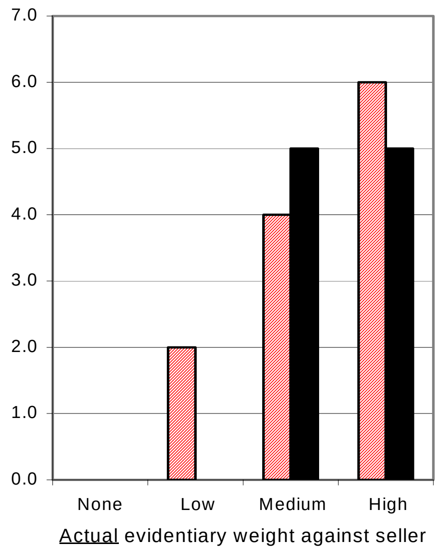

Figure 1B: Contribution to ex ante performance incentive

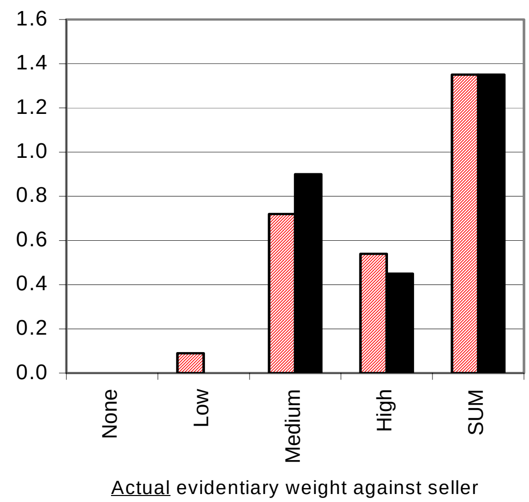

Figure 1C: Ex post evidence cost

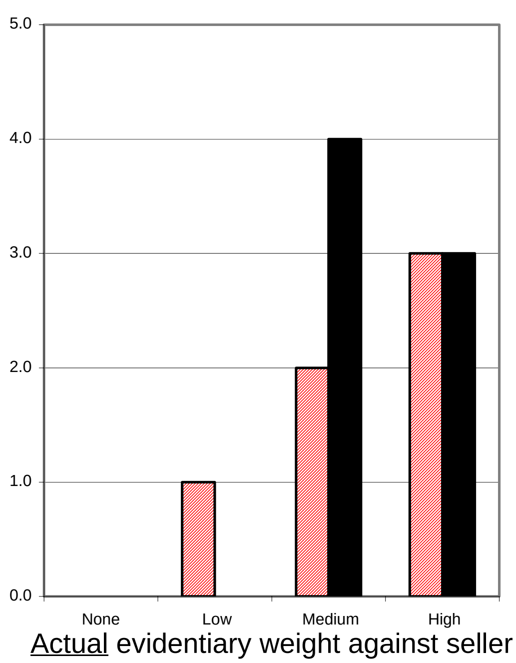

Figure 1D: Contribution to ex ante evidence cost

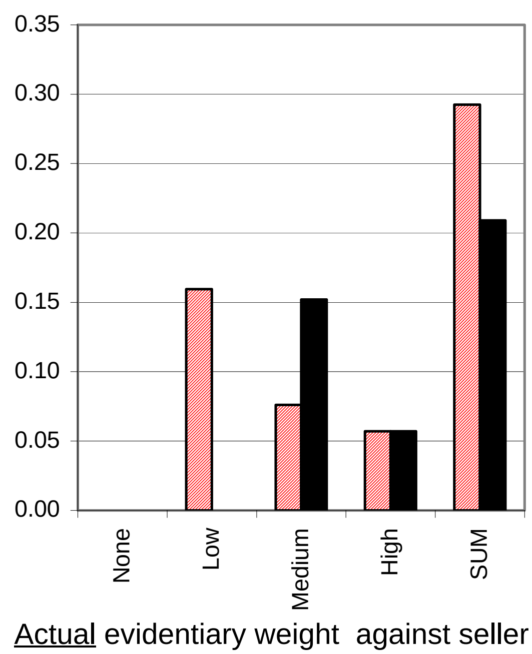

Consider now the ex ante costs of each schedule. To determine the ex ante cost of each schedule, we have to multiply the ex post cost in each state by the *ex ante probability* of that state. The ex ante probability of each state depends on the probability that the seller breaches, the probability of each state given breach (q_i) and the probability of each state given performance (p_i). The probability of breach in our model is the chance that the seller's performance costs—unknown at the time of contracting—is less than the performance incentive (i.e., the expected increase in damages from not performing). To continue our example, suppose that the cost of performance is uniformly distributed between 0 and 2. As noted, the performance incentive is the same under the two schedules that we are comparing, and one can calculate that the incentive is 1.35. Therefore, the probability of breach is roughly 10%. It follows that the ex ante probability of any given state *i* is 10% of q_i (the probability of *i* conditional on breach) and 90% of p_i (the probability of *i* conditional performance). Using this formula, one can calculate that the ex ante probabilities of the low, medium and high evidence states are roughly 16%, 4% and 2%. The low evidence state is much more likely to occur than the medium or high state. Therefore, the evidence costs conditional on that state are more likely to be incurred. The no fabrication schedule has higher evidence costs than the fabrication-inducing schedule in the low evidence state. Therefore, the no-fabrication schedule may be in fact more expensive, contrary to our earlier impression. Some arithmetic confirms that the no fabrication schedule imposes roughly .3 in ex ante evidence costs on the parties, while the fabrication schedule imposes only .2. See Figure 1D.

Moving from the no-fabrication schedule to the fabrication-inducing schedule effectively shifts a unit of evidence production from the low state to the medium state. Although the unit of evidence that was truthfully presented in the low state under the no-fabrication schedule is now fabricated at greater ex post cost, in the medium state under the fabrication-inducing schedule, the result is a decrease in ex ante evidence costs because the medium state is that much less likely to occur. At the same time, the probability *difference* (as opposed to the ex ante probability) is relatively large in the medium evidence state. Therefore, the increase in damages in that state that results from shifting from the no-fabrication schedule to the fabrication-inducing schedule is enough to offset the loss of damages in the low and high states.

As noted in the introduction, the key to the argument is the incentive efficiency ratio for each state: the ratio of the probability difference $(q_i - p_i)$ to the ex ante probability of each state (r_i).. This ratio contains important information about the efficiency of sanctioning in that state. When the ratio is volatile across states, the parties can improve efficiency by shifting sanctions into states with relatively low values for this ratio and out of states with relatively high values. This arbitrage opportunity is reflected in Figure 2 which shows that the incentive efficiency ratio in the low state in our example is significantly lower than the efficiency ratio in each of the medium and high states. Fabrication can be an essential component of this arbitrage process, even though it is more costly ex post than the truthful presentation of actual evidence. In Parts III and IV that follow, we employ a more formal model to demonstrate that any given level of performance incentive is achieved at lowest cost by a payoff schedule that induces fabrication.

Figure 2: Incentive efficiency ratio: (*q* -*p*)/*r*

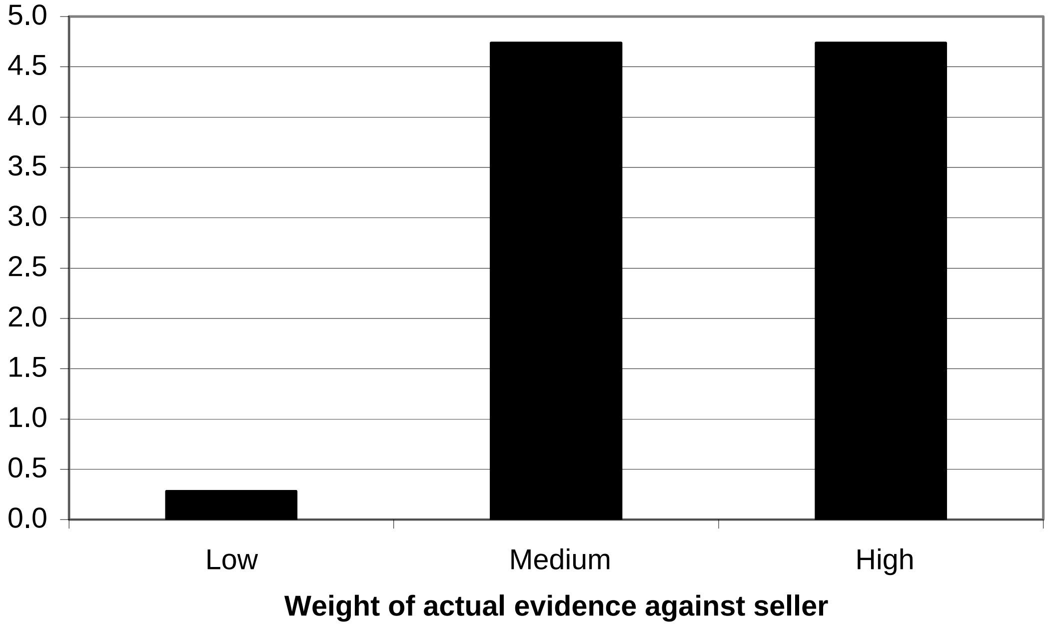

III. MODEL

A seller and a buyer form a contract. The seller chooses whether to perform under the contract. The buyer chooses how much evidence to present in court of the seller's non performance.

We model the buyer's evidence production as differential cost signaling, where signal costs are endogenous to whether the seller performed.[15] Differences in evidence costs are driven by differences in the weight of evidence actually available for production: fabrication is more expensive than truthful presentation.

In order to study the role of verifiability and fabrication in contract design, we hypothesize that the correspondence between evidence and verdict is fully within the contractual authority of the parties.

[15] Sanchirico (1995, 2000, 2001) models evidence production as endogenous cost signaling. Bernardo, Talley, and Welch (2000) employ a similar model.

A. *Timeline*

The timeline of the model is illustrated in Figure 3. In phase 1, a risk neutral seller and a risk neutral buyer form a contract. A contract consists of three elements: the seller's promise of later performance;[16] the price paid up front by the buyer; and a *litigation payoff schedule*, $t(e)$. Discussed in detail below, $t(e)$ maps evidence production *e* by the buyer onto transfers *t* from the seller to the buyer. The parties design the contract to maximize their joint expected profits, as formally defined within. (The distribution of profits between buyer and seller is determined by the contract price.)

[16] Performance might consist of the delivery of a good with particular characteristics, or the provision of a service of particular quality.

Figure 3

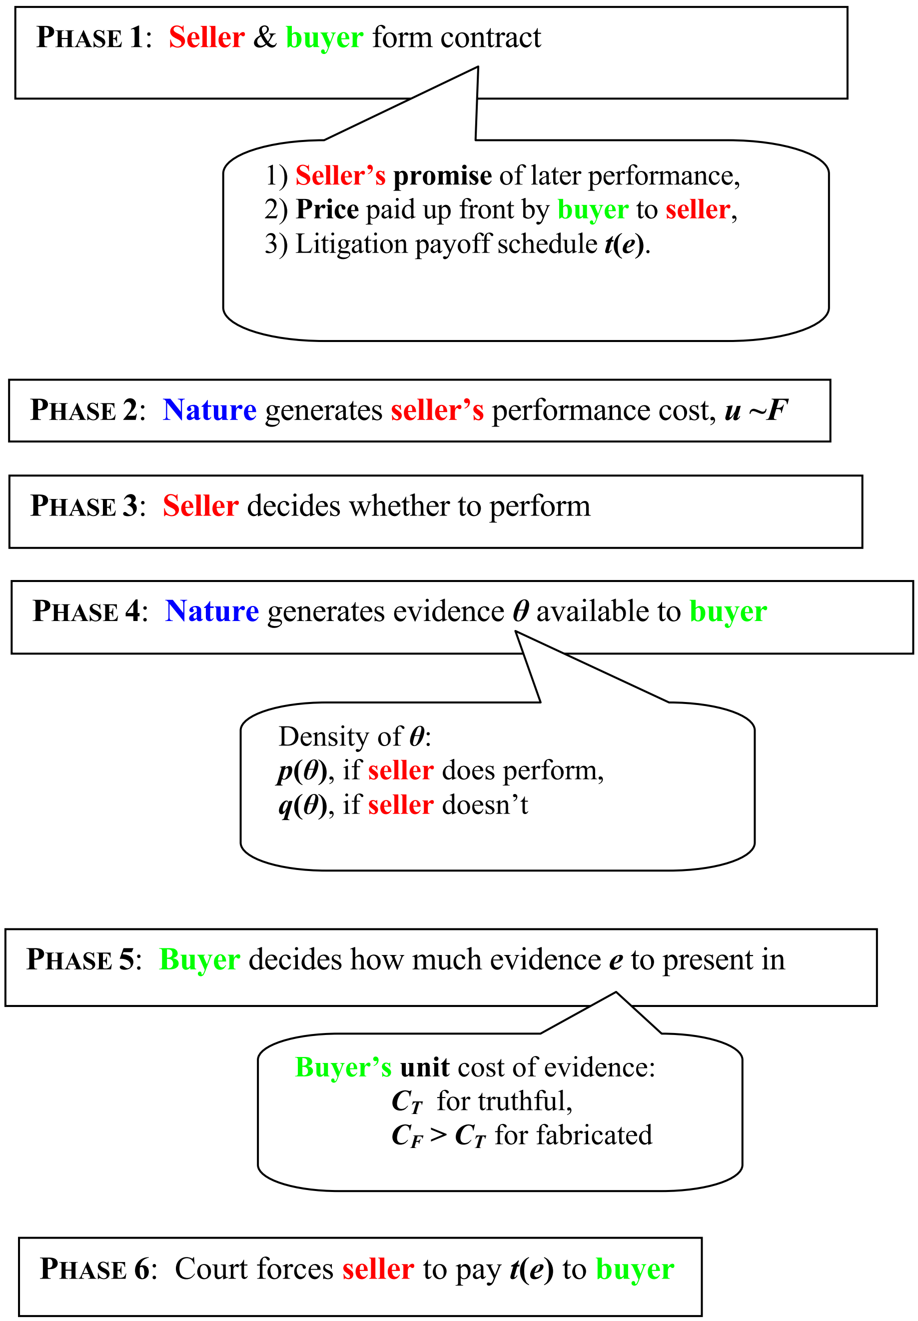
PHASE 1: Seller & buyer form contract
1) Seller's promise of later performance,
2) Price paid up front by buyer to seller,
3) Litigation payoff schedule $t(e)$.
PHASE 2: Nature generates seller's performance cost, $u \sim F$
PHASE 3: Seller decides whether to perform
PHASE 4: Nature generates evidence θ available to buyer
Density of θ:
$p(\theta)$, if seller does perform,
$q(\theta)$, if seller doesn't
PHASE 5: Buyer decides how much evidence e to present in
Buyer's unit cost of evidence:
C_T for truthful,
$C_F > C_T$ for fabricated
PHASE 6: Court forces seller to pay $t(e)$ to buyer

In phase 2 nature generates the seller's cost u of performing under the contract. The cumulative distribution of u is F. The density is f.

In phase 3, the seller decides whether or not to perform under the contract. The seller weighs the cost u of performing against the potentially positive effect that performing has on her payoffs in the litigation phase of the model, as described below. If the seller performs, the buyer enjoys value v.

In phase 4, nature determines the "weight" of evidence $\theta \in [0, \infty)$ actually available to the buyer.[17] We will refer to θ variously as the "weight of available evidence," the "buyer's type," and the "evidentiary state." θ is probabilistically dependent on whether the seller has performed under the contract. If the seller has performed, the probability density of θ is p. If the seller has not performed, the density is q. We assume that q and p are everywhere positive, and have finite expectations. If the likelihood ratio $\frac{q}{p}$ is increasing in θ, we may think of θ as the weight of available evidence of non performance.[18]

In phase 5, the buyer decides how much evidence e to present in court.[19] The buyer may present more or less evidence than is actually available. If the buyer presents more (i.e., if $e > \theta$), the

---

[17] In accord with common parlance in evidence law, we measure evidence along a single dimension, "weight." We have found that a more complicated model that distinguishes evidentiary states along more than one dimension leads to similar results. On the other hand, modeling such a multidimensional type space greatly complicates the analysis.

[18] If $\frac{q}{p}$ increases, then given any prior odds that the seller has performed, the posterior odds of performance conditional on θ are greater the greater is θ

[19] Although the litigation mechanism in this model does not include evidence production by the seller, we can think of the mechanism design problem here as that of designing a particular slice—corresponding to a fixed level of seller evidence—of an overall mechanism wherein both parties present evidence.

Because we consider only the buyer's evidence production, our mechanism does not take advantage of correlation between seller and buyer types. (On the potential advantages of exploiting correlated types, see, e.g., DREW FUDENBERG & JEAN TIROLE, GAME THEORY, 1991 (Cambridge: MIT Press) at 292-295.) Yet, in a litigation context, correlated types are not the panacea that they may (or may not) be in other contexts. See, generally, Chris Sanchirico, *Games, Information and Evidence Production: With Application to English Legal History*, 2 AM. L. & ECON. REV. 342 (2000). First, fully exploiting the advantages of correlated types requires decoupling plaintiff's recovery from defendant's damages. (Otherwise the so called "full rank assumption" may be difficult to satisfy.) Yet public budget constraints restrict the degree to which plaintiff's recovery can exceed defendant's liability. And the parties' desire to maximize joint profits will dampen their willingness to allow defendant's liability to exceed plaintiff's recovery. Second, even were the parties able fully to decouple their transfers, the fixed costs of legal process may induce them to pool a relatively large subset of states into the set of those in which the plaintiff is not induced to file suit at all. Depending on the relative size of these fixed costs, it may be in the parties' joint interest to pool to such an extent that they defeat the full rank assumption, in which case they would resort to the kind of costly evidence production that we model here. For more on this point, see id.

buyer is said to "fabricate" the shortfall $e-\theta$. The buyer bears the cost of evidence production. The unit cost for available evidence is c_T. The unit cost of fabricated evidence is c_F. The cost of presenting evidence *e* for a buyer of type θ is thus[20]

$$c(e,\theta)=\begin{cases} c_T e, & e\le\theta \\ c_T\theta+c_F(e-\theta), & e>\theta \end{cases}.$$

Unless otherwise noted, we assume $c_F > c_T$.

In choosing how much evidence *e* to present, the buyer balances the cost of evidence against the positive effect that such evidence will have on his winnings in court. In particular, he chooses *e* to maximize $t(e)-c(e,\theta)$.

In phase 6, the court "rules." The court observes neither whether the seller performs, nor the weight of evidence θ actually available to the buyer, but only the evidence *e* produced by the buyer.[21] Given *e*, the court forces the seller to pay $t(e)$ to the buyer as stipulated in the contract formed in phase 1.

B. *Intuitive mechanics of the model*

All else the same, the parties would like the seller to perform whenever *u* falls below *v*. But the seller would never perform at all without the litigation phase of the model. The litigation phase induces the seller to perform to the extent that nonperformance raises the expected transfer *t* that she must pay to the buyer. How might this occur? Nonperformance may effect expected transfers

[20] Two technical notes are evidence costs. First, although simple in form, evidence costs are not everywhere differentiable. Although this produces some technical complications (e.g., in proving that the lowest type present no evidence), the net effect of adopting this evidence cost function is to greatly simplify the analysis.

Second, evidence costs do not satisfy a strict single crossing property: a strictly higher type does not have strictly greater incremental evidence costs at all levels of evidence. For example, the cost of increasing evidence from 10 to 20 is the same for types 30 and 40 and also the same for types 3 and 4. Consequently, incentive compatibility does not imply that evidence production is even weakly increasing in type. Again this complication is more than compensated for by simplifications in other areas.

[21] If the court could directly observe the buyer's type, rather than just the evidence that the buyer chooses to send based on his type, our model would conform to the classic moral hazard/hidden action problem. For more on this, see Sanchirico, *Relying on the Information of Interested—and Potentially Dishonest—Parties, 3* AM. L. & ECON. REV. 320 (2001).

through its probabilistic impact on the weight of evidence θ available to the buyer. The seller's nonperformance may, for instance, tend to provide the buyer with more true evidence θ. The more evidence is available to the buyer, the cheaper it is for the buyer to present evidence e to the court. The cheaper it is for the buyer to present evidence, the more evidence the buyer presents. And the more evidence the buyer produces, the greater the transfer $t(e)$ that the seller must pay the buyer.

However, the cost of evidence produced by the buyer in this scenario lowers joint profits and must be traded off against the benefits of the performance incentive created for the seller. Without costly evidence production there is no incentive. If, for example, the buyer could obtain a transfer t of $100 without producing any evidence, then the buyer would always obtain this transfer, no matter what evidence θ were actually available. The seller would then have to pay $100 to the buyer, regardless of whether she performed. Consequently, the seller would have no incentive to perform. Generating a performance incentive requires that buyers of different types obtain different sized transfers from the seller. This, in turn, means that buyers of different types must present different amounts of evidence. And this, in turn, requires costly evidence production—with costs differing by type—so that, for any given transfer level, some buyer types find the transfer worth the cost of the evidence required to obtain it, and some do not.

C. *Preliminary formalities*

1. *Direct mechanisms and incentive compatibility*

In considering which litigation payoff schedule $t(e)$ to include in their contract the parties will be interested in how much evidence $e(\theta)$ each buyer type θ chooses to present and the size of the corresponding transfer $t(\theta)$ that the seller will be forced to pay to this buyer. It is more convenient—and equivalent—to imagine that the parties contract directly over $(e(\theta), t(\theta))$

limiting their attention to such *direct (litigation) mechanisms* that are *incentive compatible*. A direct mechanism $\left(e(\theta), t(\theta)\right)$ is incentive compatible if for any two types $\theta, \theta' \in [0, \infty)$,

$$t(\theta) - c\left(e(\theta), \theta\right) \geq t(\theta') - c\left(e(\theta'), \theta\right). \tag{1}$$

The requirement is that any given type would (weakly) prefer to choose the evidence and transfer "assigned" to her by the direct mechanism than the evidence and transfer assigned to any other type.[22]

2. *The performance incentive and the expected gains from performance*

If the seller performs at phase 3, the seller's expected litigation payoff under mechanism $\left(e(\theta), t(\theta)\right)$ is $-\int_0^\infty t(\theta)\, p(\theta)\, d\theta$. If the seller breaches, her expected litigation payoff is $-\int_0^\infty t(\theta)\, q(\theta)\, d\theta$. Therefore, the seller performs, if and only if

$$u - \int_0^\infty t(\theta)\, p(\theta) \geq -\int_0^\infty t(\theta)\, q(\theta) \quad \Leftrightarrow \quad u \leq \int_0^\infty t(\theta)\left(q(\theta) - p(\theta)\right) \equiv \Delta \tag{2}$$

We refer to the defined term Δ in (2) as the *performance incentive*. The performance incentive is a weighted sum across all evidentiary states of $t(\theta)$, with weights equal to the probability difference in each state. Positive transfers from the seller to the buyer in states that are more likely following nonperformance (i.e., states with $q(\theta) > p(\theta)$) increase the performance incentive.

[22] The equivalence of working with incentive compatible direct mechanisms is guaranteed by the well known result that (1) the buyer's type-contingent reaction to every litigation payoff schedule produces an incentive compatible direct mechanism $\left(e(\theta), t(\theta)\right)$; and conversely, (2) every incentive compatible direct mechanism is the buyer's type-contingent reaction to some litigation payoff schedule $t(e)$. See, e.g., Fudenberg and Tirole, supra note __ at 255-257. As is also standard, we will consider only piecewise continuously differentiable mechanisms. See, e.g., Fudenberg and Tirole (1991), p. 258. We adapt this definition to our unbounded type space as follows: within any interval of finite length there are no more than a finite number of points at which the mechanism is not continuously differentiable.

Given performance incentive Δ, the expected joint gain from performance at the time of contracting (ex ante revelation of performance costs *u*) is:

$$G(\Delta) \equiv \int_0^{\Delta} f(u)(v-u)\,du$$

(3)

Until it reaches *v*, a greater performance incentive Δ increases the expected gains from performance. Above *v*, the seller is performing even where the cost to him of doing so is greater than the benefit to the buyer. Therefore, if performance were costlessly observable by the court, the optimal litigation mechanism would produce performance incentive $\Delta = v$. The presence of litigation costs, however, may affect the optimal mechanism because such costs are not minimized at $\Delta = v$.[23]

3. *The ex ante probability of available evidence and the expected cost of litigation*

The probability that the seller will perform is the probability that the seller's cost of performance *u* is less than the performance incentive Δ. This probability is $F(\Delta)$. Therefore, the *ex ante*[24] *density* of the buyer's type θ is

$$r(\Delta,\theta) \equiv F(\Delta)\,p(\theta) + (1-F(\Delta))\,q(\theta).$$

Given performance incentive Δ and evidence schedule $e(\theta)$, expected evidence costs at the time of contracting are:

$$C(e(\bullet),\Delta) \equiv \int_0^{\infty} c(e(\theta),\theta)\,r(\theta,\Delta)\,d\theta \qquad (4)$$

[23] The performance incentive created by the optimal mechanism may in general be greater or less than v. See generally, Polinsky & Rubinfeld, *The welfare implications of costly litigation for the level of liability* 17 J. LEG. STUD. 151-164 (1988). See, also, Sanchirico, *Relying on the Information of Interested—and Potentially Dishonest—Parties, 3* AM. L. & ECON. REV. 320, ____ (2001).

[24] Specifically, this is the density at the time of contracting, ex ante revelation of the seller's performance cost, and so, a fortiori, ex ante revelation of the buyer's litigation type.

This expression averages evidence costs across all evidentiary states, weighting each state by its "probability" at the time of contracting $r(\Delta,\theta)$.

4. *The optimal contract and the sub problem of cost minimization*

Among the set of incentive compatible litigation mechanisms, the parties choose a mechanism that maximizes expected joint profits at the time of contracting. Expected joint profits equal the net gains from performance by the seller $G(\Delta)$ less the cost of evidence production, $C(e,\Delta)$.

This choice problem can be decomposed into two steps. In the first step, the parties find, for each level of performance incentive Δ, the incentive compatible litigation mechanism that generates that performance incentive at the lowest possible expected evidence cost. Write $C(\Delta)$ for this minimum cost. In the second step, the parties choose which performance incentive to generate, based on the expected gains from performance $G(\Delta)$ and expected evidence costs $C(\Delta)$, as calculated in the first step.

For the main points in this paper, it suffices to focus on the first of these two steps—the minimum cost generation of a given performance incentive. We will provide conditions under which the most efficient method of generating an arbitrarily chosen performance incentive purposefully induces fabrication by the buyer.

5. *Normalization of transfers and the superfluity of the buyer's litigation participation constraint*

Adding or subtracting a constant k from $t(\theta)$ at all types θ (i.e., *translating* t) affects neither the performance incentive Δ nor satisfaction of the incentive compatibility constraints (1). Furthermore, $t(\theta)$ enters into expected litigation costs $C(e(\theta),\Delta)$ only through Δ, which, as noted, is unaffected by translation of the transfer function. Therefore, if two mechanisms (e,t) and (e',t') differ only in that $t = t' + k$ for some constant *k*, then these mechanisms are equivalent for both overall contract design and the sub problem of cost minimization. Thus, from each *equivalence class* of mechanisms so related by transfer translation we need only evaluate one representative in the analysis to follow. It turns out to be most convenient to focus on the representative of each equivalence class whose transfers are set such that

$$t(0) - c(e(0),0) = 0 .$$

The translation invariance of the contracting problem with respect to transfers also makes a litigation participation constraint (i.e., a filing decision for the buyer) superfluous. Because evidence costs are greatest for $\theta = 0$, $t(0) - c(e(0),0) = 0$ implies that $t(0) - c(e(0),\theta) \geq 0$ for all $\theta \geq 0$. Incentive compatibility at θ then implies that $t(\theta) - c(e(\theta),\theta) \geq 0$, as well. Therefore, each equivalence class of litigation mechanisms contains a member that satisfies the additional litigation participation constraint: for all θ, $t(\theta) - c(e(\theta),\theta) \geq 0$. If we choose this representative (as we do in the analysis to follow), the buyer's threat to sue will always be credible.

IV. Analysis

Suppose that the parties wish to provide the seller with performance incentive Δ at the lowest possible evidence cost $C(e(\theta),\Delta)$. First, they would gather all the incentive compatible litigation mechanisms (e,t) that generate Δ. Then they would determine the cost of each and pick the one with lowest cost.

In this section we provide conditions under which the parties would pick a litigation mechanism that purposefully induces the buyer to fabricate evidence in some states. Along the way, we develop the intuition for why this might be so by decomposing incentive setting efficiency into two separate components, one favoring truthful evidence production the other not. In the end fabrication results when the second factor predominates. Results for various limiting cases are provided throughout to help with the intuition.

A. *Two immediate results*

Two limiting cases that help to frame our analysis can be solved by inspection. First, if fabrication is effectively prohibited (i.e., c_F is essentially infinite), then fabrication is never part of the optimal litigation mechanism. This is almost a matter of definition, but it highlights the point that most formal models of evidence production—which rule out fabrication—are ill-equipped to study the optimal evidentiary design of contracts in a world in which the cost of fabrication is often not prohibitive.

Second, fabrication is never optimally induced when $c_T = 0$ (and $c_F > 0$).[25] Under this assumption, all no fabrication mechanisms incur zero cost, whereas all mechanisms inducing fabrication incur positive cost. This second preliminary result points up the fact that the optimality of fabrication requires the existence of a positive cost for the presentation of truthful evidence. Because the presentation of truthful evidence does in practice require investigation and preparation, this assumption seems well in accord with reality.

---

[25] We assume here that the performance incentive is not so large that it cannot be generated by *some* no fabrication schedule. In general, a no fabrication constraint will strictly limit the size of the performance incentive that can be generated. (As discussed in detail below, $t(\theta) \le c_F\theta$ for all incentive compatible no fabrication mechanisms, whereas $t(\theta)$ is unlimited at every θ in the full set of incentive compatible mechanisms.) If the performance incentive exceeds this no fabrication maximum, then fabrication will in a sense be "optimal," if not necessary. The results in this paper , however, address the more interesting possibility that the parties will choose to induce fabrication even though no fabrication mechanisms are also capable of producing the performance incentive.

B. *Within state efficiency: two factors*

When $0 < c_T < c_F < \infty$, determining the most efficient means of generating performance incentive Δ is a more complicated problem. We may begin to evaluate the efficiency of a given litigation mechanism by reviewing its efficiency on a state by state basis. Thus, for any given litigation mechanism (e,t) generating performance incentive Δ, and any given state θ, consider the "price" in terms of evidence costs incurred in that state of the portion of the incentive provided in that state. That is, consider that state's contribution to expected evidence costs per its contribution to the performance incentive:

$$\frac{c(e(\theta),\theta)r}{t(\theta)(q-p)} = \underbrace{\frac{c(e(\theta),\theta)}{t(\theta)}}_{\substack{\text{ex post price of}\\ \text{incentives in}\\ \text{state } \theta}} \times \underbrace{\frac{r}{q-p}}_{\substack{\text{probabilistic}\\ \text{translator}}} ,(5)$$

(Hereafter, we drop the argument "θ" where it will not cause confusion.) As shown in (5), this price can be decomposed into two factors. The first factor $\frac{c(e(\theta),\theta)}{t(\theta)}$ is the price of transfers in state θ conditional on the occurrence of state θ. We will call this the *ex post price of incentives at state* θ. All else the same, the litigation mechanism is efficiently providing incentives in state θ, if it delivers transfers at a low price in terms of evidence costs. The second factor $\frac{r}{q-p}$ translates the ex post price of incentives into an ex ante price. This factor reflects the efficiency of the state θ taken apart from the litigation mechanism under consideration. All else the same, the parties would like to concentrate transfers in states that are significantly more likely following nonperformance, but that rarely occur. A large probability difference $(q-p)$ means that each dollar transferred from seller to buyer in that state has a large impact on the seller's performance incentive. A small ex ante probability (r) means that each dollar of evidence costs incurred by the buyer in that state subtracts little from ex ante profits because that state rarely occurs.

C. *Efficiency over all states*

Ultimately, the parties will choose a mechanism that delivers performance incentive Δ at the lowest *overall* price

$$\frac{C\left(e(\bullet),\Delta\right)}{\Delta}=\frac{\int_0^\infty c\left(e(\theta),\theta\right)r}{\Delta}. \tag{6}$$

This overall price (6) can conveniently recast as the weighted average of the state by state price in (5)

$$\int_0^\infty \frac{t(\theta)(q-p)}{\Delta}\frac{c\left(e(\theta),\theta\right)}{t(\theta)}\frac{r}{q-p} \tag{7}$$

The weight $\frac{t(\theta)(q-p)}{\Delta}$ at each state θ is the proportion of the total incentive Δ contributed at state θ under the mechanism.[26] Thus, a litigation mechanism delivers performance incentive Δ at low overall price if most of its performance incentive is delivered in states with low $\frac{c\left(e(\theta),\theta\right)}{t(\theta)}\frac{r}{q-p}$.

D. *Fabrication and the ex post price of incentives* $\frac{c\left(e(\theta),\theta\right)}{t(\theta)}$

A surprising amount can be learned about the circumstances under which fabrication is optimal by careful examination of the weighted average in (7). In this section we analyze the first component $\frac{c\left(e(\theta),\theta\right)}{t(\theta)}$, the ex post price of incentives.

[26] The weights do add to one $\int_0^\infty t\left(\theta\right)\left(q-p\right)=1$, though they may not all be non negative. It may be helpful in what follows to imagine that $q>p$ everywhere (and p has unit mass at 0).

Litigation mechanisms that do not induce fabrication would seem to be able to deliver incentives at a lower ex post price because truthful evidence production is cheaper than fabrication. In fact, litigation mechanisms not inducing fabrication do have an edge with regard to ex post price, but the reason is more complicated than this.

Consider the two simple *threshold mechanisms* depicted in Figure 1. Under both mechanisms, types below some threshold present no evidence while types above the threshold all present the same fixed level of evidence. Under both mechanisms transfers are zero for types that present no evidence and a fixed positive amount for all types that present the fixed amount of evidence.[27] In both cases, the transfer level is set so that the threshold type is indifferent between presenting no evidence and presenting the fixed level of evidence. This level of transfer is necessary and sufficient for incentive compatibility for such simple threshold mechanisms.[28]

The dashed evidence schedule induces fabrication; it crosses the 45 degree line. The solid schedule does not. How do the schedules' $\frac{c(e(\theta),\theta)}{t(\theta)}$ ratios compare? By choice of their respective fixed transfer levels, both achieve a $\frac{c(e(\theta),\theta)}{t(\theta)}$ ratio of 1 at their respective threshold types. To the right of the threshold $\frac{c(e(\theta),\theta)}{t(\theta)}$ remains constant at one for the no fabrication schedule. As we move to the right in the fabrication mechanism, however, $\frac{c(e(\theta),\theta)}{t(\theta)}$ decreases, as the cost of producing e' decreases while the transfer remains fixed at t'. Eventually, after the fabrication mechanism crosses back under the 45 degree line, $\frac{c(e(\theta),\theta)}{t(\theta)}$ reaches a constant low level of

---

[27] Both direct mechanisms correspond to litigation payoff schedules $t(e)$ that award a fixed transfer amount if the buyer presents at least some threshold level of evidence.

[28] If the threshold type has no positive incentive to present the evidence, then neither does any lower type, whose evidence costs are higher. If the threshold type did have a strictly positive incentive to present the evidence then so would some slightly lower type, thus violating incentive compatibility. A similar pair of arguments applies regarding the threshold type's having no positive incentive to present zero evidence. Therefore, the indifference of the threshold type (no positive incentive either way) is necessary and sufficient for incentive compatibility.

$\frac{c_T e'}{c_T \theta' + c_F(e'-\theta')} = \frac{c_T e'}{c_F e' - (c_F - c_T)\theta'}$, which is less than 1.[29] Thus, $\frac{c(e(\theta),\theta)}{t(\theta)}$ is attains lower levels for the fabrication mechanism than for the no fabrication mechanism.

The intuition here is two fold. First, the issue is not just the cost $c(e(\theta),\theta)$ of evidence, but the cost *per* transfer $t(\theta)$. Higher costs, as caused, for example, by fabrication, allow for higher transfers. Second, of the two mechanisms in the figure only the fabrication mechanism takes advantage of the fact that higher buyer types, who need to fabricate less to present the threshold evidence level, have lower evidence costs. Incentive compatibility is satisfied by setting transfers equal to a threshold type's evidence costs. Higher types then get the same transfer at lower evidence cost, producing a $\frac{c(e(\theta),\theta)}{t(\theta)}$ ratio less than 1.

No fabrication mechanisms can also take advantage of declining evidence costs across the fabrication barrier. Consider, for example, a simple no fabrication mechanisms in which all buyer types present all their available evidence and no more (i.e., $e(\theta)=\theta$) and wherein transfers $t(\theta)$ always equal $c_F\theta$.[30] The full mechanism is depicted by the two arrows in Figure 2. The reader can confirm that this mechanism is incentive compatible.[31]

As shown in the inset box in Figure 2 $\frac{c(e(\theta),\theta)}{t(\theta)}$ equals $\frac{c_T}{c_F}$ at *all* types under this mechanism. This ratio $\frac{c_T}{c_F}$ is lower than the lowest $\frac{c(e(\theta),\theta)}{t(\theta)}$ attained by either threshold schedule in Figure 1.

---

[29] It is also larger than $\frac{c_T}{c_F}$, for future reference.

[30] This direct mechanism corresponds to the litigation payoff schedule $t(e) = c_F e$.

[31] For each buyer type θ, the increase in the transfer for producing more evidence than θ is c_F per unit. The increase in evidence costs is also c_F per unit, because type θ is already producing at the limit of all its available evidence. Therefore, the increase in litigation payoffs is zero for presenting any additional evidence. There is also no gain (indeed there is a positive loss) from presenting less evidence, since θ would lose c_F per unit in transfers and save only c_T per unit in evidence costs.

Indeed, as the proof of the next proposition shows, $\frac{c_T}{c_F}$ is the lowest value that $\frac{c(e(\theta),\theta)}{t(\theta)}$ can attain at any type across all incentive compatible litigation mechanisms.

PROPOSITION 1: *If a litigation mechanism is incentive compatible,*[32] *then at all states* θ*,* $\frac{c(e(\theta),\theta)}{t(\theta)} \geq \frac{c_T}{c_F}$.

As the reader can check, this lower bound is attainted at every θ also by truncated versions of the mechanism in Figure 2, i.e., mechanisms wherein $e(\theta)=\theta$ and $t(\theta)=c_F\theta$ up to some level $\overline{\theta}$, after which both evidence and transfers remain constant at $\overline{\theta}$ and $c_F\overline{\theta}$ respectively. We will call schedules of this sort *canonical no fabrication schedules*. The following proposition, whose proof is omitted, gathers several observations about canonical no fabrication schedules.

PROPOSITION 2: *For every canonical no fabrication schedule,* $\frac{c(e(\theta),\theta)}{t(\theta)} = \frac{c_T}{c_F}$ *at every type* $\theta \in [0,\infty)$. *Moreover, every canonical no fabrication schedule is incentive compatible, and every performance incentive* Δ *that can be generated with some no fabrication schedule can be generated with a canonical no fabrication schedule.*

Conversely, if an incentive compatible schedules does induce fabrication at a particular type, it cannot be at the lower bound for $\frac{c(e(\theta),\theta)}{t(\theta)}$ at that type.

[32] Recall that we have, without loss of generality, restricted attention litigation mechanisms that satisfy $t(0)-c(e(0),0)=0$. See our discussion of normalization at the end of the previous section.

PROPOSITION 3: *Let* (e,t) *be an incentive compatible litigation mechanism. Wherever* $e(\theta)>\theta$*, there* $\frac{c(e(\theta),\theta)}{t(\theta)}>\frac{c_T}{c_F}$.

Taken together, the preceding three propositions establish that with regard to the ex post price of incentives, $\frac{c(e(\theta),\theta)}{t(\theta)}$, no fabrication mechanisms—in particular, canonical no fabrication schedules—are indeed superior to mechanisms that induce fabrication. We can extend this observation to a result about a limiting case in which canonical no fabrication schedules are optimal by revisiting the weighted average in (7).

Consider the limiting case wherein there are no "false positives" in the natural generation of evidence: i.e., the seller's performance precludes the existence of any strictly positive amount of true evidence $\theta>0$. Formally, $p(\theta)=0$ at all positive θ, and the probability of $\theta=0$ given performance is 1. In this case, $\frac{r}{q-p}$, which generally varies in θ, reduces to the constant $1-F(\Delta)$ and the weighted average in (7) is proportional to

$$\int_0^\infty \frac{t(\theta)q}{\Delta}\frac{c(e(\theta),\theta)}{t(\theta)}, \qquad (8)$$

the weighted average of the ratios $\frac{c(e(\theta),\theta)}{t(\theta)}$ taken alone. Therefore, among all litigation mechanisms generating the performance incentive Δ, the canonical no fabrication schedule generating Δ must have minimum cost. Because $\frac{c(e(\theta),\theta)}{t(\theta)}$ is always at its lower bound in these mechanism, the weighted average (8) must also be at this lower bound. Conversely, fabrication is decidedly suboptimal in this no-false-positives case because every mechanism with fabrication puts positive weight on types where $\frac{c(e(\theta),\theta)}{t(\theta)} > \frac{c_T}{c_F}$. We have thus established:[33]

PROPOSITION 4: *If there are no false positives in available evidence—i.e., if* $p(\theta)=0$ *at all positive* θ*, and the probability of* $\theta = 0$ *following performance is 1—then fabrication is suboptimal. In particular, the canonical no fabrication mechanism that generates the incentive* Δ *does so at minimal cost. And no incentive compatible mechanism that induces fabrication with positive ex ante probability generates the incentive at minimal cost.*

E. *Probabilistic efficiency and the possibility of optimal fabrication*

The previous corollary is conceptually edifying because the assumption of no false positives in truly available evidence serves to isolate the importance of $\frac{c(e(\theta),\theta)}{t(\theta)}$ in the efficient provision of performance incentives. But as the weighted average in (7) makes clear, $\frac{c(e(\theta),\theta)}{t(\theta)}$ is not the only consideration in the more empirically interesting case in which false positives do occur and $\frac{r}{q-p}$ varies over θ.

As we shall see, variation in $\frac{r}{q-p}$ gives rise to the possibility that fabrication will be optimal. En route to explaining why this is, we construct a sufficient condition for when variation in $\frac{r}{q-p}$ dominates in comparing the efficiency of two litigation mechanisms that provide the same

[33] The analysis is a bit more complicated than this due to the irregular structure of the probability measure over types given performance. The full argument is included in the proof of the corollary.

performance incentive. To construct this condition, we need to bound variation in $\frac{c(e(\theta),\theta)}{t(\theta)}$. We already have a lower bound for this ratio (Proposition 1). Now we provide an upper lower bound.

PROPOSITION 5: *If a litigation mechanism is incentive compatible, then at all states* θ, $\frac{c(e(\theta),\theta)}{t(\theta)} \le 1$.

As illustrated by the example in Figure 1, this upper bound is tight: it is precisely attained at the threshold type in both mechanisms depicted in that figure.

Now suppose that we are comparing the cost of two incentive compatible litigation mechanisms (e,t) and (e',t') that generate performance incentive Δ. Because the two mechanisms generate the same incentive, asking whether (e',t') is less expensive is the same as asking whether it has the lower weighted average in (7):

$$\int_0^\infty \frac{t'(\theta)(q-p)}{\Delta}\frac{c(e'(\theta),\theta)}{t'(\theta)}\frac{r}{q-p} < \int_0^\infty \frac{t(\theta)(q-p)}{\Delta}\frac{c(e(\theta),\theta)}{t(\theta)}\frac{r}{q-p}? \tag{9}$$

Substituting our upper bound for $\frac{c(e'(\theta),\theta)}{t'(\theta)}$ on the left, and our lower bound on the right, for (9) to hold it suffices that

$$\int_0^\infty \frac{t'(\theta)(q-p)}{\Delta}\frac{r}{q-p} < \frac{c_T}{c_F}\int_0^\infty \frac{t(\theta)(q-p)}{\Delta}\frac{r}{q-p}$$

Both integrals here are weighted averages of $\frac{r}{q-p}$ taken alone, with the weight at each state equal to the extent to which the respective mechanism concentrates the incentive in that state. Each integral is lower, the more the respective mechanism delivers its incentive in states with low $\frac{r}{q-p}$—that is, states whose occurrence from an ex ante perspective (r) is rare relative to the extent to which it is more likely to occur following non performance than performance (q-p). The expression as a whole tells us that if the weighted average $\frac{r}{q-p}$ under (e',t') is not just less than that for (e,t) but less than the fraction $\frac{c_T}{c_F}$ of the weighted average for (e,t), then we may conclude that (e',t') delivers the performance incentive at lower cost, without bothering to examine the relative sizes of each mechanism's $\frac{c(e(\theta),\theta)}{t(\theta)}$ ratios. The following proposition summarizes this result using the notation

$$\rho \equiv \int_0^\infty \frac{t(\theta)(q-p)}{\Delta} \frac{r}{q-p}$$

for the incentive-weighted average probabilistic "price" of the incentive.

PROPOSITION 6: *Given two incentive compatible litigation mechanisms (e,t) and (e',t'), both of which produce performance incentive Δ, in order to conclude that (e',t') is strictly less costly than (e,t) it suffices to find that $\frac{\rho'}{\rho} < \frac{c_T}{c_F}$.*

Proposition 6 allows us to consider yet another limiting case. When fabricating is no more costly than truthfully presenting ($c_T = c_F$), the weighted average probabilistic prices, the ρ's, run the show. The optimal mechanism for producing Δ is the one that is able to most concentrate its transfers at the lowest possible $\frac{r}{q-p}$'s. Thus the case where $c_T = c_F$, wherein the ρ's only determine optimality, is reciprocal to the case of no false positives considered above, wherein $\frac{c(e(\theta),\theta)}{t(\theta)}$ was the only issue.

More generally, Proposition 6 can be used to explain why fabrication might be optimal. Suppose that $\frac{r}{q-p}$ varies widely across states. Then some states are markedly less efficient than others at delivering incentives. That is, while some states that occur often from an ex ante perspective and are roughly as likely to occur following performance as non performance, others rarely occur in any event, but are substantially less rare after performance. This great variability in the probabilistic efficiency of states produces, in turn, an intense desire to load transfers $t(\theta)$ into the efficient, low $\frac{r}{q-p}$ states and out of the inefficient, high $\frac{r}{q-p}$ states. But incentive compatibility requires that evidence production $e(\theta)$ be loaded along with $t(\theta)$. Evidence and transfers must go roughly hand in hand: imbuing an efficient, low $\frac{r}{q-p}$ state with high transfers but not also high evidence production will induce the inefficient, high $\frac{r}{q-p}$ states to grab the same high transfer. Fabrication enters the picture because the desire to load transfers into low $\frac{r}{q-p}$ states may be so intense that it worth letting evidence in these states—which must rise in tandem with transfers—break through the no fabrication bound. That is, in order to increase transfers $t(\theta)$ in these states, we may be willing to raise $e(\theta)$ over θ, even though evidence costs are higher above this barrier. Naturally, this is more likely to be attractive, the lower the premium on fabricated evidence: in particular the closer $\frac{c_T}{c_F}$ is to one.

F. *Special case: increasing likelihood ratio*

To say more about the ρ and the optimality of fabrication, we must add more structure to the relationship between the two probability densities *q* and *p*. Let us then consider the most reasonable regularity: an increasing likelihood ratio $\frac{q}{p}$. As noted earlier, an increasing likelihood ratio corresponds to the case where higher evidentiary states are more informative of non performance. We show, however, that if the likelihood ratio increases with sufficient rapidity, fabrication will be optimal.

1. *The likelihood ratio and* $\frac{r}{q-p}$

The first step is to see the close relationship between the likelihood ratio and $\frac{r}{q-p}$. Some algebra reveals that

$$\frac{r}{q-p} = F(\Delta)\frac{1}{1-\frac{p}{q}} + \left(1-F(\Delta)\right)\frac{1}{\frac{q}{p}-1}.$$

Therefore, wherever $q > p$, $\frac{r}{q-p}$ is strictly *de*creasing in $\frac{q}{p}$. Furthermore, under the assumption that $\frac{q}{p}$ is strictly increasing (and given that *q* and *p* are densities), *q* and *p* must cross precisely once, with *q* cutting *p* from below. Therefore, after some point θ^*, $q > p$, and for lower types than this $q < p$. Therefore, to the right of θ^*, $\frac{r}{q-p}$ is strictly decreasing.

2. *Lower bound on* ρ *for no fabrication mechanisms*

The second step is to establish a lower bound on ρ among all incentive compatible no fabrication mechanisms generating a given performance incentive. This is the work of the following proposition:

PROPOSITION 7: *Suppose that the likelihood ratio* $\frac{q}{p}$ *is strictly increasing. For all incentive compatible no fabrication mechanisms generating performance incentive* Δ*,*

$$\rho \geq \int_{\theta(\Delta)}^{\infty} \frac{c_F\theta(q-p)}{\Delta}\frac{r}{q-p},$$

(10)

where $\theta(\Delta)$ *satisfies* $\int_{\theta(\Delta)}^{\infty} c_F\theta(q-p) = \Delta$.

This proposition is much simpler than it looks. Focus on states to the right of θ^*. There, the fact that $\frac{q}{p}$ is strictly increasing implies that $\frac{r}{q-p}$ is strictly decreasing. When $\frac{r}{q-p}$ is strictly decreasing, ρ is reduced by loading more transfers onto higher types. But an incentive compatible no fabrication schedule is constrained in the extent to which it can load transfers onto high types. *Every* incentive compatible schedule must have $t(\theta) \le c_F e(\theta)$. Otherwise, $\theta = 0$ would be induced to present evidence as intended for θ. For no fabrication incentive compatible mechanisms, which must have $e(\theta) \le \theta$, this specifically means $t(\theta) \le c_F \theta$. Thus, transfers can grow without bound, but at any given type, transfers are bounded from above.

Now broaden the set of mechanisms under consideration to include all those satisfying this last inequality, and not just no fabrication incentive compatible mechanisms. Clearly, minimizing ρ over this containing set means starting $t(\theta) = c_F \theta$ at the highest θ which will still allow us to attain Δ. This is shown in Figure 3, where it also becomes clear that the $\theta(\Delta)$ in the statement of the proposition is precisely this highest possible type. Similarly, the bound in the proposition is the ρ for the transfer schedule shown in the figure.

To be sure, the transfer schedule in the figure, and implicitly, in the statement of the proposition, may not actually be attainable with a no fabrication incentive compatible mechanism. But we can be sure that this transfer schedule's ρ is no greater than that of any no fabrication incentive compatible mechanism, since all such mechanisms must satisfy $t(\theta) \le c_F \theta$ *a fortiori*.

3. *Lowest ρ attainable allowing fabrication*

The next step is investigate the lowest ρ attainable when we allow for fabrication, but still require generation of a positive incentive Δ. Since ρ is a weighted average of $\frac{r}{q-p}$, this can be no lower than the lowest value for $\frac{r}{q-p}$ over states where $q > p$, which under the assumption of an increasing likelihood ratio is $\lim_{\theta\to\infty}\frac{r}{q-p}$. In fact, this lowest possible ρ is attainable (in the limit) while still providing the requisite incentive Δ. Indeed, it is attainable by the kind of simple threshold schedules illustrated in Figure 1.

As noted above, a simple threshold mechanism—which fully defined by the threshold type $\bar{\theta} > 0$, the fixed evidence level, $\bar{e}$ and the fixed transfer level $\bar{t}$—is incentive compatible so long as $\bar{t} = c\left(\bar{e},\bar{\theta}\right)$, i.e., so long as the threshold type is indifferent between producing no evidence and producing the fixed level $\bar{e}$. Thus, $\bar{t}$ can be made arbitrarily large so long as $\bar{e}$ is correspondingly increased. Such a mechanism provides the requisite incentive Δ if $\int_{\bar{\theta}}^{\infty}\bar{t}\left(q-p\right)=\Delta$. Therefore, for any threshold type $\bar{\theta}$, however large, there exists a transfer level $\bar{t}$ large enough to deliver incentive Δ with that threshold type. Moreover, there also exists an evidence level $\bar{e}$ that in combination with $\bar{\theta}$ and $\bar{t}$ defines an incentive compatible litigation mechanism. This choice of larger and larger thresholds is illustrated in Figure 4.

The ρ for such a simple threshold mechanism is $\int_{\bar{\theta}}^{\infty}\frac{\bar{t}(q-p)}{\Delta}\frac{r}{q-p}$. Given that the likelihood ratio is increasing (and $\bar{\theta}$ is greater than θ^*), this weighted average is no larger than $\frac{r(\bar{\theta})}{q(\bar{\theta})-p(\bar{\theta})}$, the largest (and first) item over which the average is taken. Therefore, within the set of threshold mechanisms generating Δ, ρ approaches its absolute minimum $\lim_{\theta\to\infty}\frac{r}{q-p}$ as $\bar{\theta}$ increases (and $\bar{t}$ and $\bar{e}$ increase in tandem).

All told, then, we have proven:

PROPOSITION 8: *Suppose the likelihood ratio* $\frac{q}{p}$ *is increasing. We can find an incentive compatible mechanism generating performance incentive* Δ *whose* ρ *is arbitrarily close to* $\lim_{\theta\to\infty}\frac{r}{q-p}$*, which is in turn a lower bound for* ρ *among all litigation mechanisms incentive compatible or not.*

We now put everything together to produce a result on optimal fabrication when the likelihood ratio is increasing. Suppose that we want to produce incentive Δ at minimum cost and suppose that the following condition holds:

$$\lim_{\theta\to\infty}\frac{r}{q-p} < \frac{c_T}{c_F}\int_{\theta(\Delta)}^{\infty}\frac{c_F\theta(q-p)}{\Delta}\frac{r}{q-p} \tag{11}$$

Then by Proposition 8, there is some incentive compatible mechanism $\left(\hat{e},\hat{t}\right)$ whose $\hat{\rho}$ satisfies

$$\hat{\rho} < \frac{c_T}{c_F}\int_{\theta(\Delta)}^{\infty}\frac{c_F\theta(q-p)}{\Delta}\frac{r}{q-p}$$

Furthermore, we know from Proposition 7 that $\rho \geq \int_{\theta(\Delta)}^{\infty}\frac{c_F\theta(q-p)}{\Delta}\frac{r}{q-p}$ for every no fabrication mechanism generating Δ. Therefore, we may conclude two things. First, the incentive compatible mechanism $\left(\hat{e},\hat{t}\right)$ must be inducing fabrication. Otherwise $\hat{\rho}$ would exceed $\int_{\theta(\Delta)}^{\infty}\frac{c_F\theta(q-p)}{\Delta}\frac{r}{q-p}$, let alone $\frac{c_T}{c_F}\int_{\theta(\Delta)}^{\infty}\frac{c_F\theta(q-p)}{\Delta}\frac{r}{q-p}$. Second, compared to any incentive compatible no fabrication schedule, $\left(\hat{e},\hat{t}\right)$ satisfies $\frac{\hat{\rho}}{\rho} > \frac{c_T}{c_F}$. By Proposition 6, this implies that the fabrication inducing schedule $\left(\hat{e},\hat{t}\right)$ must deliver incentive Δ at lower cost than any no fabrication schedule. We conclude, therefore, that (11) is a sufficient condition for the optimality of fabrication. Rearranging (11) into a more interpretable form produces the follow proposition:

Proposition 9: *Suppose that the likelihood ratio* $\frac{q}{p}$ *is strictly increasing. Given performance incentive* Δ, *consider the weighted average of the percentage drop in* $\frac{r}{q-p}$ *from* $\theta(\Delta)$ *to* ∞

$$\int_{\theta(\Delta)}^{\infty} w_\theta \frac{\frac{r}{q-p} - \inf \frac{r}{q-p}}{\inf \frac{r}{q-p}}$$

where weights are $w_\theta = \int_{\theta(\Delta)}^{\infty} \frac{c_F \theta(q-p)}{\int_{\theta(\Delta)}^{\infty} c_F \theta(q-p)}$.[34] *If the weighted average exceeds the fabrication premium* $\frac{c_F - c_T}{c_T}$, *then fabrication is optimal. That is, it is possible to generate performance incentive* Δ *at strictly lower cost with a mechanism that induces fabrication than with a mechanism that does not.*

G. *Technical discussion*

1. *Wealth constraints*

The construction used to prove the optimality of fabrication in the increasing likelihood case may not be feasible with wealth constraints. If $t(\theta)$ must be less than some amount *W*, we will be prevented from attaining $\rho = \inf \frac{r}{q-p}$ with a simple threshold schedule. But while the particular construction we employ is defeated, the point of analysis is not. If wealth constraints are not too severe relative to the speed at which $\frac{r}{q-p}$ falls, then the lowest attainable ρ given the wealth constraint will still satisfy the sufficient condition Proposition 6. (What is more, the lower bound on ρ for non fabrication schedules will also be higher, as we will have to truncate the transfer schedule in Figure 3.)

[34] These are truly weights in the sense that they add to one and are nonnegative.

2. *Existence of a solution*

When $\frac{r}{q-p}$ goes all the way to zero (and there is no wealth constraint), the condition in Proposition 6 always holds. But Proposition 6 also implies that there is no solution the problem of producing a given Δ at minimal cost. Put another way we may produce Δ at arbitrarily low cost with step functions at higher and higher thresholds. This technical issue also does not defeat our main point. Even if there is technically no solution to the problem of cost minimization, one can *approach* zero cost (i.e., there will be "ε - solutions"), and fabrication schedules will still be superior to no fabrication schedules. Proposition 9 is worded so as to allow for this possibility.

3. *Boundedness*

We have constructed a model wherein variation in $\frac{c(e(\theta),\theta)}{t(\theta)}$ is bounded whereas variation in $\frac{r}{q-p}$. We do not, however, wish to assert that this configuration of variability is a deep structural characteristic of the problem that we are studying. Rather, it is the modeling technique that we have employed to point out what is a deep structural fact: the possibility some amount of fabrication may be part of the optimal means of providing performance incentives.

Figure 1

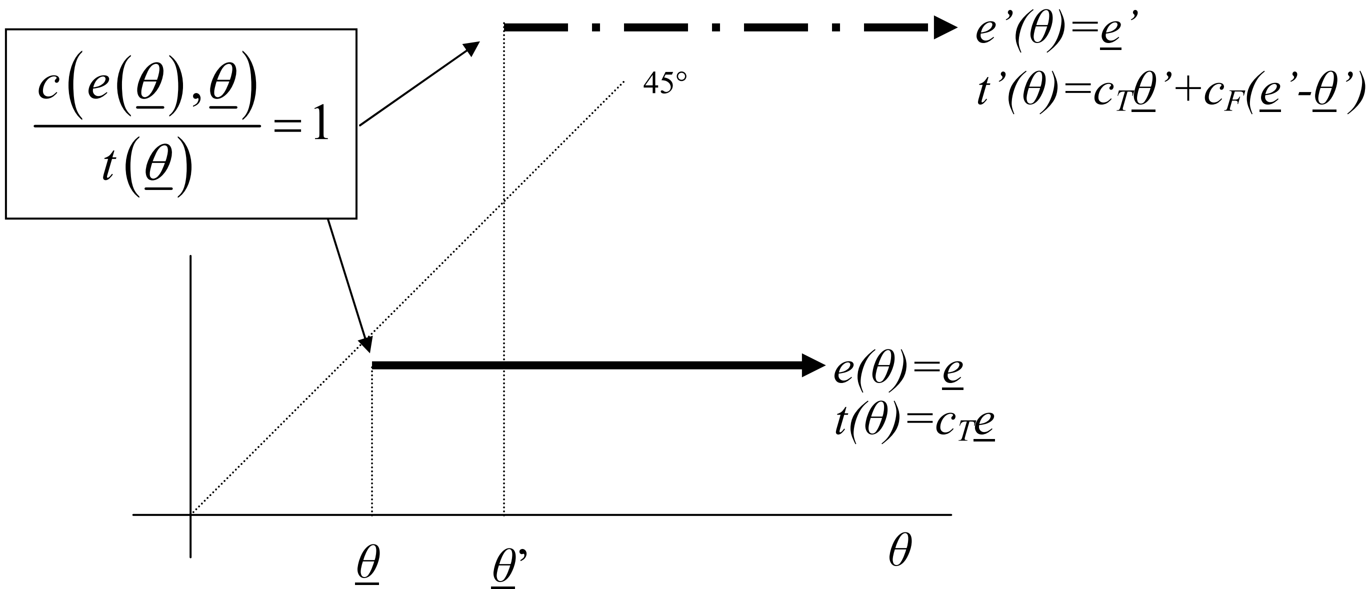

$\frac{c(e(\underline{\theta}),\underline{\theta})}{t(\underline{\theta})}=1$
45°
$e'(\theta)=\underline{e}'$
$t'(\theta)=c_T\underline{\theta}'+c_F(\underline{e}'-\underline{\theta}')$
$e(\theta)=\underline{e}$
$t(\theta)=c_T\underline{e}$
$\underline{\theta}$
$\underline{\theta}'$
θ

Figure 2

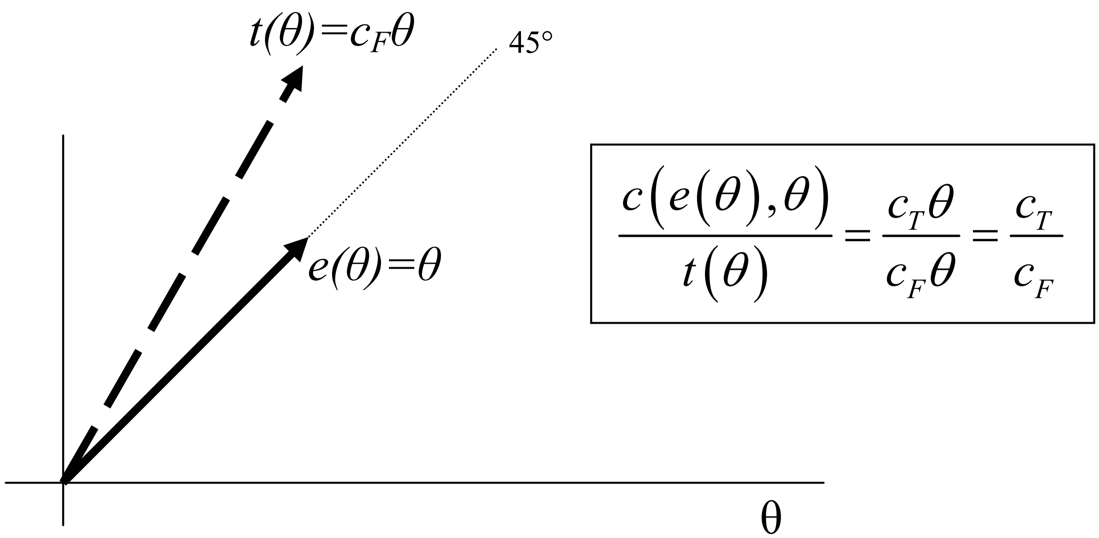

$t(\theta)=c_F\theta$
45°
$e(\theta)=\theta$
θ
$\frac{c(e(\theta),\theta)}{t(\theta)}=\frac{c_T\theta}{c_F\theta}=\frac{c_T}{c_F}$

Figure 3

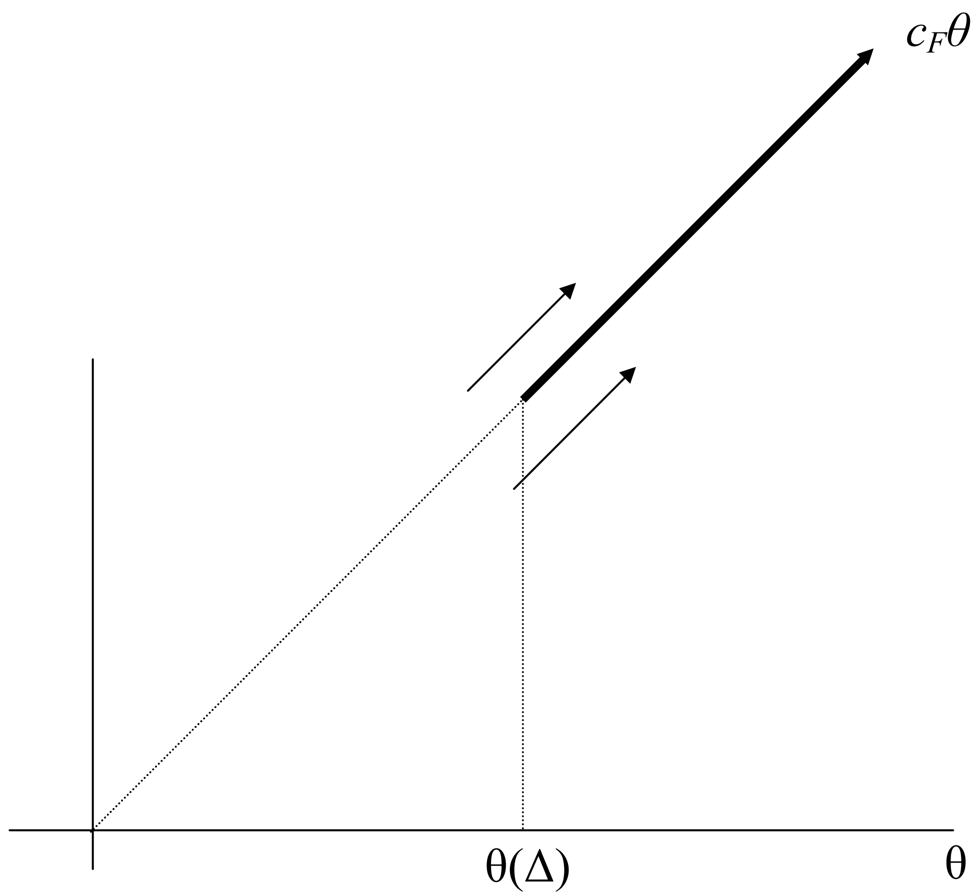

Figure 4

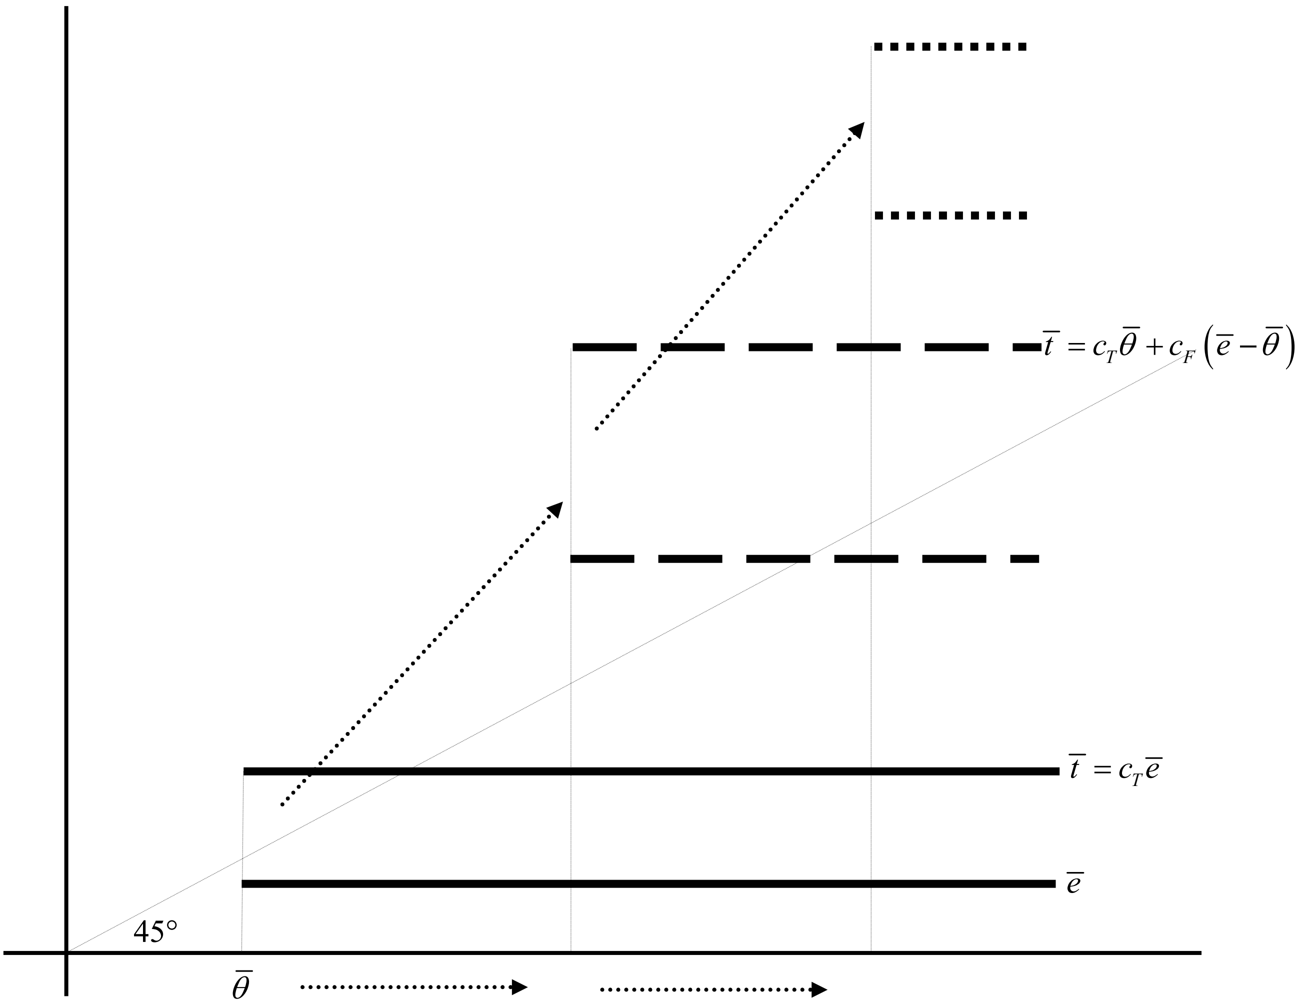

$\overline{t} = c_T\overline{\theta} + c_F\left(\overline{e} - \overline{\theta}\right)$
$\overline{t} = c_T\overline{e}$
$\overline{e}$
45°
$\overline{\theta}$

V. Implications and extensions

Our project was motivated by the received wisdom in contract theory that parties only contract over actions and contingencies that are verifiable. There is considerable ambiguity as to what verifiability means. We argue that the concept itself fails both in theoretical relevance and in empirical support. Courts in civil actions make determinations of complex facts on the basis of balance of probabilities. They are often mistaken in their fact finding because they are misled by one or both parties. Consistent with the literature on accuracy in adjudication, this paper is premised on the assumption that contracting parties are less concerned about whether the court learns the truth, and more concerned about increasing the value of their contract, which depends on performance incentives and enforcement costs rather than the future verification of truth. As a general matter, therefore, we hope to shift the focus of incomplete contract models away from verifiability. In particular, whether a term of performance is contractible should be determined in part by the ex ante incentive efficiency ratio, rather than solely by reference to the court's ability to determine the truth. This approach may reveal that a wider range of terms are contractible than contract theory currently suggests.

Parties in the real world frequently contract over performance standards and contingencies that are patently not verifiable: in the sense that they contemplate costly litigation in which there is a good chance that a court will not find the truth. For example, they often contract for performance that is "reasonable" rather than specifying easily verifiable dimensions of performance and omitting nonverifiable dimensions. This contract term invites the introduction of costly fabricated evidence at trial. However, our analysis suggests that this may indeed improve incentive efficiency. [51]

Our analysis also sheds new light on the all-or-nothing character of burdens of proof.[52] In particular, a plaintiff is not entitled to any remedy until she satisfies the burden of proof threshold. And, once that threshold is passed, the plaintiff receives no additional remedy for surpassing it by a larger margin. Suppose that the threshold entails e^*. It is clear that, depending on the cost of fabrication, at least some of the plaintiffs with truthful evidence that falls short of e^* are likely to be induced to fabricate up to that level and receive the remedy. Assume for the sake of argument that the plaintiff who has no evidence has no incentive to fabricate to the proof threshold. The plaintiffs with truthful evidence greater than e^* cannot increase their recovery by spending more on evidence presentment, and therefore they withhold the balance of their actual evidence. As a result of these evidentiary incentives, the court has much less accurate information as to how much actual evidence exists and consequently, the probability that the alleged event occurred. Without the benefit of our analysis, one might think that this evidentiary pooling would undermine deterrence. Even if it did not, deterrence seems to be attained at a higher evidence cost than a mechanism that deterred fabrication and rewarded the plaintiff incrementally for each unit of truthful evidence she presented. However, we demonstrate in this paper that the burden-of-proof mechanism may in fact lower the evidence cost of the deterrence if the evidentiary states above e* have lower incentive

[51] See, generally, Triantis, Vague Contract Terms, supra note --.

[52] In our model, the court is vulnerable to being misled by fabricated evidence in a predictable manner. In contrast, in other analyses, the court's determination is subject to random error. E.g., Richard Craswell and John E. Calfee, Deterrence and Uncertain Legal Standards, J. Law, Econ. & Organ. 279 (1986).

efficiency ratios than the states below e*. Thus, in at least some cases, this may be a justification for the nonlinear payoff schedule associated with burdens of proof.

The litigation payoff schedule used in our analysis is a modeling construct. Parties may try to condition liquidated damages directly on evidence presentment. This is more likely to succeed if their disputes are arbitrated because the state's rules of evidence and procedure are generally mandatory. Nevertheless, Kaplow suggests that law makers have considerable discretion through the definition of substantive legal obligations.[53] Contracting parties, in particular, can specify distinct or grouped performance obligations and can assign separate or aggregate liquidated damages for the breach of the respective obligations. In addition, the parties may express the obligations more or less like rules or standards.

Our analysis may be extended in several directions. First, we have shown how the parties can exploit variations over evidence states in the incentive efficiency ratio. They might alternatively be able to exploit variations in the difference between the cost of fabricated and truthful evidence production. For example, it might be that the marginal cost of truthful presentation increases faster than the marginal cost of fabrication. Second, we have focused on the evidentiary choices and costs of promisee plaintiffs. We believe that a model where the promisor sues to recover the contract price and must prove performance, would similarly indicate that there are conditions under which a price schedule that induces fabrication by the promisor yields the most efficient performance incentive. Third, the model may be further extended to consider the joint evidentiary strategies of both sides.[54] Fourth, future models should also incorporate the effects of settlement and contract renegotiation.

---

[53] Kaplow, Accuracy in Adjudication, supra note --, at 308, 333, 358 n.358, 363n159.

[54] See generally, Sanchirico, Games, Information, and Evidence Production, supra note ___.

VI. Appendix of Proofs

Proof of Proposition 1: By definition of $c(\bullet)$

$$c\left(e(\theta),\theta\right)\geq c_T e(\theta). \qquad (12)$$

Second, the definition of $c(\bullet)$, incentive compatibility at $\theta=0$, and normalization imply

$$\begin{aligned} &t(\theta)-c_F e(\theta)=t(\theta)-c\left(e(\theta),0\right)\leq t(0)-c\left(e(0),0\right)=0 \\ &\Rightarrow \quad e(\theta)\geq \frac{1}{c_F}t(\theta) \end{aligned} \qquad (13)$$

Combining (12) and (13) yields the result. QED

Proof of Proposition 3: Suppose on the contrary that $c_F c\left(e(\theta),\theta\right)\leq c_T t(\theta)$. As per (13), $t(\theta)\leq c_F e(\theta)$. Combining, we have $c\left(e(\theta),\theta\right)\leq c_T e(\theta)$, which is impossible if $e(\theta)>\theta$. QED

Proof of Proposition 4: By normalization and incentive compatibility $t(\theta)-c\left(e(\theta),\theta\right)\geq 0$. This implies $t(\theta)\geq 0$. Proposition 1 then implies

$$\int_0^\infty \frac{t(\theta)q}{\Delta}\frac{c\left(e(\theta),\theta\right)}{t(\theta)}\geq \frac{c_T}{c_F}\int_0^\infty \frac{t(\theta)q}{\Delta} \qquad (14)$$

With the probability of $\theta=0$ equal to 1 after performance, $\Delta=\int_0^\infty \frac{t(\theta)q}{\Delta}-t(0)$. Therefore,

$$\frac{c_T}{c_F}\int_0^\infty \frac{t(\theta)q}{\Delta}=\frac{c_T}{c_F}\frac{\Delta+t(0)}{\Delta} \qquad (15)$$

Again using $t(\theta)\geq 0$,

$$\frac{c_T}{c_F}\frac{\Delta+t(0)}{\Delta}\geq\frac{c_T}{c_F}. \qquad (16)$$

Combining (14), (15), and (16)

$$\int_0^\infty \frac{t(\theta)q}{\Delta}\frac{c(e(\theta),\theta)}{t(\theta)}\geq\frac{c_T}{c_F}.$$

The canonical litigation mechanism generating Δ achieves this lower bound. But by Proposition 3, and the fact that $t(\theta)\geq 0$, a schedule inducing fabrication with positive ex ante probability cannot. QED

Proof of Proposition 5: By incentive compatibility at θ, the definition of $c(\bullet)$, and the normalization $t(0)-c(e(0),0)=0$

$$t(\theta)-c(e(\theta),\theta)\geq t(0)-c(e(0),\theta)\geq t(0)-c(e(0),0)=0 \Rightarrow t(\theta)\geq c(e(\theta),\theta)$$

QED