MODULAR CONTRACTS

Jason Roderick Donaldson* Giorgia Piacentino† Zichen Zhao‡

April 22, 2026

VERY PRELIMINARY

STILL INCOMPLETE

Abstract

Real-world contracts are written as collections of clauses that courts interpret and evaluate separately. Standard theories, however, suppress much of the internal structure of contracts, representing them as mappings from, typically, output to transfers. In this paper, we develop a theory of contracts starting with the (formal) language they are written in. Clauses, which serve as the units of evaluation by courts, are modeled as groups of sentences in the language. Courts assess clauses based on data, which might not be perfectly observed but might be shuffled, permutations of the true data. We show that, as a result, how sentences are grouped into clauses can determine contractual outcomes, explaining phenomena familiar in practice but difficult to capture in standard models, including loopholes, complexity risk, contractual landmines, incompleteness, and evidentiary conservatism. The framework speaks to how modular structure matters for how contracts are interpreted and how informational and legal risks are allocated.

*USC and CEPR.

†USC, CEPR, ECGI, and NBER.

‡USC.

1 Introduction

Current theories view contracts as incentive schemes and sharing rules, modeled as mappings from, typically, output into transfers. In the real world, however, contracts are expressed in language, not such abstract mappings. That language is divided into modules—clauses—that courts evaluate separately, deciding, for example, whether each covenant in a credit agreement is satisfied one by one (Smith (2006)).[1]

Evaluating clauses separately can lead to unexpected, even perverse phenomena. Examples include loopholes, in which the outcome can change with the order in which clauses are evaluated; errors due to complexity, in which long and intricate clauses lead to mistakes; contractual landmines, in which clauses in the same contract contradict one another; contractual incompleteness, in which relevant contingencies are omitted; and evidentiary conservatism, in which the burden of truth is high relative to what a Bayesian might expect.

Why are contracts modular, with clauses evaluated in isolation from one another? What leads to loopholes, complexity risk, landmines, incompleteness, and conservatism?

In this paper we develop a theory of contracts written in a formal language in which sentences are constructed from atomic premises as in Battigalli and Maggi (2002). In our model, these premises are deemed true or false according to associated data. The data describing each premise are modeled as sequences of zeros and ones, with sufficiently many ones indicating a premise is satisfied. But the data are not observed perfectly; each observed premise could be a permutation of the true one, an assumption we call "shuffled data."

A contract is defined by the units of evaluation: clauses, somewhat akin to Jakob-

[1] See also Schwartz and Scott (2010); Gilson, Sabel and Scott (2014).

sen (2020). The contracting problem is therefore the problem of grouping premises into clauses to implement a policy rule we call the target. Within each clause, the relevant premises are first conjoined and then evaluated, so that each clause is deemed either satisfied (1) or violated (0). The evaluations of the clauses are then themselves conjoined to determine the outcome of the contract (1 or 0).

For concreteness, take corporate debt. The premises could correspond to a leverage covenant $\mathbf{a}_j$ and an earnings covenant $\mathbf{a}_k$, with the target policy that they both be satisfied. To evaluate the target, you must know $a_j^i a_k^i$ for each state i. That would be straightforward to compute if you observed the true data perfectly, i.e. if you knew a_j^i and a_k^i for each i. In practice, however, the data could be shuffled, as leverage observations $\hat{a}_j^i$ come from the balance sheet and earnings observations $\hat{a}_k^i$ from the income statement. The reports being compiled separately, you could be unsure that pairs of observations correspond to the same underlying state.

The contract with the shortest clauses, in which each clause contains a single atomic premise, is called premise-based, and denoted by $\Phi^{\mathcal{A}}$. The contract with a single long clause corresponding to the target policy is called target-based, and denoted by $\Phi^{\mathcal{T}}$. The difference between these contracts is the order of the evaluation and conjunction of clauses, denoted by $[\![\cdot]\!]$ and by $\wedge$, respectively. For illustration, suppose the target policy is $\mathbf{a}_j \wedge \mathbf{a}_k$ for premises $\mathbf{a}_j$ and $\mathbf{a}_k$. The premise-based contract is evaluated as $[\![\Phi^{\mathcal{A}}]\!] = [\![\mathbf{a}_j]\!] \wedge [\![\mathbf{a}_k]\!]$ and the target-based as $[\![\Phi^{\mathcal{T}}]\!] = [\![\mathbf{a}_j \wedge \mathbf{a}_k]\!]$ (where the conjunction of $\mathbf{a}_j$ and $\mathbf{a}_k$ is taken componentwise).

We define a loophole as a set of premises for which the two contracts yield different outcomes when applied to the true data: Changing the order in which premises are evaluated changes the conclusion. For example, suppose $\mathbf{a}_j = \texttt{110}$ and $\mathbf{a}_k = \texttt{101}$ and suppose clauses are deemed true if they include two or more ones: $[\![a_j^1 a_j^2 a_j^3]\!] = \texttt{1}$ iff

$a_j^1 + a_j^2 + a_j^3 \geq \mathtt{2}$. Absent shuffle risk—so the observations are the true data—the premise based contract is evaluated as $[\![\Phi^{\mathcal{A}}]\!] = [\![\mathtt{110}]\!] \wedge [\![\mathtt{101}]\!] = \mathtt{1} \wedge \mathtt{1} = \mathtt{1}$ and the target based as $[\![\Phi^{\mathcal{T}}]\!] = [\![\mathtt{110} \wedge \mathtt{101}]\!] = [\![\mathtt{100}]\!] = \mathtt{0}$. Loopholes capture undesirable deviations of the premise-based contract from the target policy, even absent shuffled data.

Our first main result is that loopholes persist, namely that the premise-based contract can be preferable to the target-based contract even though it is subject to loopholes. The reason is that the premise-based contract is not susceptible to shuffled data. To see why, continue the example above, supposing now that the observations are not perfect: $\hat{\mathbf{a}}_j = \mathtt{110}$ and $\hat{\mathbf{a}}_k = \mathtt{110}$ (so $\mathbf{a}_j$ is not shuffled and $\mathbf{a}_k$ is). The evaluation of the premise-based contract is unchanged as it depends on only the marginals—the number of ones and zeros in each premise—so shuffling has no effect. The evaluation of the target-based contract, in contrast, is flipped. We show that whenever there are loopholes the target-based contract is susceptible to evaluation errors from shuffled data. Because loopholes are relatively unlikely, the drafters of contracts could prefer to accept them, knowing any attempt to eliminate them exposes them to evaluation errors.

Our second main result is the flip side of our first. If dependencies among premises are important, drafters are unwilling to rely on the marginal distributions alone, but need to keep track of the whole joint. They prefer the target-based contract to the premise-based one. That contract is more complex—it tracks all the joint information—and more susceptible to data shuffling as a result. Thus our framework speaks to the costs of complexity as well as when it may be the lesser evil, explaining, e.g., regulatory and tax codes in which complexity is known to lead to compliance errors.

Our third main result concerns redundancy. We call a contract redundant if the same premise appears in multiple clauses. For example, we call $[\![\mathbf{a}_j \wedge \mathbf{a}_k]\!] \wedge [\![\mathbf{a}_k \wedge \mathbf{a}_\ell]\!]$ redundant because $\mathbf{a}_k$ appears in both clauses. We show that such redundant contracts can be preferable to both premise- and target-based contracts. The reason is that it keeps track of pairwise dependencies—between $\mathbf{a}_k$ and both $\mathbf{a}_j$ and $\mathbf{a}_\ell$—without being susceptible to data shuffling that could lead to errors in evaluating $\mathbf{a}_j$ and $\mathbf{a}_\ell$ jointly.

Our fourth result concerns incompleteness. We call a contract incomplete if it omits a premise appearing in the target. We show that incomplete contracts could be preferred to target-based ones, simply because they provide a way to avoid shuffle risk.

For our fifth main result, on conservatism, we extend the model to allow the required burden of truth—the ratio of ones to zeros you need to deem something true—to depend on the clause itself. We show that premise-based contracts are optimally conservative, in that the burden of truth for each clause is higher than for the target itself. That helps explain why contracts, being divided into clauses, typically must satisfy legal standards such as "reasonable doubt" or "probable cause," i.e. why the ostensible burden of truth is so high.

2 Model

Here we set up the model of contracts, represented as collections of clauses in a formal language, evaluated based on imperfect observations of true data.

Atomic premises and data. Let

$$\mathcal{A} \coloneqq \{1, \ldots, J\} \tag{1}$$

be the index set of atomic premises. Each premise $j \in \mathcal{A}$ is represented by a random vector of *data*

$$\mathbf{a}_j = \left(a_j^1, \ldots, a_j^I\right) \in \{0,1\}^I. \tag{2}$$

Clauses and target policy. A *clause* is a nonempty set of indices $\mathcal{M} \subseteq \mathcal{A}$. For any clause $\mathcal{M} \subseteq \mathcal{A}$, write

$$\phi_{\mathcal{M}} \coloneqq \bigwedge_{j \in \mathcal{M}} \mathbf{a}_j. \tag{3}$$

as the conjoined data.[2] We assume that the *target policy* representing the ground truth is one of the clauses denoted as $\mathcal{T} \subseteq \mathcal{A}$ and associated with $\phi_{\mathcal{T}}$.

Observations (shuffled data). We assume that the underlying data may not be perfectly observable, instead, they may be shuffled across premises. For each evaluation of a clause and each premise j appearing in that clause, the *observations* $\hat{\mathbf{a}}_j$ satisfy

$$\hat{a}_j^i = \begin{cases} a_j^i & \text{with probability } 1-e \\ a_j^{\pi_j(i)} & \text{with probability } e \end{cases} \qquad i = 1, \ldots, I, \tag{4}$$

where $e \in [0,1]$ is the shuffle probability and, conditional on a shuffle, π_j is drawn uniformly from the permutations of $\{1, \ldots, I\}$. These draws are independent across

[2] We use conjunction as the starting point because many contract terms are written as such: The legal effect follows only if all stated requirements are met, see for example American Law Institute (1981); Judicial Council of California (2025). Although the baseline model does not include negation as a primitive, the negation of any premise can be included in the set of admissible premises, say $\mathbf{a}_k = \neg \mathbf{a}_j \in \mathcal{A}$.

premises within a clause and across clause evaluations.

The observed clause $\mathcal{M}$ is

$$\hat{\phi}_{\mathcal{M}} := \bigwedge_{j \in \mathcal{M}} \hat{\mathbf{a}}_j, \tag{5}$$

with ith component

$$\hat{\phi}^i_{\mathcal{M}} := \bigwedge_{j \in \mathcal{M}} \hat{a}^i_j, \qquad i = 1, \dots, I. \tag{6}$$

Evaluation. For any $\phi \in \{0,1\}^I$, define

$$[\![\phi]\!] := \begin{cases} 1 & \text{if } \#\{i : \phi^i = 1\} \geq \theta I \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

where $\theta \in [0,1]$ is the threshold share, so the threshold count is θI. In particular, $[\![\phi_{\mathcal{T}}]\!]$ is the target outcome, while $[\![\hat{\phi}_{\mathcal{M}}]\!]$ is the outcome produced by observed clause $\mathcal{M}$.

For singleton clauses, shuffling has no effect on evaluation: if $\mathcal{M} = \{j\}$, then

$$[\![\hat{\phi}_{\{j\}}]\!] = [\![\phi_{\{j\}}]\!], \tag{8}$$

because shuffling only permutes data entries and therefore preserves the number of ones in $\mathbf{a}_j$. More generally, shuffling preserves the marginal distribution of each premise but can change the joint evaluation of a multi-premise clause.[3]

[3]Related work studies inference when variables are observed but correct correspondence is unknown, including the broken-sample and shuffled-data problems; see DeGroot, Feder and Goel (1971); DeGroot and Goel (1980); Hsu, Shi and Sun (2017); Pananjady, Wainwright and Courtade (2017). The same indexing problem also appears in regulated data environments that require accurate linkage across systems, dates, or entities; see U.S. Securities and Exchange Commission (2012); Basel Committee on Banking Supervision (2013); European Central Bank (2024); Office of the Comptroller of the Currency (2017).

Contracts. A *contract* is a family of clause index-sets,

$$\Phi := \{\mathcal{M}_m\}_{m=1}^{M}, \qquad \mathcal{M}_m \subseteq \mathcal{A} \text{ for each } m. \tag{9}$$

We use the same bracket notation for the contract-level evaluation, here defined by

$$[\![\Phi]\!] := \bigwedge_{m=1}^{M} \left[\!\left[\hat{\phi}_{\mathcal{M}_m}\right]\!\right]. \tag{10}$$

Redundancy is allowed: The same premise label $j \in \mathcal{A}$ may appear in several distinct clause index-sets $\mathcal{M}_m$.

Canonical contracts. Two canonical contracts appear frequently below. The *premise-based* contract is

$$\Phi^{\mathcal{A}} := \big\{\{j\}\big\}_{j\in\mathcal{T}}, \tag{11}$$

that is, one singleton clause for each target premise. The *target-based* contract is

$$\Phi^{\mathcal{T}} := \{\mathcal{T}\}, \tag{12}$$

that is, a single clause containing the full target index set.

Objective function. Our objective is to choose Φ to minimize expected cost of departures from the target outcome $[\![\phi_{\mathcal{T}}]\!]$, defined as a weighted sum of the probabilities of type-I and type-II errors:

$$C(\Phi) := c_{\mathrm{I}}\,\mathsf{Pr}[[\![\Phi]\!] > [\![\phi_{\mathcal{T}}]\!]] + c_{\mathrm{II}}\,\mathsf{Pr}[[\![\Phi]\!] < [\![\phi_{\mathcal{T}}]\!]], \tag{13}$$

where $c_{\mathrm{I}}, c_{\mathrm{II}} \geq 0$ weight false positives and false negatives.

3 Results

Below we develop our main results. First we separate loopholes from shuffle risk. Then we characterize the basic tradeoff between premise-based and target-based contracts. Next we study three responses to that tradeoff: Redundancy, which repeats premises across clauses; incompleteness, which omits premises that are too costly due to shuffle risk; and conservatism, which raises clause-level thresholds when premise-based evaluation is too permissive. Together, these results show how imperfect observations of data across clauses can generate several familiar features of real contracts.

3.1 Loopholes and Shuffle Risk

We begin by separating two sources of differences between premise-based and target-based contracting. The first is logical, present even with perfectly observed data. The second arises from shuffled data.

We define the *loophole states* as the set of data realizations where the two canonical contracts disagree even with perfect observations,

$$\mathsf{LH} := \{[\![\Phi^{\mathcal{A}}]\!] \neq [\![\Phi^{\mathcal{T}}]\!] \text{ when } e = 0\}\,. \tag{14}$$

Since the target-based contract's observation coincides with the target clause when $e = 0$, this is equivalently

$$\mathsf{LH} = \{[\![\Phi^{\mathcal{A}}]\!] \neq [\![\phi_{\mathcal{T}}]\!]\}\,. \tag{15}$$

We show that premise-based contract accepts in weakly more states than the target-based contract.

Lemma 1 (Premise-based acceptance weakly dominates target-based acceptance)**.** *Under a common threshold share θ, for every realization of observed premises,*

$$[\![\Phi^{\mathcal{A}}]\!] \geq [\![\Phi^{\mathcal{T}}]\!]. \tag{16}$$

Proof sketch. The full proof is in Appendix Section B.1. The key point is that the observed target clause $\hat{\phi}_{\mathcal{T}}$ is componentwise bounded above by each observed singleton premise, so if the target-based clause clears the threshold, every singleton clause clears it as well. □

The result can be stated as

$$[\![\Phi^{\mathcal{A}}]\!] = 0 \quad \Longrightarrow \quad [\![\Phi^{\mathcal{T}}]\!] = 0, \tag{17}$$

In particular, when $e = 0$,

$$[\![\phi_{\mathcal{T}}]\!] \leq [\![\Phi^{\mathcal{A}}]\!], \tag{18}$$

so premise-based contracts can differ from the target policy only by false positives. In other words, by Lemma 1, premise-based contracts can only admit false positives relative to the target, so

$$\mathsf{LH} = \left\{[\![\Phi^{\mathcal{A}}]\!] = 1, [\![\phi_{\mathcal{T}}]\!] = 0\right\}. \tag{19}$$

We define the *shuffle-risk states* SR as the set of data realizations for which shuffled observations create a positive probability that the two canonical contracts disagree,

$$\mathsf{SR} \coloneqq \left\{\mathsf{Pr}\left[[\![\Phi^{\mathcal{A}}]\!] \neq [\![\Phi^{\mathcal{T}}]\!] \,\middle|\, \{\mathbf{a}_j\}_{j\in\mathcal{T}}\right] > 0\right\}. \tag{20}$$

The two notions, defining LH and SR, are related, but not the same; in fact, loophole states constitute a subset of the shuffle risk states, and a proper subset under some conditions.

Lemma 2. *Fix a shuffle probability $e > 0$, if*

$$|\mathcal{T}| \geq 2, \qquad I \geq 2, \qquad 1 \leq \theta I \leq I - 1, \tag{21}$$

Then

$$\mathsf{LH} \subsetneq \mathsf{SR}. \tag{22}$$

Proof sketch. The basic argument of the proof (in Appendix Section B.2) is the following: On loophole states, premise-based evaluation is unchanged by shuffling, while target-based evaluation can still disagree with it with positive probability, so $\mathsf{LH} \subseteq \mathsf{SR}$. Strictness follows because shuffling can also create different outcomes on some non-loophole states. □

A loophole is a deviation from the target policy caused by replacing a joint clause with separate atomic clauses. Shuffle risk is larger: It also includes states where the target-based contract becomes fragile because conjunction depends on reliable observations across premises, while the premise-based contract remains unaffected because atomic clauses depend only on marginals.

Example 1 (Two-state loophole example ($I = 3$, $\theta = 2/3$))**.** Let $\mathcal{T} = \{1, 2\}$ and $\theta = 2/3$. Consider two realizations.

Loophole realization. Let $\mathbf{a}_1 = \mathtt{110}, \mathbf{a}_2 = \mathtt{101}$. Then $[\![\phi_{\{1\}}]\!] = [\![\phi_{\{2\}}]\!] = \mathtt{1}$, so $[\![\Phi^{\mathcal{A}}]\!] = \mathtt{1}$, but $\phi_{\mathcal{T}} = \mathtt{100}, [\![\phi_{\mathcal{T}}]\!] = \mathtt{0}$. Hence this realization lies in LH.

No-loophole realization. Let $\mathbf{a}_1 = \mathbf{a}_2 = \mathtt{110}$. Then $\phi_{\mathcal{T}} = \mathtt{110}, [\![\Phi^{\mathcal{A}}]\!] = [\![\phi_{\mathcal{T}}]\!] = \mathtt{1}$.

Table 1: Example 1: loophole and benchmark realizations

Realization	$\mathbf{a}_1$	$\mathbf{a}_2$	$[\![\phi_{\{1\}}]\!]$	$[\![\phi_{\{2\}}]\!]$	$[\![\Phi^{\mathcal{A}}]\!]$	$[\![\phi_{\mathcal{T}}]\!]$
Loophole	110	101	1	1	1	0
Non-loophole	110	110	1	1	1	1

■

The set LH captures the kind of loophole that has led to litigation recently, after J. Crew and Serta exploited them in controversial debt restructurings.[4] In the J. Crew case, the borrower relied on a sequence of separately permitted investment baskets to move valuable intellectual-property collateral first into a non-guarantor restricted subsidiary and then into an unrestricted subsidiary, where it could support new financing. In the Serta case, a majority group of lenders used amendment authority and an exchange structure to roll existing debt into new superpriority tranches, leaving nonparticipating lenders effectively subordinated. In both cases, the key feature is not that any single clause was obviously violated when viewed in isolation. It is that a sequence of individually permissible steps opened a transaction path that the joint restriction was supposed to block. States in LH formalize exactly that logic: Each component clause can be satisfied separately, even though the full conjunction fails and the overall restructuring route should not be available.

This result is closely related to the doctrinal-paradox and judgment-aggregation literature, in which premise-based and conclusion-based procedures can disagree on logically connected propositions.[5] It also connects to recent work on loan agreement

[4] See Eat (2018); Ser (2024)

[5] See Kornhauser and Sager (1993); Saari (2001); List and Pettit (2002); Dietrich (2007); Katz (2010); List (2012); Mongin (2012); Miyashita (2021).

loopholes, which studies how separately drafted permissions and exceptions combine into restructuring pathways in practice (Buccola and Nini, 2024; Ayotte and Badawi, 2025).[6]

3.2 Target-based Contract under Shuffled Data

The result in this section shows that the target-based contract removes loopholes, but does so by evaluating conjoined premises on observed data. As we see in the previous result, the very states that generate loopholes under premise-based contract fall within the set of shuffle-risk states. As a consequence, target-based contracting closes the loophole but requires accepting greater fragility to imperfect observations.

First, we introduce notation for the acceptance probability of the target clause, namely the random variable

$$\mathsf{p}_{\mathcal{T}}\left(\{\mathbf{a}_j\}_{j\in\mathcal{T}}\right) := \mathsf{Pr}\left[\left[\!\left[\hat{\phi}_{\mathcal{T}}\right]\!\right] = 1 \,\middle|\, \{\mathbf{a}_j\}_{j\in\mathcal{T}}\right]. \tag{23}$$

We suppress the dependence on $\{\mathbf{a}_j\}_{j\in\mathcal{T}}$ when no confusion arises.

Proposition 1 (Target-based errors under shuffled data). *$\Phi^{\mathcal{T}}$ is error-free if and only if*

$$\mathsf{p}_{\mathcal{T}} = 0 \text{ a.s. on } \{[\![\phi_{\mathcal{T}}]\!] = 0\}, \quad \mathsf{p}_{\mathcal{T}} = 1 \text{ a.s. on } \{[\![\phi_{\mathcal{T}}]\!] = 1\}. \tag{24}$$

In particular, this holds when $e = 0$ or $\mathsf{Pr}(\mathsf{SR}) = 0$.

The result states that, while target-based contract addresses the issue on loophole states, it suffers from shuffle risk unless observations are perfect or the shuffle risk states are empty, that is, shuffling does not alter the conclusions because of sufficient

[6] For practitioner and case-based discussions of such transactions, see Oce (2024); Ser (2024). For a classic form-versus-substance analog, see Gre (1935).

ones or zeros present in the data.

Example 2 (Two-state target-based contract ($I = 3$, $\theta = 2/3$)). Let $\mathcal{T} = \{1, 2\}$ and $\theta = 2/3$. Consider the two realizations from Example 1.

Low realization. Let $\mathbf{a}_1 = \mathtt{110}$, $\mathbf{a}_2 = \mathtt{101}$. Then $[\![\phi_{\mathcal{T}}]\!] = \mathtt{0}$. A direct enumeration shows that the clause accepts with probability one third whenever at least one of the two premises is shuffled. Since each premise is shuffled independently with probability e, $\mathsf{p}_{\mathcal{T}} = (2e - e^2)/3$. Therefore $\mathsf{Pr}\big[[\![\Phi^{\mathcal{T}}]\!] = \mathtt{1} \mid \mathbf{a}_1, \mathbf{a}_2\big] = (2e - e^2)/3$.

High realization. Let $\mathbf{a}_1 = \mathbf{a}_2 = \mathtt{110}$. Then $[\![\phi_{\mathcal{T}}]\!] = \mathtt{1}$. The same enumeration gives acceptance probability one when neither premise is shuffled and one third whenever at least one premise is shuffled, so $\mathsf{p}_{\mathcal{T}} = (1 - e)^2 + (2e - e^2)/3 = 1 - (2(2e - e^2))/3$. Hence $\mathsf{Pr}\big[[\![\Phi^{\mathcal{T}}]\!] = \mathtt{0} \mid \mathbf{a}_1, \mathbf{a}_2\big] = (2(2e - e^2))/3$.

i	$\mathbf{a}_1$	$\mathbf{a}_2$	$\mathbf{a}_1 \wedge \mathbf{a}_2$
1	1	1	1
2	1	0	0
3	0	1	0
	$[\![\mathbf{a}_1]\!] = 1$	$[\![\mathbf{a}_2]\!] = 1$	$[\![\mathbf{a}_1 \wedge \mathbf{a}_2]\!] = 0$

Panel A: true data

i	$\hat{\mathbf{a}}_1$	$\hat{\mathbf{a}}_2$	$\hat{\mathbf{a}}_1 \wedge \hat{\mathbf{a}}_2$
1	1	1	1
2	1	1	1
3	0	0	0
	$[\![\hat{\mathbf{a}}_1]\!] = 1$	$[\![\hat{\mathbf{a}}_2]\!] = 1$	$[\![\hat{\mathbf{a}}_1 \wedge \hat{\mathbf{a}}_2]\!] = 1$

Panel B: one shuffled premise

Figure 1: A permutation can preserve marginals while changing the joint. In Panel A, premise-based evaluation accepts while the target-based clause rejects. In Panel B, a shuffle in one premise leaves the marginal evaluations unchanged but changes the conjunction, which is the core source of shuffle risk.

■

Recent evidence from leveraged-loan markets is consistent with this logic. Credit agreements often preserve headline restrictions while embedding carve-outs, baskets, and other modular permissions that can be combined into states the broad covenant appears to forbid. Those interactions are not always fully internalized ex ante. Ivashina and Vallée (2025) show that weak credit covenants are pervasive and that contracts with greater embedded optionality were repriced after the J. Crew episode, consistent with lenders not fully internalizing cross-clause interactions ex ante. Buccola and Nini (2024) provide complementary evidence from dropdown and uptier transactions: Following the Serta dispute, contracts were revised quickly to restrict uptiers, whereas the response to J. Crew was much more limited, suggesting that some flexibility associated with loophole risk is retained because its net value remains positive. Sophisticated markets, in other words, do not eliminate loopholes whenever they become salient; they often retain modular permissions and patch only the margins that turn out to be especially costly.[7]

This evidence also clarifies why target-based contracting is both appealing and fragile. A target-based clause closes J. Crew- and Serta-style loopholes by evaluating the relevant conjunction directly rather than each premise in isolation. But once implementation requires combining multiple conditions, greater logical precision comes with greater exposure to error whenever the underlying information is difficult to align, verify, or interpret jointly. This is the sense in which the model nests the doctrinal-paradox intuition inside an environment with imperfect observation. Katz-style loopholes capture the different outcomes by premise-based and target-based evaluation, while the shuffled-data analysis explains why loopholes may persist even when the drafter can articulate the target rule: For $e > 0$, eliminating exploitable gaps can make implementation more fragile.

[7] For a complementary law-and-economics account of why sophisticated parties may still leave exploitable seams in complex debt contracts, see Ayotte and Badawi (2025).

3.3 Complexity Risk: Premise-based versus Target-based

The previous two subsections show that: Premise-based contracting avoids shuffle risk errors on atomic clauses but suffers from loopholes. Target-based contracting closes loopholes but is exposed to shuffle risk on the target clause. Consequently, complexity is valuable when the loophole cost is large enough to justify that additional risk, shown in the following result.

Corollary 1 (Complexity risk)**.** *If*

$$c_{\mathrm{I}}\,\mathsf{Pr}(\mathsf{LH}) \text{ is sufficiently large, in the sense that} \tag{25}$$

$$c_{\mathrm{I}}\mathsf{E}[(1-\mathsf{p}_{\mathcal{T}})\mathbf{1}\{\mathsf{LH}\}] > c_{\mathrm{II}}\mathsf{E}[(1-\mathsf{p}_{\mathcal{T}})\mathbf{1}\{\mathsf{SR}\setminus\mathsf{LH}\}]\,. \tag{26}$$

then

$$\Phi^{\mathcal{T}} \succ \Phi^{\mathcal{A}}. \tag{27}$$

In particular, when $e = 0$,

$$\Phi^{\mathcal{T}} \succ \Phi^{\mathcal{A}} \iff \mathsf{Pr}(\mathsf{LH}) > 0. \tag{28}$$

Proof sketch. The argument of the full proof in Appendix Section B.4: The premise-based contract differs from the target only on loophole states, so its cost is the loophole cost. The target-based contract removes those loopholes but introduces shuffle-driven false positives and false negatives. The condition is exactly that the loophole cost avoided by $\Phi^{\mathcal{T}}$ exceeds the new shuffle-risk cost it creates. □

If loopholes are sufficiently likely or sufficiently costly, the drafter should prefer the more complex target-based contract. The benefit of complexity is that it retains joint information; the cost is that long clauses are more fragile under shuffled observations.

Example 3 (Two-state cutoffs)**.** In the two-state world from Examples 1 and 2, let the loophole realization occur with probability α. Then $\mathsf{Pr}_{\mathrm{I}}(\Phi^{\mathcal{A}}) = \alpha, \mathsf{Pr}_{\mathrm{II}}(\Phi^{\mathcal{A}}) = 0, C(\Phi^{\mathcal{A}}) = c_{\mathrm{I}}\alpha$.

For the target-based contract, $\mathsf{Pr}_{\mathrm{I}}(\Phi^{\mathcal{T}}) = \alpha(2e - e^2)/3$, $\mathsf{Pr}_{\mathrm{II}}(\Phi^{\mathcal{T}}) = 2(1-\alpha)(2e - e^2)/3$, so

$$C(\Phi^{\mathcal{T}}) = c_{\mathrm{I}}\alpha\frac{2e - e^2}{3} + c_{\mathrm{II}}(1-\alpha)\frac{2(2e - e^2)}{3}. \tag{29}$$

Hence

$$\Phi^{\mathcal{T}} \succ \Phi^{\mathcal{A}} \quad \Longleftrightarrow \quad \alpha > \frac{2c_{\mathrm{II}}(2e - e^2)}{c_{\mathrm{I}}(3 - 2e + e^2) + 2c_{\mathrm{II}}(2e - e^2)}. \tag{30}$$

The cutoff decreases in c_{I} and increases in e. ■

Evidence on contract text points in the direction that preserving joint contingencies is increasingly costly. Definitions such as EBITDA have become extraordinarily long and permissive, precisely because they serve as common inputs into multiple clauses, so that drafters write ever richer target-like rules into a single contractual object (Badawi et al., 2021). More broadly, Coates (2016) documents sustained growth in the length and linguistic complexity of M&A agreements, and Ganglmair and Wardlaw (2017) link greater contractual complexity in loan agreements to both richer specification and greater renegotiation. Outside private contracting, Beuve, Moszoro and Saussier (2019) show that public contracts become more rigid and formalized when third-party challenge risk is high, and adjacent tax evidence finds that complexity itself generates compliance mistakes (Adhikari, Alm and Harris, 2021). See also Eggleston, Posner and Zeckhauser (2000) on complexity as an endogenous feature of legal design and Loughran and McDonald (2014) on textual complexity in financial disclosure. Taken together, this evidence fits the idea that target-like clauses are often chosen because they preserve economically relevant dependencies, even though they are harder to draft, administer, and verify.[8]

3.4 Redundancy

The previous two subsections compared the two canonical extremes: the premise-based contract $\Phi^{\mathcal{A}}$ uses only singleton clauses, while the target-based contract $\Phi^{\mathcal{T}}$ uses a single long clause. We now study contracts that lie between these two extremes. The question is whether one can improve on both by using several non-atomic clauses, and in particular by allowing some overlap across them.

Recall that a contract is a family of clause index sets,

$$\Phi = \{\mathcal{M}_m\}_{m=1}^{M}. \tag{31}$$

A contract is *nonredundant* if its clauses do not overlap,

$$\mathcal{M}_m \cap \mathcal{M}_{m'} = \varnothing \qquad \text{for all } m \neq m', \tag{32}$$

and it is *redundant* otherwise. This distinction matters because overlap can create several independent opportunities to reject a bad data profile. Premise-based contracting is too weak because it leaves loopholes open, while target-based contracting can be too fragile because one long clause is exposed to shuffle risk. A redundant contract can occupy the middle ground: It may screen undesirable profiles more than once without concentrating in a single clause.

[8]Additional evidence includes discussions of tax-code complexity, accounting-reporting complexity, and operational failures in administering complex financial agreements; see Taxpayer Advocate Service (2023); Hoitash and Hoitash (2018); Chychyla, Leone and Minutti-Meza (2019); Cit (2021); Benzarti (2020); Zwick (2021); Behn, Haselmann and Vig (2022).

The first formal result isolates the mechanism through which redundancy helps.

Proposition 2 (Screening by multiple rejecting clauses)**.** *Assume evaluations are independent across clauses. Fix any underlying data profile and suppose exactly r clauses satisfy*

$$[\![\phi_{\mathcal{M}_m}]\!] = \mathsf{0}, \tag{33}$$

while the remaining clauses satisfy

$$[\![\phi_{\mathcal{M}_m}]\!] = \mathsf{1}. \tag{34}$$

If every clause that should reject falsely accepts with probability $O(e)$, and every clause that should accept falsely rejects with probability $O(e)$, then

$$r \geq 1 \quad \Longrightarrow \quad \mathsf{Pr}\left([\![\Phi]\!] = \mathsf{1} \,\middle|\, \{\mathbf{a}_j\}_{j \in \bigcup_{m=1}^{M} \mathcal{M}_m}\right) = O(e^r), \tag{35}$$

while

$$r = 0 \quad \Longrightarrow \quad \mathsf{Pr}\left([\![\Phi]\!] = \mathsf{0} \,\middle|\, \{\mathbf{a}_j\}_{j \in \bigcup_{m=1}^{M} \mathcal{M}_m}\right) = O(e). \tag{36}$$

Proof sketch. The argument of the proof in Section B.5 is: If r clauses should reject, false acceptance requires all r of them to accept, so it is $O(e^r)$. If all clauses should accept, one clause-level error is enough for false rejection, so that remains $O(e)$. Redundancy helps when bad profiles are screened more than once. □

The proposition identifies the force behind redundancy. What matters is how many clauses independently reject the same bad data profile. If only one clause should reject, false acceptance is first-order in e; if two clauses should reject, it is second-order; and each additional rejecting clause contributes another factor of e. By contrast, on a good data profile one mistaken rejection is enough to block the whole contract, so false rejection remains first order. The corollary therefore identifies a middle ground between the two canonical contracts: Relative to premise-based contracting, overlap closes loopholes by screening bad profiles more than once; relative to target-based contracting, it avoids loading all of the contract onto one long clause sensitive to shuffle risk.

Corollary 2 (Dominance of redundant contract over canonical contracts)**.** *Let $\Phi^{\mathcal{R}}$ be a redundant contract. Suppose the following hold.*

1. *Whenever the target policy should reject, at least two clauses of $\Phi^{\mathcal{R}}$ reject under perfect observation.*

2. *Whenever the target policy should accept, every clause of $\Phi^{\mathcal{R}}$ should accept under perfect observation.*

3. *Clause-level mistakes are first-order: every clause that should reject falsely accepts with probability $O(e)$, and every clause that should accept falsely rejects with probability $O(e)$.*

4. *On loophole data profiles the target-based contract falsely accepts with probability at least proportional to e.*

Then, if $e > 0$ is sufficiently small and $c_{\mathrm{I}} \,\mathsf{Pr}(\mathsf{LH})$ is sufficiently large relative to c_{II},

$$\Phi^{\mathcal{R}} \succ \Phi^{\mathcal{A}} \qquad \textit{and} \qquad \Phi^{\mathcal{R}} \succ \Phi^{\mathcal{T}}. \tag{37}$$

Proof sketch. The argument of the proof in Section B.6: Under the stated conditions, bad profiles require at least two clause-level failures to be falsely accepted, so false positives are second-order. Good profiles still suffer only first-order false negatives.

Hence redundancy can dominate premise-based contracting by reducing loophole error and target-based contracting by avoiding first-order shuffle risk. □

The first condition is what gives redundancy its bite: At least two clauses must reject whenever the target policy should reject, so false acceptance becomes second-order rather than first-order. The second condition rules out contracts that are overly stringent even under perfect observation. The third says that clause-level implementation mistakes are not themselves too severe. The fourth identifies the relevant comparison with $\Phi^{\mathcal{T}}$: On loophole states, a single target-based clause still makes first-order false-acceptance mistakes under shuffling, whereas the redundant contract requires more than one implementation failure.

Observed covenant packages in debt markets look much like this structure. Bräuning, Ivashina and Ozdagli (2025) show that high-yield debt commonly combines incurrence covenants with maintenance covenants tied to the same underlying financial ratio, and that incurrence triggers have substantial real effects. In our language, that is a concrete form of redundancy: The same underlying premise is repeated across distinct clauses with different enforcement consequences, so the contract preserves dependence on a common signal without relying on one all-purpose clause. A useful complementary reference is Nikolaev (2010), who shows that debt contracts with more extensive covenant use are associated with timelier loss recognition, consistent with the broader idea that parties often rely on multiple overlapping screens rather than a single test.

Example 4 (Redundancy trade-off)**.** Fix $I = 3$ and $\theta = 2/3$. Consider the target set $\mathcal{T} = \{1, 2, 3\}$ and compare

$$\Phi^{\mathcal{T}} := \{\{1, 2, 3\}\}, \qquad \Phi^{\mathcal{R}} := \{\{1, 3\}, \{2, 3\}\}. \tag{38}$$

The second contract is redundant because premise 3 appears twice.

For the reject data profile $\mathbf{a}_1 = \mathbf{a}_2 = \mathtt{110}, \mathbf{a}_3 = \mathtt{101}$, the target-based contract has one rejecting clause, whereas the redundant contract has two rejecting clauses. Direct enumeration gives

$$\mathsf{Pr}\left[[\![\Phi^{\mathcal{T}}]\!] = \mathsf{1} \,\middle|\, \mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3\right] = \frac{e}{3} - \frac{e^2}{3} + \frac{e^3}{9}, \tag{39}$$

while

$$\mathsf{Pr}\left[[\![\Phi^{\mathcal{R}}]\!] = \mathsf{1} \,\middle|\, \mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3\right] = \left(\frac{2e - e^2}{3}\right)^2 = \frac{4e^2}{9} + O(e^3). \tag{40}$$

Thus redundancy pushes false acceptance from first order to second order on this bad profile.

For the accept data profile $\mathbf{a}_1 = \mathbf{a}_2 = \mathbf{a}_3 = \mathtt{110}$, both contracts should accept under perfect observation, and false rejection is first-order in either case:

$$\mathsf{Pr}\left[[\![\Phi^{\mathcal{T}}]\!] = \mathsf{0} \,\middle|\, \mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3\right] = 2e - \frac{4e^2}{3} + \frac{2e^3}{9}, \tag{41}$$

while

$$\mathsf{Pr}\left[[\![\Phi^{\mathcal{R}}]\!] = \mathsf{0} \,\middle|\, \mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3\right] = \frac{8e}{3} + O(e^2). \tag{42}$$

So the redundant contract remains fragile on good profiles, but it is substantially more robust on bad ones. ■

A related literature studies contractual landmines and boilerplate persistence. That work shows how interactions among standardized clauses can generate hidden hazards and how such language can survive through template reuse.[9] Our redundancy

[9] See (Choi, Gulati and Scott, 2017; Gelpern, Gulati and Zettelmeyer, 2019; Choi et al., 2023; Scott, Choi and Gulati, 2024a,b); for broader discussions of modular boilerplate and the system-level consequences of clause reuse, see Radin (2006); Rakoff (2006).

result highlights the other side of this mechanism: Repeated or overlapping clauses need not be a drafting mistake. When observed data are shuffled, they can improve evaluations by making the same bad profile fail in more than one clause.

3.5 Incompleteness and Relevance

We now study when it is optimal to leave target-relevant material out of the contract. As before, adding a clause has two effects: It can screen out false positives, but it can also create new false negatives when the clause must be evaluated on shuffled data. Incompleteness is therefore not merely a drafting failure as it can be an optimal response to shuffle risk.

Say that a contract $\Phi^{\mathcal{I}}$ is *incomplete with respect to $\mathcal{T}$* if

$$j \notin \bigcup_{m=1}^{M} \mathcal{M}_m \qquad \text{for some } j \in \mathcal{T}. \tag{43}$$

Before we present the incompleteness result, we first notice that adding a clause weakly decreases the chance of a positive conclusion. Take any baseline contract Φ and any additional clause $\mathcal{N} \subseteq \mathcal{A}$, and define the new contract

$$\Phi' := \Phi \cup \{\mathcal{N}\}. \tag{44}$$

Then

$$[\![\Phi']\!] = [\![\Phi]\!] \wedge [\![\hat{\phi}_{\mathcal{N}}]\!]. \tag{45}$$

So adding a clause can only turn acceptances into rejections.

Proposition 3 (Adding a clause: screening versus shuffle risk)**.** *Given a baseline* *contract* Φ*, it is strictly preferred to* Φ' *whenever*

$$c_{\mathrm{I}} \,\mathsf{Pr}\Big[[\![\Phi]\!] = 1,\ [\![\phi_{\mathcal{T}}]\!] = 0,\ [\![\hat{\phi}_{\mathcal{N}}]\!] = 0\Big] < c_{\mathrm{II}} \,\mathsf{Pr}\Big[[\![\Phi]\!] = 1,\ [\![\phi_{\mathcal{T}}]\!] = 1,\ [\![\hat{\phi}_{\mathcal{N}}]\!] = 0\Big]. \quad (46)$$

Proof sketch. The argument of the proof in Section B.7: Adding a clause only turns some baseline acceptances into rejections. This is beneficial where the target should reject and harmful where the target should accept. The condition simply says that the weighted gain in false positives exceeds the weighted increase in false negatives. □

A new clause is useful only if it screens enough bad baseline acceptances without vetoing too many good baseline acceptances. That is the common logic behind incompleteness and relevance. Omitting a clause is therefore not necessarily a mistake. If the omitted clause is sufficiently noisy under shuffled data, if false negatives are sufficiently costly, and if the baseline false-positive risk is limited, then enforcing the additional clause can make the contract worse rather than better. In this sense, incompleteness arises not only because parties fail to foresee contingencies, but also because conditioning on an additional premise may require matching, verification, or state identification that is too unreliable.

Evidence on incompleteness points in the same direction. Nyarko (2021) shows that important clauses are often omitted for organizational and path-dependent reasons, while Demerjian (2011) finds that the use of balance-sheet covenants declines when accounting measurement becomes less reliable. In procurement, Bajari, Houghton and Tadelis (2014) show that contracts remain incomplete when adaptation is valuable and ex ante specification is costly, and Eisenberg and Miller (2009) document that even relatively standard jurisdictional clauses are often left unspecified. These findings are consistent with our interpretation of incompleteness as a response not only to unforeseeable contingencies, but also to difficulty in verifying or aligning the information needed to condition on a premise cleanly. Omitting a premise can therefore be a rational way to avoid errors rather than a simple failure of foresight. More broadly, this is the same trade-off as in Holmström–Milgrom-style multitask settings: when measurement is distortionary, it can be optimal to place less contractual weight on the affected dimension.[10]

Corollary 3 (When uninformative off-target clauses should not be used)**.** *Fix $j \notin \mathcal{T}$ and define*

$$\Phi^{(j)} := \Phi \cup \{\{j\}\}. \tag{47}$$

Suppose there exists $q \in (0,1)$ such that

$$\mathsf{Pr}\big(\llbracket \phi_{\{j\}} \rrbracket = 1 \,\big|\, \llbracket \Phi \rrbracket,\ \llbracket \phi_{\mathcal{T}} \rrbracket\big) = q \qquad a.s. \tag{48}$$

In particular, if

$$c_{\mathrm{II}}\, \mathsf{Pr}[\llbracket \Phi \rrbracket = 1,\ \llbracket \phi_{\mathcal{T}} \rrbracket = 1] \geq c_{\mathrm{I}}\, \mathsf{Pr}_{\mathrm{I}}(\Phi), \tag{49}$$

then adding the off-target atom weakly increases cost.

Example 5 (Bundled-clause incompleteness)**.** Fix $I = 3$, $\theta = 2/3$, and $\mathcal{T} = \{1, 2, 3\}$. Suppose the feasible menu includes only

$$\Phi_0 := \{\{1\}\}, \qquad \Phi_1 := \{\{1\}, \{2, 3\}\}. \tag{50}$$

Thus the relevant pair $\{2, 3\}$ can either be imposed as a bundle or omitted altogether.

[10] See also Allen and Gale (1992); Holmstrom and Milgrom (1991); Townsend (1979); Grossman and Hart (1986); Hart and Moore (1990).

Suppose there are two realizations of the data, a low realization and a high realization, occurring with probabilities α and $1-\alpha$, respectively. In the high realization $\mathbf{a}_1 = \mathbf{a}_2 = \mathbf{a}_3 = \texttt{110}$, and $[\![\phi_\mathcal{T}]\!] = \texttt{1}$. In the low realization $\mathbf{a}_1 = \texttt{110}$, $\mathbf{a}_2 = \texttt{101}$, $\mathbf{a}_3 = \texttt{011}$, and $[\![\phi_\mathcal{T}]\!] = \texttt{0}$. In both realizations, $[\![\phi_{\{1\}}]\!] = \texttt{1}$.

A direct enumeration gives, for the bundled clause $\{2,3\}$,

$$\mathsf{Pr}\left[[\![\hat{\phi}_{\{2,3\}}]\!] = \texttt{1} \,\middle|\, \text{low realization}\right] = \frac{2e-e^2}{3}, \tag{51}$$

and

$$\mathsf{Pr}\left[[\![\hat{\phi}_{\{2,3\}}]\!] = \texttt{0} \,\middle|\, \text{high realization}\right] = \frac{2(2e-e^2)}{3}. \tag{52}$$

Therefore $\mathsf{Pr}_{\mathrm{I}}(\Phi_0) = \alpha$, $\mathsf{Pr}_{\mathrm{II}}(\Phi_0) = 0$. Meanwhile, $\mathsf{Pr}_{\mathrm{I}}(\Phi_1) = \alpha(2e-e^2)/3, \mathsf{Pr}_{\mathrm{II}}(\Phi_1) = (1-\alpha)(2(2e-e^2))/3$.

The incomplete contract Φ_0 is preferred to Φ_1 if and only if

$$\alpha < \frac{2c_{\mathrm{II}}(2e-e^2)}{c_{\mathrm{I}}(3-2e+e^2)+2c_{\mathrm{II}}(2e-e^2)}. \tag{53}$$

In other words, when the bad realization is sufficiently rare and shuffle risk is sufficiently important, it is optimal to omit a target-relevant bundled clause. ■

3.6 Conservatism within Premise-based Contracts

We now study an extension in which the contract remains premise-based, but the premise-level thresholds are allowed to differ from the target threshold. The question is when it is optimal to impose a stricter burden of proof at the premise level than the target itself requires.

Fix $I \geq 2$ and a nonempty target set $\mathcal{T}$ with $|\mathcal{T}| \geq 2$. Let the target clause be

$$\phi_{\mathcal{T}} = \bigwedge_{j \in \mathcal{T}} \mathbf{a}_j. \tag{54}$$

Let the target threshold share be $\theta_{\mathcal{T}} \in \{0, 1/I, \ldots, 1\}$,[11] so that the target outcome is

$$[\![\phi_{\mathcal{T}}]\!] = \mathbf{1}\{s_{\mathcal{T}} \geq \theta_{\mathcal{T}} I\}, \qquad s_{\mathcal{T}} \coloneqq \sum_{i=1}^{I} \phi_{\mathcal{T}}^i. \tag{55}$$

For each premise $j \in \mathcal{T}$, define the marginal score

$$s_j \coloneqq \sum_{i=1}^{I} a_j^i. \tag{56}$$

Consider the family of premise-based contracts with free premise-level threshold shares

$$[\![\Phi^{\mathcal{A}}_{(\theta_j)_{j \in \mathcal{T}}}]\!] \coloneqq \bigwedge_{j \in \mathcal{T}} \mathbf{1}\{s_j \geq \theta_j I\}, \qquad (\theta_j)_{j \in \mathcal{T}} \in \{0, 1/I, \ldots, 1\}^{|\mathcal{T}|}. \tag{57}$$

A premise-based implementation is *conservative relative to the target threshold share* if

$$\theta_j > \theta_{\mathcal{T}} \qquad \text{for every } j \in \mathcal{T}. \tag{58}$$

Restrict attention to symmetric premise-level thresholds $\theta_j = \theta$ for all $j \in \mathcal{T}$, and write

$$\Phi^{\mathcal{A}}_{\theta} \coloneqq \Phi^{\mathcal{A}}_{(\theta, \ldots, \theta)}. \tag{59}$$

[11] For fixed I, this grid restriction is without loss of generality. Since all scores are integer-valued, any $\theta \in [0, 1]$ induces the same acceptance rule as $\bar{\theta} \coloneqq \lceil \theta I \rceil / I \in \{0, 1/I, \ldots, 1\}$:

$$\mathbf{1}\{s \geq \theta I\} = \mathbf{1}\{s \geq \lceil \theta I \rceil\} = \mathbf{1}\{s \geq \bar{\theta} I\}.$$

Thus only the induced integer cutoff matters, and the restriction to grid values is only for notational simplicity.

Define the acceptance event

$$\mathsf{ACC}_\theta := \{[\![\Phi_\theta^{\mathcal{A}}]\!] = \mathsf{1}\}\,, \tag{60}$$

and, for $\theta \in \{0, 1/I, \ldots, (I-1)/I\}$, the marginal acceptance band

$$\mathsf{M}_\theta := \mathsf{ACC}_\theta \setminus \mathsf{ACC}_{\theta+1/I}. \tag{61}$$

Thus M_θ is the set of realizations that are accepted at threshold share θ but rejected at threshold share $\theta + 1/I$.

Proposition 4 (Marginal condition for conservatism)**.** *Fix $\theta \in \{0, 1/I, \ldots, (I-1)/I\}$ and suppose $\mathsf{Pr}(\mathsf{M}_\theta) > 0$. Raising the symmetric premise-level threshold share from θ to $\theta + 1/I$ strictly reduces cost if and only if*

$$\mathsf{Pr}([\![\phi_{\mathcal{T}}]\!] = \mathsf{1} \,|\, \mathsf{M}_\theta) < \frac{c_{\mathrm{I}}}{c_{\mathrm{I}} + c_{\mathrm{II}}}. \tag{62}$$

Moreover, if $\theta_{\mathcal{T}} < 1$ and (62) holds at $\theta = \theta_{\mathcal{T}}$, then every optimal symmetric premise-level threshold share satisfies

$$\theta^\star \geq \theta_{\mathcal{T}} + \frac{1}{I}. \tag{63}$$

Hence the optimal symmetric premise-based contract is conservative relative to the target threshold share.

Proof sketch. The argument of the proof in Section B.9: Raising the premise threshold only matters on the marginal set M_θ. On that set it removes false positives but also creates false negatives. The stated condition is exactly that the first effect dominates the second. □

Intuitively, the event M_θ consists of realizations that barely satisfy the premise-based standard at threshold share θ. Condition (62) says that one should tighten the atomic standard whenever these marginal acceptances are sufficiently likely to be false positives. This is where conservatism comes from in the model. Premise-based evaluation is insulated from shuffle risk because it uses only marginals, but that same insulation can make it too permissive: Weak premise-level evidence may pass clause by clause even when it is not sufficiently informative about the target clause. Raising the atomic threshold then sacrifices some true positives in order to screen out weak acceptances near the margin. In that sense, the optimal premise-level standard can be stricter than the target threshold itself.

Observed contracts often look exactly like this. Debt covenants are frequently written and enforced with slack, tripwires, and threshold adjustments that create asymmetric trigger regions rather than neutral cutoffs (Dichev and Skinner, 2002; Murfin, 2012). Roberts and Sufi (2009) show that control rights shift well before formal payment default, and Li, Vasvari and Wittenberg-Moerman (2016) document covenant structures with moving earnings-based thresholds over the life of the loan.

This result is also related to the literature on burdens of proof and to the accounting literature on conservatism.[12] In that literature, the common theme is that decision rules often require stronger evidence on each issue than the underlying target criterion would suggest. The next example shows this in the simplest possible environment.

Example 6 (Two-realization illustration with $I = 3$)**.** For transparency, keep $|\mathcal{T}| = 2$, with $\mathcal{T} = \{1, 2\}$, and fix $\theta_{\mathcal{T}} = 2/3$. Suppose there are two realizations of the data, a low realization and a high realization, occurring with probabilities α and $1 - \alpha$,

[12]See Kaplow (2012); Basu (1997); Win (1970); Add (1979); San (1982); Watts (2003).

respectively.

In the low realization, $\mathbf{a}_1 = \texttt{110}$, $\mathbf{a}_2 = \texttt{101}$. So $\phi_{\mathcal{T}} = \texttt{100}$, $[\![\phi_{\mathcal{T}}]\!] = \texttt{0}$.

In the high realization, $\mathbf{a}_1 = \texttt{110}$, $\mathbf{a}_2 = \texttt{110}$. So $\phi_{\mathcal{T}} = \texttt{110}$, $[\![\phi_{\mathcal{T}}]\!] = \texttt{1}$.

In both realizations, $s_1 = s_2 = 2$. Hence $[\![\Phi^{\mathcal{A}}_{2/3}]\!] = \texttt{1}$ in both realizations, hence $C(\Phi^{\mathcal{A}}_{2/3}) = c_{\mathrm{I}}\alpha$.

By contrast, $[\![\Phi^{\mathcal{A}}_{1}]\!] = \texttt{0}$ in both realizations, hence $C(\Phi^{\mathcal{A}}_{1}) = c_{\mathrm{II}}(1-\alpha)$.

Therefore

$$\Phi^{\mathcal{A}}_{1} \succ \Phi^{\mathcal{A}}_{2/3} \quad \Longleftrightarrow \quad c_{\mathrm{II}}(1-\alpha) < c_{\mathrm{I}}\alpha \quad \Longleftrightarrow \quad \alpha > \frac{c_{\mathrm{II}}}{c_{\mathrm{I}} + c_{\mathrm{II}}}. \tag{64}$$

So the optimal symmetric premise-based implementation is conservative:

$$\theta^{\star} = 1 > \theta_{\mathcal{T}} = \frac{2}{3} \tag{65}$$

whenever $\alpha > (c_{\mathrm{II}})/(c_{\mathrm{I}} + c_{\mathrm{II}})$.

■

4 Conclusion

We study contracts as they are written and applied in practice: in language, through clauses, and on imperfect data. Ostensibly equivalent rules need not perform the same way, because clause architecture matters once clauses are evaluated separately and observations can be shuffled versions of the data. The model shows how premise-based contracts can remain attractive even though they leave loopholes, whereas target-based contracts can close those loopholes but become more fragile. It establishes how redundancy and incompleteness can be optimal responses to noisy information as well as how clause-by-clause contracting can be optimally conservative. Once contracts are written and evaluated clause by clause, these features—loopholes, complexity risk, landmines, incompleteness, and conservatism—need not arise as drafting mistakes, but can instead be rational responses to the way information is observed and used.

A Notation

Notation	Meaning
$\mathcal{A} = \{1, \ldots, J\}$	Index set of atomic premises
J	Number of atomic premises
I	Number of data entries in each premise vector
$\mathbf{a}_j$	Data describing premise j
a_j^i	Data entry i of premise j
$\mathcal{M} \subseteq \mathcal{A}$	Clause index set
$\mathcal{T} \subseteq \mathcal{A}$	Target index set
$\hat{\mathbf{a}}_j$	Observation of premise $\mathbf{a}_j$ under possible shuffling
π_j	Random permutation of $\{1, \ldots, I\}$ used when premise j is shuffled
$e \in [0, 1]$	Shuffle probability
$\hat{\phi}_{\mathcal{M}} = \bigwedge_{j \in \mathcal{M}} \hat{\mathbf{a}}_j$	Observed clause
$\hat{\phi}^i_{\mathcal{M}} = \bigwedge_{j \in \mathcal{M}} \hat{a}^i_j$	Entry i of observed clause $\hat{\phi}_{\mathcal{M}}$
θ	Acceptance-threshold share (the count threshold is θI)
$[\![\phi]\!]$	Threshold evaluation: equals 1 iff $\#\{i : \phi^i = 1\} \geq \theta I$
M	Number of clauses in the contract
$\Phi = \{\mathcal{M}_m\}_{m=1}^M$	Contract as a family of clause index-sets
$[\![\Phi]\!] = \bigwedge_{m=1}^M \left[\!\left[\hat{\phi}_{\mathcal{M}_m}\right]\!\right]$	Contract-level outcome
$\Phi^{\mathcal{A}} = \{\{j\}\}_{j \in \mathcal{T}}$	Premise-based contract
$\Phi^{\mathcal{T}} = \{\mathcal{T}\}$	Target-based contract
$c_{\mathrm{I}}, c_{\mathrm{II}} \geq 0$	Weights on false positives and false negatives
$C(\Phi)$	Expected cost of contract Φ

Table 2: Notation

B Omitted Proofs

B.1 Proof of Lemma 1

Proof. For each i and each $j \in \mathcal{T}$,

$$\hat{\phi}^i_{\mathcal{T}} = \prod_{k \in \mathcal{T}} \hat{a}^i_k \leq \hat{a}^i_j. \tag{66}$$

Summing over $i = 1, \dots, I$ yields

$$\sum_{i=1}^{I} \hat{\phi}^i_{\mathcal{T}} \leq \sum_{i=1}^{I} \hat{a}^i_j \qquad \text{for every } j \in \mathcal{T}. \tag{67}$$

Therefore, if the target-based contract accepts, then

$$\sum_{i=1}^{I} \hat{\phi}^i_{\mathcal{T}} \geq \theta I \quad \Longrightarrow \quad \sum_{i=1}^{I} \hat{a}^i_j \geq \theta I \qquad \text{for every } j \in \mathcal{T}. \tag{68}$$

Hence

$$[\![\Phi^{\mathcal{T}}]\!] = \mathtt{1} \quad \Longrightarrow \quad \left[\!\!\left[\hat{\phi}_{\{j\}}\right]\!\!\right] = \mathtt{1} \qquad \text{for every } j \in \mathcal{T}, \tag{69}$$

which implies

$$[\![\Phi^{\mathcal{T}}]\!] = \mathtt{1} \quad \Longrightarrow \quad [\![\Phi^{\mathcal{A}}]\!] = \mathtt{1}. \tag{70}$$

This proves

$$[\![\Phi^{\mathcal{A}}]\!] \geq [\![\Phi^{\mathcal{T}}]\!]. \tag{71}$$

When $e = 0$, we have $[\![\Phi^{\mathcal{T}}]\!] = [\![\phi_{\mathcal{T}}]\!]$, so

$$[\![\Phi^{\mathcal{A}}]\!] \geq [\![\phi_{\mathcal{T}}]\!]. \tag{72}$$

Thus premise-based contracting can differ from the target only by false positives. □

B.2 Proof of Lemma 2

Proof. Take any realization in LH. By definition,

$$[\![\Phi^{\mathcal{A}}]\!] = \mathsf{1}, \qquad [\![\phi_{\mathcal{T}}]\!] = \mathsf{0}. \tag{73}$$

The premise-based contract depends only on atomic clauses, so by singleton invariance from Equation (8),

$$[\![\Phi^{\mathcal{A}}]\!] = \bigwedge_{j\in\mathcal{T}} [\![\phi_{\{j\}}]\!] \tag{74}$$

regardless of shuffling.

Under the shuffle model, there is positive probability that every observed premise equals its underlying premise. Hence

$$\mathsf{Pr}[\hat{\mathbf{a}}_j = \mathbf{a}_j \text{ for all } j \in \mathcal{T} \mid \{\mathbf{a}_j\}_{j\in\mathcal{T}}] > 0. \tag{75}$$

On that event,

$$[\![\Phi^{\mathcal{T}}]\!] = \left[\!\!\left[\hat{\phi}_{\mathcal{T}}\right]\!\!\right] = [\![\phi_{\mathcal{T}}]\!] = \mathsf{0}. \tag{76}$$

Combining (73), (74), and (76), we obtain

$$\mathsf{Pr}[[\![\Phi^{\mathcal{A}}]\!] \neq [\![\Phi^{\mathcal{T}}]\!] \mid \{\mathbf{a}_j\}_{j\in\mathcal{T}}] > 0. \tag{77}$$

Therefore the realization lies in SR, proving $\mathsf{LH} \subseteq \mathsf{SR}$.

We now prove strictness under (21). Choose two distinct indices $j, k \in \mathcal{T}$. Let

$$\mathbf{v} \coloneqq (\underbrace{1, \ldots, 1}_{\lceil \theta I \rceil}, \underbrace{0, \ldots, 0}_{I - \lceil \theta I \rceil}) \in \{0, 1\}^I. \tag{78}$$

Define a realization by setting

$$\mathbf{a}_\ell = \mathbf{v} \qquad \text{for every } \ell \in \mathcal{T}. \tag{79}$$

Then

$$[\![\Phi^{\mathcal{A}}]\!] = \mathsf{1}, \qquad [\![\phi_{\mathcal{T}}]\!] = \mathsf{1}, \tag{80}$$

so this realization is not in LH.

Next, let π be a permutation that swaps one coordinate where $\mathbf{v}$ equals 1 with one coordinate where $\mathbf{v}$ equals 0. Since $1 \le \theta I \le I - 1$, such coordinates exist. Consider the observation profile in which

$$\hat{\mathbf{a}}_j = \mathbf{v}, \qquad \hat{\mathbf{a}}_k = \mathbf{v}^\pi, \qquad \hat{\mathbf{a}}_\ell = \mathbf{v} \quad \text{for all } \ell \in \mathcal{T} \setminus \{j, k\}. \tag{81}$$

This profile has positive conditional probability because $e > 0$ and the shuffle draws assign positive probability to each permutation.

By construction, the vector $\mathbf{v}^\pi$ agrees with $\mathbf{v}$ on exactly $\lceil \theta I \rceil - 1$ coordinates where $\mathbf{v}$ has a 1. Hence

$$\hat{\phi}_{\mathcal{T}} = \bigwedge_{\ell \in \mathcal{T}} \hat{\mathbf{a}}_\ell \tag{82}$$

has exactly $\lceil \theta I \rceil - 1$ ones, so

$$[\![\Phi^{\mathcal{T}}]\!] = \left[\!\!\left[\hat{\phi}_{\mathcal{T}}\right]\!\!\right] = \mathsf{0}. \tag{83}$$

At the same time, atom evaluations are unaffected by shuffling, so

$$[\![\Phi^{\mathcal{A}}]\!] = \mathsf{1}. \tag{84}$$

Therefore

$$\mathsf{Pr}[[\![\Phi^{\mathcal{A}}]\!] \neq [\![\Phi^{\mathcal{T}}]\!] \mid \{\mathbf{a}_\ell\}_{\ell\in\mathcal{T}}] > 0, \tag{85}$$

so this realization lies in $\mathsf{SR} \setminus \mathsf{LH}$. Hence (22) holds. □

B.3 Proof of Proposition 1

Proof. For the target-based contract,

$$\mathsf{Pr}_{\mathrm{I}}(\Phi^{\mathcal{T}}) = \mathsf{E}[\mathsf{p}_{\mathcal{T}}\, \mathbf{1}\{[\![\phi_{\mathcal{T}}]\!] = \mathsf{0}\}]\,, \tag{86}$$

and

$$\mathsf{Pr}_{\mathrm{II}}(\Phi^{\mathcal{T}}) = \mathsf{E}[(1-\mathsf{p}_{\mathcal{T}})\, \mathbf{1}\{[\![\phi_{\mathcal{T}}]\!] = \mathsf{1}\}]\,. \tag{87}$$

Since $[\![\Phi^{\mathcal{T}}]\!] = \left[\!\!\left[\hat{\phi}_{\mathcal{T}}\right]\!\!\right]$,

$$\mathsf{Pr}_{\mathrm{I}}(\Phi^{\mathcal{T}}) = \mathsf{Pr}\left[\left[\!\!\left[\hat{\phi}_{\mathcal{T}}\right]\!\!\right] = \mathsf{1},\ [\![\phi_{\mathcal{T}}]\!] = \mathsf{0}\right]. \tag{88}$$

Taking conditional expectations with respect to the underlying data yields (86). The false negative identity is identical:

$$\mathsf{Pr}_{\mathrm{II}}(\Phi^{\mathcal{T}}) = \mathsf{Pr}\left[\left[\!\!\left[\hat{\phi}_{\mathcal{T}}\right]\!\!\right] = \mathsf{0},\ [\![\phi_{\mathcal{T}}]\!] = \mathsf{1}\right] = \mathsf{E}[(1-\mathsf{p}_{\mathcal{T}})\mathbf{1}\{[\![\phi_{\mathcal{T}}]\!] = \mathsf{1}\}]\,. \tag{89}$$

□

B.4 Proof of Corollary 1

Proof. For the premise-based contract, singleton invariance implies that shuffling never changes any atomic clause evaluation. Hence $\Phi^{\mathcal{A}}$ can differ from the target outcome only on loophole states, and only by false positives. Therefore

$$C(\Phi^{\mathcal{A}}) = c_{\mathrm{I}}\,\mathsf{Pr}(\mathsf{LH}). \tag{90}$$

From the target-based error decomposition,

$$C(\Phi^{\mathcal{T}}) = c_{\mathrm{I}}\mathsf{E}[\mathsf{p}_{\mathcal{T}}\mathbf{1}\{\mathsf{LH}\}] + c_{\mathrm{II}}\mathsf{E}[(1-\mathsf{p}_{\mathcal{T}})\mathbf{1}\{\mathsf{SR}\setminus\mathsf{LH}\}]\,. \tag{91}$$

Thus (26) implies

$$C(\Phi^{\mathcal{T}}) < C(\Phi^{\mathcal{A}}), \tag{92}$$

which is exactly (27).

If $e = 0$, then

$$\mathsf{p}_{\mathcal{T}} = [\![\phi_{\mathcal{T}}]\!] \qquad \text{a.s.} \tag{93}$$

Hence

$$C(\Phi^{\mathcal{T}}) = 0, \tag{94}$$

while

$$C(\Phi^{\mathcal{A}}) = c_{\mathrm{I}}\,\mathsf{Pr}(\mathsf{LH}), \tag{95}$$

which proves (28). □

B.5 Proof of Proposition 2

Proof. By independent evaluations across clauses,

$$\mathsf{Pr}\Big(\llbracket\Phi\rrbracket = \mathtt{1} \,\Big|\, \{\mathbf{a}_j\}_{j\in\bigcup_{m=1}^{M}\mathcal{M}_m}\Big) = \prod_{m=1}^{M} \mathsf{p}_{\mathcal{M}_m}. \tag{96}$$

If exactly r clauses should reject, then exactly r factors on the right-hand side are $O(e)$, while the remaining factors are $1 - O(e)$. Therefore

$$\mathsf{Pr}\Big(\llbracket\Phi\rrbracket = \mathtt{1} \,\Big|\, \{\mathbf{a}_j\}_{j\in\bigcup_{m=1}^{M}\mathcal{M}_m}\Big) = O(e^r), \tag{97}$$

which proves (35).

If $r = 0$, then every clause should accept, so

$$\mathsf{Pr}\Big(\llbracket\Phi\rrbracket = \mathtt{0} \,\Big|\, \{\mathbf{a}_j\}_{j\in\bigcup_{m=1}^{M}\mathcal{M}_m}\Big) = 1 - \prod_{m=1}^{M} \mathsf{p}_{\mathcal{M}_m}. \tag{98}$$

Using the union bound,

$$1 - \prod_{m=1}^{M} \mathsf{p}_{\mathcal{M}_m} \leq \sum_{m=1}^{M} \big(1 - \mathsf{p}_{\mathcal{M}_m}\big) = O(e), \tag{99}$$

which proves (36). □

B.6 Proof of Corollary 2

Proof. By the first and third assumptions, Proposition 2 implies that on every data profile with

$$\llbracket\phi_{\mathcal{T}}\rrbracket = \mathtt{0}, \tag{100}$$

the redundant contract falsely accepts with probability $O(e^2)$. Hence

$$\mathsf{Pr}_{\mathrm{I}}(\Phi^{\mathcal{R}}) = O(e^2). \tag{101}$$

By the second and third assumptions, Proposition 2 implies that on every data profile with

$$[\![\phi_{\mathcal{T}}]\!] = \mathtt{1}, \tag{102}$$

the redundant contract falsely rejects with probability $O(e)$. Hence

$$\mathsf{Pr}_{\mathrm{II}}(\Phi^{\mathcal{R}}) = O(e). \tag{103}$$

Therefore

$$C(\Phi^{\mathcal{R}}) = O(c_{\mathrm{I}}e^2) + O(c_{\mathrm{II}}e). \tag{104}$$

For the premise-based benchmark,

$$C(\Phi^{\mathcal{A}}) = c_{\mathrm{I}}\,\mathsf{Pr}(\mathsf{LH}), \tag{105}$$

so $\Phi^{\mathcal{R}} \succ \Phi^{\mathcal{A}}$ for e sufficiently small and $c_{\mathrm{I}}\,\mathsf{Pr}(\mathsf{LH})$ sufficiently large relative to c_{II}.

For the target-based contract, the fourth assumption implies that its false-positive cost on loophole states is first-order in e, while (101) shows that the corresponding false-positive cost for $\Phi^{\mathcal{R}}$ is only second-order. Since $\Phi^{\mathcal{R}}$ also has only first-order false negative by (103), it follows that $\Phi^{\mathcal{R}} \succ \Phi^{\mathcal{T}}$ for e sufficiently small and $c_{\mathrm{I}}\,\mathsf{Pr}(\mathsf{LH})$ sufficiently large relative to c_{II}. $\square$

B.7 Proof of Proposition 3

Proof. The two contracts differ only on the event

$$\left\{ [\![\Phi]\!] = \mathsf{1},\ [\![\hat{\phi}_{\mathcal{N}}]\!] = \mathsf{0} \right\}. \tag{106}$$

On that event, adding the new clause flips an acceptance into a rejection. This removes false positives exactly on the subevent where the target outcome is 0, and creates false negatives exactly on the subevent where the target outcome is 1. That yields the exact cost difference

$$C(\Phi') - C(\Phi) = -\, c_{\mathrm{I}} \, \mathsf{Pr}\left[[\![\Phi]\!] = \mathsf{1},\ [\![\phi_{\mathcal{T}}]\!] = \mathsf{0},\ [\![\hat{\phi}_{\mathcal{N}}]\!] = \mathsf{0} \right] + c_{\mathrm{II}} \, \mathsf{Pr}\left[[\![\Phi]\!] = \mathsf{1},\ [\![\phi_{\mathcal{T}}]\!] = \mathsf{1},\ [\![\hat{\phi}_{\mathcal{N}}]\!] = \mathsf{0} \right], \tag{107}$$

and the preference condition follows immediately. □

B.8 Proof of Corollary 3

Proof. By (48) and iterated expectations, conditional on the baseline outcome and the target outcome, the added atomic clause accepts with probability q.

$$\mathsf{Pr}_{\mathrm{I}}(\Phi^{(j)}) = q \ \mathsf{Pr}_{\mathrm{I}}(\Phi), \tag{108}$$

while every baseline true positive vetoed by the added clause becomes a new false negative, giving

$$\mathsf{Pr}_{\mathrm{II}}(\Phi^{(j)}) = \mathsf{Pr}_{\mathrm{II}}(\Phi) + (1 - q)\, \mathsf{Pr}[[\![\Phi]\!] = \mathsf{1},\ [\![\phi_{\mathcal{T}}]\!] = \mathsf{1}]\,. \tag{109}$$

The cost formula follows directly

$$C(\Phi^{(j)}) - C(\Phi) = (1-q)\left(c_{\mathrm{II}} \,\mathsf{Pr}[[\![\Phi]\!] = \mathsf{1},\ [\![\phi_{\mathcal{T}}]\!] = \mathsf{1}] - c_{\mathrm{I}} \,\mathsf{Pr}_{\mathrm{I}}(\Phi)\right). \tag{110}$$

□

B.9 Proof of Proposition 4

Proof. The two contracts $\Phi^{\mathcal{A}}_{\theta}$ and $\Phi^{\mathcal{A}}_{\theta+1/I}$ differ only on the event M_θ. On that event,

$$[\![\Phi^{\mathcal{A}}_{\theta}]\!] = \mathsf{1} \qquad \text{and} \qquad [\![\Phi^{\mathcal{A}}_{\theta+1/I}]\!] = \mathsf{0}. \tag{111}$$

Therefore

$$C(\Phi^{\mathcal{A}}_{\theta+1/I}) - C(\Phi^{\mathcal{A}}_{\theta}) = -\, c_{\mathrm{I}} \,\mathsf{Pr}(\mathsf{M}_\theta, [\![\phi_{\mathcal{T}}]\!] = \mathsf{0}) + c_{\mathrm{II}} \,\mathsf{Pr}(\mathsf{M}_\theta, [\![\phi_{\mathcal{T}}]\!] = \mathsf{1}). \tag{112}$$

Since $\mathsf{Pr}(\mathsf{M}_\theta) > 0$, dividing by $\mathsf{Pr}(\mathsf{M}_\theta)$ shows that

$$C(\Phi^{\mathcal{A}}_{\theta+1/I}) < C(\Phi^{\mathcal{A}}_{\theta}) \tag{113}$$

if and only if

$$c_{\mathrm{II}} \,\mathsf{Pr}([\![\phi_{\mathcal{T}}]\!] = \mathsf{1} \,|\, \mathsf{M}_\theta) < c_{\mathrm{I}} \,\mathsf{Pr}([\![\phi_{\mathcal{T}}]\!] = \mathsf{0} \,|\, \mathsf{M}_\theta). \tag{114}$$

Using

$$\mathsf{Pr}([\![\phi_{\mathcal{T}}]\!] = \mathsf{0} \,|\, \mathsf{M}_\theta) = 1 - \mathsf{Pr}([\![\phi_{\mathcal{T}}]\!] = \mathsf{1} \,|\, \mathsf{M}_\theta), \tag{115}$$

this is equivalent to (62).

It remains to prove (63). For any $\theta \leq \theta_{\mathcal{T}}$, target acceptance implies premise-level acceptance. Indeed, if $[\![\phi_{\mathcal{T}}]\!] = \mathtt{1}$, then

$$s_{\mathcal{T}} \geq \theta_{\mathcal{T}} I. \tag{116}$$

Moreover, for every $j \in \mathcal{T}$ and every state i,

$$\phi^i_{\mathcal{T}} \leq a^i_j. \tag{117}$$

Summing over i gives

$$s_{\mathcal{T}} \leq s_j \qquad \text{for every } j \in \mathcal{T}. \tag{118}$$

Hence

$$[\![\phi_{\mathcal{T}}]\!] = \mathtt{1} \quad \Longrightarrow \quad s_j \geq s_{\mathcal{T}} \geq \theta_{\mathcal{T}} I \geq \theta I \qquad \text{for every } j \in \mathcal{T}. \tag{119}$$

Therefore

$$\mathsf{Pr}_{\mathrm{II}}(\Phi^{\mathcal{A}}_{\theta}) = 0 \qquad \text{for every } \theta \leq \theta_{\mathcal{T}}. \tag{120}$$

Moreover, the acceptance events are nested:

$$\mathsf{ACC}_{\theta+1/I} \subseteq \mathsf{ACC}_{\theta}. \tag{121}$$

Therefore, for $\theta \leq \theta_{\mathcal{T}}$, raising the threshold share weakly reduces false positives and does not create false negatives, so

$$C(\Phi^{\mathcal{A}}_{\theta+1/I}) \leq C(\Phi^{\mathcal{A}}_{\theta}). \tag{122}$$

Thus no optimal symmetric threshold share can lie below $\theta_{\mathcal{T}}$ once

$$C(\Phi^{\mathcal{A}}_{\theta_{\mathcal{T}}+1/I}) < C(\Phi^{\mathcal{A}}_{\theta_{\mathcal{T}}}) \tag{123}$$

holds. This proves (63). □

References

(1935). Gregory v. helvering. 293 U.S. 465 (1935).

(1970). In re winship. 397 U.S. 358 (1970).

(1979). Addington v. texas. 441 U.S. 418 (1979).

(1982). Santosky v. kramer. 455 U.S. 745 (1982).

(2018). Eaton vance mgt. v. wilmington sav. fund socy. No. 654397/2017, NY: Supreme Court.

(2021). In re citibank august 11, 2020 wire transfers. 520 F. Supp. 3d 390 (S.D.N.Y. 2021).

(2024). In re serta simmons bedding, llc. No. 23-20181, Court of Appeals, 5th Circuit.

(2024). Ocean trails clo vii v. mln topco ltd. 2024 NY Slip Op 06660 (App. Div. 1st Dep't Dec. 31, 2024).

Adhikari, B., J. Alm, and T. F. Harris (2021). Small business tax compliance under third-party reporting. *Journal of Public Economics 203*, 104514.

Allen, F. and D. Gale (1992). Measurement distortion and missing contingencies in optimal contracts. *Economic Theory 2*(1), 1–26.

American Law Institute (1981). *Restatement (Second) of Contracts*. St. Paul, MN: American Law Institute. §224.

Ayotte, K. and A. B. Badawi (2025). Loopholes in complex contracts. *American Law and Economics Review*.

Badawi, A. B., S. D. Dyreng, E. de Fontenay, and R. W. Hills (2021). Contractual complexity in debt agreements: The case of EBITDA. Working paper.

Bajari, P., S. Houghton, and S. Tadelis (2014). Bidding for incomplete contracts: An empirical analysis of adaptation costs. *American Economic Review 104*(4), 1288–1319.

Basel Committee on Banking Supervision (2013, January). Principles for effective risk data aggregation and risk reporting. Technical report, Bank for International Settlements.

Basu, S. (1997). The conservatism principle and the asymmetric timeliness of earnings. *Journal of Accounting and Economics 24*(1), 3–37.

Battigalli, P. and G. Maggi (2002). Rigidity, discretion, and the costs of writing contracts. *American Economic Review 92*(4), 798–817.

Behn, M., R. Haselmann, and V. Vig (2022). The limits of model-based regulation. *The Journal of Finance 77*(3), 1635–1684.

Benzarti, Y. (2020). How taxing is tax filing? using revealed preferences to estimate compliance costs. *American Economic Journal: Economic Policy 12*(4), 38–57.

Beuve, J., M. W. Moszoro, and S. Saussier (2019). Political contestability and public contract rigidity: An analysis of procurement contracts. *Journal of Economics & Management Strategy 28*(2), 316–335.

Bräuning, F., V. Ivashina, and A. Ozdagli (2025). High-yield debt covenants and their real effects. *Review of Financial Studies*.

Buccola, V. S. J. and G. Nini (2024). The loan market response to dropdown and uptier transactions. *The Journal of Legal Studies 53*(2), 489–526.

Choi, S. J., M. Gulati, M. Jennejohn, and R. E. Scott (2023). Contract production in M&A markets. *University of Pennsylvania Law Review 171*, 1881–1924.

Choi, S. J., M. Gulati, and R. E. Scott (2017). The black hole problem in commercial boilerplate. *Duke Law Journal 67*(1), 1–77.

Chychyla, R., A. J. Leone, and M. Minutti-Meza (2019). Complexity of financial reporting standards and accounting expertise. *Journal of Accounting and Economics 67*(1), 226–253.

Coates, J. C. I. (2016). Why have M&A contracts grown? evidence from twenty years of deals. Discussion Paper 889, Harvard Law School, John M. Olin Center for Law, Economics, and Business. Also circulated as ECGI Law Working Paper No. 322/2016.

DeGroot, M. H., P. I. Feder, and P. K. Goel (1971). Matchmaking. *The Annals of Mathematical Statistics 42*(2), 578–593.

DeGroot, M. H. and P. K. Goel (1980). Estimation of the correlation coefficient from a broken random sample. *The Annals of Statistics 8*(2), 264–278.

Demerjian, P. R. (2011). Accounting standards and debt covenants: Has the "balance sheet approach" led to a decline in the use of balance sheet covenants? *Journal of Accounting and Economics 52*(2–3), 178–202.

Dichev, I. D. and D. J. Skinner (2002). Large-sample evidence on the debt covenant hypothesis. *Journal of Accounting Research 40*(4), 1091–1123.

Dietrich, F. (2007). A generalised model of judgment aggregation. *Social Choice and Welfare 28*(4), 529–565.

Eggleston, K., E. A. Posner, and R. Zeckhauser (2000). The design and interpretation of contracts: Why complexity matters. *Northwestern University Law Review 95*(1), 91–132.

Eisenberg, T. and G. P. Miller (2009). The flight to new york: An empirical study of choice of law and choice of forum clauses in publicly-held companies' contracts. *Cardozo Law Review 30*(4), 1475–1532.

European Central Bank (2024, May). Guide on effective risk data aggregation and risk reporting. Technical report, European Central Bank.

Ganglmair, B. and M. Wardlaw (2017). Measuring contract completeness: A text based analysis of loan agreements. Working paper.

Gelpern, A., M. Gulati, and J. Zettelmeyer (2019). If boilerplate could talk: The work of standard terms in sovereign bond contracts. *Law & Social Inquiry 44*(3), 617–646.

Gilson, R. J., C. F. Sabel, and R. E. Scott (2014). Text and context: Contract interpretation as contract design. *Cornell Law Review 100*(1), 23–97.

Grossman, S. J. and O. D. Hart (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy 94*(4), 691–719.

Hart, O. and J. Moore (1990). Property rights and the nature of the firm. *Journal of Political Economy 98*(6), 1119–1158.

Hoitash, R. and U. Hoitash (2018). Measuring accounting reporting complexity with XBRL. *The Accounting Review 93*(1), 259–287.

Holmstrom, B. and P. Milgrom (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, and Organization 7*(Special Issue), 24–52.

Hsu, D., K. Shi, and X. Sun (2017). Linear regression without correspondence. In *Advances in Neural Information Processing Systems 30*.

Ivashina, V. and B. Vallée (2025). Weak credit covenants. *Management Science 71*(12), 9998–10021.

Jakobsen, A. M. (2020). A model of complex contracts. *American Economic Review 110*(5), 1243–1273.

Judicial Council of California (2025). Judicial council of california civil jury instructions (caci) no. 322: Occurrence of agreed condition precedent. Official jury instruction, 2025 edition. Available in the Judicial Council of California Civil Jury Instructions.

Kaplow, L. (2012). Burden of proof. *The Yale Law Journal 121*(4), 738–859.

Katz, L. (2010). A theory of loopholes. *The Journal of Legal Studies 39*(1), 1–31.

Kornhauser, L. A. and L. G. Sager (1993). The one and the many: Adjudication in collegial courts. *California Law Review 81*(1), 1–59.

Li, N., F. P. Vasvari, and R. Wittenberg-Moerman (2016). Dynamic threshold values in earnings-based covenants. *Journal of Accounting and Economics 61*(2–3), 605–629.

List, C. (2012). The theory of judgment aggregation: An introductory review. *Synthese 187*(1), 179–207.

List, C. and P. Pettit (2002). Aggregating sets of judgments: An impossibility result. *Economics and Philosophy 18*(1), 89–110.

Loughran, T. and B. McDonald (2014). Measuring readability in financial disclosures. *The Journal of Finance 69*(4), 1643–1671.

Miyashita, M. (2021). Premise-based vs conclusion-based collective choice. *Social Choice and Welfare 57*(2), 361–385.

Mongin, P. (2012). The doctrinal paradox, the discursive dilemma, and logical aggregation theory. *Theory and Decision 73*(3), 315–355.

Murfin, J. (2012). The supply-side determinants of loan contract strictness. *The Journal of Finance 67*(5), 1565–1601.

Nikolaev, V. V. (2010). Debt covenants and accounting conservatism. *Journal of Accounting Research 48*(1), 137–176.

Nyarko, J. (2021). Stickiness and incomplete contracts. *The University of Chicago Law Review 88*(1), 1–80.

Office of the Comptroller of the Currency (2017). *Asset-Based Lending.* Comptroller's Handbook. Version 1.1, January 27, 2017; references to reputation risk removed as of March 20, 2025.

Pananjady, A., M. J. Wainwright, and T. A. Courtade (2017). Linear regression with shuffled data: Statistical and computational limits of permutation recovery. *IEEE Transactions on Information Theory 64*(5), 3286–3300.

Radin, M. J. (2006). Boilerplate today: The rise of modularity and the waning of consent. *Michigan Law Review 104*(5), 1223–1234.

Rakoff, T. D. (2006). The law and sociology of boilerplate. *Michigan Law Review 104*(5), 1235–1248.

Roberts, M. R. and A. Sufi (2009). Control rights and capital structure: An empirical investigation. *The Journal of Finance 64*(4), 1657–1695.

Saari, D. (2001). *Decisions and elections: explaining the unexpected*. Cambridge University Press.

Schwartz, A. and R. E. Scott (2010). Contract interpretation redux. *The Yale Law Journal 119*, 926–964.

Scott, R. E., S. J. Choi, and G. M. Gulati (2024a). Contractual landmines. *Yale Journal on Regulation 41*, 307–376.

Scott, R. E., S. J. Choi, and M. Gulati (2024b). Commercial boilerplate: A review and research agenda. *Annual Review of Law and Social Science 20*, 201–218.

Smith, H. E. (2006). Modularity in contracts: Boilerplate and information flow. *Michigan Law Review 104*(5), 1175–1222.

Taxpayer Advocate Service (2023). 2022 annual report to congress: Most serious problem 2 — complexity of the tax code: The complexity of the tax code burdens taxpayers and the IRS alike. Technical report, Taxpayer Advocate Service.

Townsend, R. M. (1979). Optimal contracts and competitive markets with costly state verification. *Journal of Economic Theory 21*(2), 265–293.

U.S. Securities and Exchange Commission (2012). Consolidated audit trail. Securities Exchange Act Release No. 34-67457; Rule 613.

Watts, R. L. (2003). Conservatism in accounting part i: Explanations and implications. *Accounting Horizons 17*(3), 207–221.

Zwick, E. (2021). The costs of corporate tax complexity. *American Economic Journal: Economic Policy 13*(2), 467–500.