Obviously Strategy Proof Liability Rules

Omer Pelled* Michael Goldboim

January 14, 2026

Abstract

Legal systems worldwide have gravitated toward comparative fault despite the theoretical equivalence of multiple liability regimes in achieving efficient care levels. This Article argues that the superiority of comparative fault lies in its game-theoretic robustness. Specifically, the comparative fault regime's ability to make compliance a dominant strategy. We analyze when liability rules can be strategy-proof (compliance is always optimal regardless of others' actions) and obviously strategy-proof (this dominance is transparent to decision-makers). We establish fundamental limits: when total liability cannot exceed actual harm, no rule can make compliance globally dominant once parties' actions interact as complements or substitutes. However, a Proportional Excess-Benefit (PEB) rule—splitting losses based on each actor's saved compliance costs—is unique in achieving strategy-proofness for independent externalities and on relevant one-sided domains where mistakes are most costly. The rule becomes obviously strategy-proof when compliant parties bear no residual liability. We demonstrate how comparative fault implements PEB principles in bilateral accidents, explaining its widespread adoption. For multiple tortfeasors, PEB offers implementable guidance superior to traditional apportionment schemes. The framework extends across tort contexts from bilateral accidents to mass torts, providing robust incentives even when parties cannot predict others' behavior.

JEL Codes: C72, D82, K13

Keywords: Comparative Negligence, Contributory Negligence, Strict Liability, Obviously Strategy-Proof Mechanism

1 Introduction

Most contemporary legal systems allocate accident losses through some form of comparative fault (Edelman 2007; Haddock and Curran 1985). The literature thoroughly examined alternative liability regimes, and found pure negligence, strict liability with a contributory negligence defense, and comparative fault as essentially interchangeable regimes with regard to care incentives: each is

*Bar-Ilan University, Faculty of Law. We thank Ariel Porat, Avraham Tabbach, and the participants of the Israeli Law and Economics Association annual conference for their valuable insights

capable of inducing first-best care levels under broad parametric conditions (R. D. Cooter and Ulen 1986; Miceli 1997, pp.17–20).

The choice of regime thus turned to secondary considerations or frictions, including incentivizing parties about unverifiable parameters that are outside the negligence inquiry (activity level), reducing error costs in the assignment of liability (Bar-Gill and Ben-Shahar 2003; R. D. Cooter and Ulen 1986), and dealing with the judgment-proof problem (Stremitzer and Tabbach 2014), to name a few.

However, the basic argument, that the regimes are interchangeable with regard to care, raises two interconnected puzzles: why have real-world legal systems gravitated toward comparative fault even when simpler fault-based regimes could theoretically suffice, and why do parties in practice often fail to select optimal care levels predicted by theoretical models?

Empirical evidence from laboratory experiments and field studies provides a valuable clue. People often choose a different strategy than the one predicted by an equilibrium analysis. This happens for two key reasons: players may lack the cognitive resources to predict others' behavior accurately, and struggle to figure out the best response to every possible action their counterparts might take, or they may assume others are similarly limited and will deviate from theoretical predictions, making it sensible to respond to these anticipated mistakes rather than textbook equilibrium play. When an equilibrium action happens to be a dominant strategy, players are much more likely to follow it, since they don't need to worry about predicting others' behavior. However, even dominant strategies can fail to guide behavior when players with limited cognitive resources can't easily recognize that the strategy is always best (Van der Linden 2018). This suggests that for a legal rule to create a broad change in behavior, what matters isn't just whether a strategy is theoretically dominant, but whether its dominance is obvious to people making decisions under uncertainty.

This insight motivates adapting two key concepts from mechanism design to tort law: strategy-proofness (SP) and obvious strategy-proofness (OSP)(Li 2017). A liability rule is strategy-proof when compliance is a dominant strategy equilibrium - i.e., complying with the legal requirement is each actor's best choice, regardless of what others do. It is obviously strategy-proof when compliance is transparently the best choice—when even the worst possible outcome from following the standard beats the best possible outcome from any violation.[1] We first establish that any rule which assigns full liability to a subset of violating actors—even when multiple actors fail to comply—cannot be strategy-proof. This rules out what we term "first fault" approaches, which encompass many common doctrines, including simple negligence, negligence with contributory negligence defenses, strict liability with contributory negligence defense, and the last clear chance doctrine. We then derive conditions under which strategy-proofness can be achieved. We show that when the externality players create are complements or substitutes, a general strategy-proof liability rule is impossible. We find, however, a class of rules that are SP compatible for independent externalities. These rules,

[1] The analysis of accidents is more convoluted, since actions affect the probability of an accident occurring, but actors often avoid accidents regardless of their actions, making the best outcome for any action one with no liability. We adopt a narrower definition of obviousness, based on the expected costs. We discuss the stochastic nature of accidents, and what that might mean to the obviousness of dominant strategy in the conclusion.

which we call Proportional Excess Benefit rules (PEB), divide harm among violators based on how much each gained by deviating from the optimal standard. When externalities are complements (substitutes), PEB rules are SP if players can act in a restricted domain, and can only comply or exceed (fall short) the efficient level. Crucially, whenever the liability rule is SP, achieving obvious strategy-proofness requires that compliant actors bear no share of residual harm—the joint externality that remains even when all players comply with their legal standards.

The analysis has important practical implications across multiple legal contexts. In bilateral accidents, the proportional-deficit rule closely corresponds to comparative fault. Current application of comparative fault always leaves residual harm with the victim or the injurer, making comparative fault systems strategy-proof but not obviously strategy-proof. OSP becomes achievable only when an outside source—such as social insurance—bears the residual harm when both parties comply. When OSP is unattainable, the proportional-deficit approach can still create an obviously dominant strategy for one party by assigning all residual harm to the other, providing a principled way to choose between negligence and strict liability versions of comparative fault. In multi-party scenarios where victims cannot affect the externality, OSP mechanisms become attainable by having victims bear the residual harm when all injurers comply with optimal standards. Finally, in sequential accidents involving multiple stages, the analysis clarifies how liability rules affect parties who may know they bear no liability or seek information about others' actions. The OSP property ensures these results hold regardless of players' information sets or beliefs about others' behavior.

2 Intuitive Explanation for SP and OSP Liability Rules

To understand the architecture of strategy-proof liability, begin with a structural weakness of traditional regimes that assign the entire loss to a single non-compliant actor. These rules—simple negligence, strict liability with contributory negligence, last clear chance, and other all-or-nothing doctrines—achieve efficiency only in a perfect equilibrium where every actor correctly anticipates that all others will comply. Outside that equilibrium, the "winner-takes-all" assignment creates a *liability immunity* for everyone except the designated wrongdoer, which invites rent-seeking: actors invest not only in precaution, but in predicting or influencing others' mistakes to avoid paying. A strategy-proof (SP) rule eliminates this informational waste by making compliance a dominant choice *regardless* of what others do.

Consider a bilateral accident: injurer and victim each face a due-care cost of 10. If both comply, harm is 0; if both deviate, harm is 60; if exactly one deviates, harm is 40. Under a rule that assigns full liability to a single deviator (the negligent party, or the "last mover," or the "most negligent"), incentives are fragile. For instance, suppose that the injurer observes that the victim already deviated. In this case, the injurer can deviate "for free:" as the victim is already fully liable due to his blatant negligence, the injurer faces no penalty for also deviating. This interdependence imposes a heavy epistemic and economic burden: to act optimally, parties must expend resources not on preventing accidents, but on monitoring and predicting other players' actions to exploit their strategic errors.

Under a rule that assigns full liability to a single deviator (e.g., the negligent injurer, the negligent victim, the "last" actor to have a chance to avoid the accident, or the "most" negligent one), incentives become fragile. Suppose the injurer observes that the victim has failed to invest in care. Under many traditional rules, the victim's blatant negligence acts as a liability shield for the injurer; because the victim is already fully liable, the injurer faces no penalty for also deviating. The injurer effectively saves 10 at no cost. This interdependence imposes a heavy epistemic and economic burden: to act optimally, parties must expend resources not on preventing accidents, but on monitoring and predicting other players' actions to exploit their strategic errors. A strategy-proof (SP) rule eliminates this informational waste by making compliance a dominant choice regardless of what others do.

Constructing such a rule requires a precise calibration of liability to the economic benefits of non-compliance. The key economic fact is simple: benefits gained by all deviating players—in the form of saved precaution costs, must be strictly lower than the resulting increase in social harm. In the running example, both parties save 10 each if they deviate (a pot of 20), but their deviation causes 60 in harm. The harm dwarfs the savings.

Incentivizing compliance requires a liability mechanism that leverages this inequality. If we split the realized harm among deviators in proportion to the private benefit each got from deviating, the "price" of deviation automatically exceeds the private saving. In the example, both saved 10, so they split 60 equally and each faces 30. For Player A, the choice stops depending on Player B: comply and pay 10 in care, or deviate and expect to pay 30—compliance wins either way. That is the basic architecture of a strategy-proof rule: tie each person's liability to their own saved costs so that the penalty always outweighs the gain.

Strictly linking the liability share to the proportion of saved costs, produces another robust implication concerning imperfect standards. This proportional architecture is also robust to judicial errors in setting the due-care benchmark. When the standard of care is set inefficiently high, meaning that saved costs from deviating outweigh the negative externality, each actor is better of deviating to the optimal care level than complying, regardless of what other do.

There is, however, a real technological limit. Consider a scenario of technological substitutes, where one person's care makes the other's redundant, or complements, where one person's negligence renders the other's care useless. If Player A's deviation causes the marginal effectiveness of Player B's care to drop to zero, it creates a scenario where Player B should ideally stop investing in care to avoid waste. A general liability rule cannot simultaneously encourage Player B to comply (to be strategy-proof) and encourage Player B to react optimally to A's deviation (to be efficient). This tension, as will we shall prove below, means no liability rule can be strategy-proof on the entire action space when actions interact. The practical fix is to target the side of mistakes that matters most: with complements, guard against overshooting (the "too high" region); with substitutes, guard against undershooting (the "too low" region). On those one-sided regions, the same proportional logic restores the dominance of compliance.

Finally, strategy-proofness and obvious strategy-proofness are not the same. An SP rule ensures that Compliance is better than Deviation, but it does not guarantee that Compliance is risk-free. Crucially, a liability rule remains strategy-proof regardless of who bears the "residual harm," that is the cost of accidents that occur even when everyone complies. An SP liability rule, we show, is characterized solely by how it divides harm between several deviating parties. Once all parties comply, the rule can assign residual harm arbitrarily. However, for a rule to be OSP, the worst possible outcome of compliance must be better than the best possible outcome of deviation. The only way to make dominance *obvious* is to strip out residual downside at full compliance for all players. The social planner can achieve any residual harm allocation while maintaining OSP liability through a form of mandatory insurance - forcing the residual bearer to pay for insurance that cover only non-negligent accidents, complying becomes an obvious dominant strategy for all players.

We formally model these intuitions in the next sections. After introducing some primitive notions, we introduce a highly minimal axiomatic framework on liability rules which places a mild requirements on how harm is divided. Building on that, we formalize SP to mean that, for each party, complying weakly dominates any deviation no matter what others do; OSP strengthens this by comparing worst-case compliance to best-case deviation. We then prove a global limit: if actions have cross-effects (complements or substitutes), no liability rule is SP on the entire action space. We show that this global failure is fixable by focusing on the *one-sided regions* where mistakes matter most: with complements, the "too-high" orthant; with substitutes, the "too-low" orthant (no restriction under independence). On each admissible region we record each party's relevant departure from the standard via a nonnegative, measurable mapping, anchored at zero when that party is at the standard. From the definitions of SP, OSP, and the requirements we derive *necessary* constraints any SP rule must satisfy on that region. We then provide an *existence* witness: the Proportional Excess-Benefit (PEB) rule, which instantiates the measurable mapping with each party's saved compliance cost. We prove that PEB is SP on the full domain under separability, and SP on the appropriate one-sided region under cross-effects. We also prove that OSP holds exactly when compliant parties bear no residual harm at the corner. Finally, we show that, on each admissible region, the requirments of liability rules together with the SP/OSP constraints *uniquely* determine the share map. Specializing this map to excess-benefit recovers PEB uniquely as a corollary.

3 The Model

We study an externality–generating activity carried out by n parties, $N = \{1, \ldots, n\}$. Each party i chooses an activity level $a_i \in \mathcal{A}_i \subseteq \mathbb{R}$, where $\mathcal{A}_i$ is a nonempty subset of $\mathbb{R}$ with its usual order and topology. We write $A = (a_1, \ldots, a_n) \in \mathcal{A} := \prod_i \mathcal{A}_i$. Activities generate private benefits and a joint external loss (social harm).

Each party's gross benefit $b_i : \mathcal{A}_i \to \mathbb{R}$ satisfies

$$b_i'(a_i) > 0, \qquad b_i''(a_i) < 0,$$

and the harm function $H : \mathcal{A} \to \mathbb{R}_+$ is twice continuously differentiable, strictly increasing in each argument, and convex. Cross–partials

$$H_{ij}(A) := \frac{\partial^2 H}{\partial a_i \, \partial a_j}(A)$$

encode the technology:

$$H_{ij} < 0 \text{ (substitutes)}, \qquad H_{ij} = 0 \text{ (independent)}, \qquad H_{ij} > 0 \text{ (complements)}.$$

We normalize $H(0) = 0$, and define the social optimum as the maximizer of aggregate welfare

$$W(A) := \sum_{i=1}^{n} b_i(a_i) - H(A),$$

which, by concavity of the b_i's and convexity of H, has a unique maximizer $A^* = (a_1^*, \ldots, a_n^*)$, and the first–order conditions are

$$b_i'(a_i^*) = H_{A^i}(A^*) \qquad (i = 1, \ldots, n), \tag{1}$$

with $H_{A^i} := \partial H / \partial a_i$. We interpret A^* as the *due care* standard. An *environment* is the pair (H, A^*), and we refer to environments by their technological relationship: independent (separable), complements, or substitutes.

3.1 Liability Rules

A *liability rule* is a mapping

$$\Lambda : \mathcal{A} \to \mathbb{R}_+^n, \qquad A \mapsto \lambda(A) = (\lambda_1(A), \ldots, \lambda_n(A)).$$

Define the *total assessment*

$$T(A) \ := \ \sum_{i=1}^{n} \lambda_i(A),$$

and, when $T(A) > 0$, the induced *shares*

$$s_i(A) \ := \ \frac{\lambda_i(A)}{T(A)} \in [0, 1], \qquad s(A) \in \Delta^{n-1}.$$

The portion of harm not assigned to players is borne by a *residual bearer* and equals

$$r(A) \ := \ H(A) - T(A).$$

The vector $\lambda(A)$ records the harm assessed to players, $(T(A))$ is their aggregate, and $s(A) = \lambda(A)/T(A)$ determines how $T(A)$ is split across the players.

The intelligibility of Λ rests on several structural constraints. We impose four structural axioms on liability rules. They constrain the *division* of any chosen assessment of harm, but do not otherwise normalize $T(A)$. Throughout, (H, A^*) is fixed (with A^* defined by (1)).

A1 (Feasibility). For every $A \in \mathcal{A}$,

$$\lambda_i(A) \ \geq \ 0 \quad \text{for all } i, \qquad \sum_{i=1}^{n} \lambda_i(A) \ \leq \ H(A).$$

A1 is a minimal constraint on the allocation of liability: no party is assigned a negative liability, and aggregate liability never exceeds realized harm. It follows that the residual harm is $r(A) = H(A) - T(A) \geq 0$, and each player's payoff is

$$u_i(A) \ = \ b_i(a_i) \ - \ \lambda_i(A). \tag{2}$$

Now, note that while the total $T(A)$ and the shares $s(A) = \lambda(A)/T(A)$ only describe what liability is allocated and how it is divided, they do not specify which features of the activity profile A the division may depend on. To make the division constraints precise, we fix a party–indexed, nonnegative measurement of deviations from the standard A^*.

Definition 1 (Measurement map). *Let $q : \mathcal{A} \to \mathbb{R}^n_+$ be a Borel–measurable mapping, written*

$$A \ \mapsto \ q(A) \ = \ \big(q_1(A), \dots, q_n(A)\big),$$

with each coordinate $q_i(A) \geq 0$.

Less formally, q is defined as follows: for each activity profile A, the vector $q(A)$ assigns to every party i a nonnegative *number* $q_i(A)$ that captures the party's measured deviation from A^* for purposes of dividing $T(A)$. Thus, q specifies which party-specific measured quantities the division of liability may depend on, whereas s is the division itself.

Now, in each state A, a feasible total assessment $T(A) \in [0, H(A)]$ is chosen (the aggregate liability to assess). Individual assignments are then $\lambda(A) = T(A) \cdot s(A)$, where the share vector $s(A) \in \Delta^{n-1}$ may depend on party–indexed measured quantities $q(A)$ (Definition 1). A core incentive requirement is this: holding fixed the total T and all other parties' measured quantities $\{q_k\}_{k \neq i}$, an increase in party i's own measured deviation q_i must not *reduce* her assignment λ_i; and when $T > 0$ it must *strictly* increase it. Otherwise, an actor could worsen her measured deviation while everything outside her control for the division (the aggregate T and others' q_k) remains unchanged, yet pay less—creating a moral-hazard pathology. Hence:

A2 (Own–measure monotonicity). There exists a measurable mapping q (defined as in Definition 1) such that for any $A, B \in \mathcal{A}$ and any party i,

$$T(B) = T(A), \quad q_k(B) = q_k(A) \ \ \forall k \neq i, \quad q_i(B) \geq q_i(A) \implies \lambda_i(B) \geq \lambda_i(A),$$

with $\lambda_i(B) > \lambda_i(A)$ whenever $q_i(B) > q_i(A)$ and $T(A) > 0$.

Informally, A2 ensures that—holding the total and others' measured quantities fixed—a higher own measured quantity cannot yield a lower assignment. The requirement is minimal: it imposes no restriction on the policy choice $T(\cdot)$, no functional form on the share map $s(\cdot)$, and no cross–party comparative statics beyond the fixed-others clause. It simply rules out perverse own-deviation incentives at the level of $(T, q) \mapsto \lambda$.

Now, the measurement map $q(\cdot)$ represents party–indexed *additive* units of the same party's deviation from A^*. Hence, This representation must be invariant under arbitrary relabeling of a single party's behavior, yet remain identifiable at the party level and compatible with feasibility. We meet this observation by imposing the following axiomatic requirement:

A3 (Clone–neutrality). Fix $A \in \mathcal{A}$ with total $T(A) \geq 0$ and let q be as in Definition 1. Replace party j by two *clones* j_a, j_b to form $\widetilde{A}$ satisfying

$$T(\widetilde{A}) = T(A), \qquad q_k(\widetilde{A}) = q_k(A) \ \ \forall k \notin \{j_a, j_b\}, \qquad q_{j_a}(\widetilde{A}), q_{j_b}(\widetilde{A}) \geq 0, \quad q_{j_a}(\widetilde{A}) + q_{j_b}(\widetilde{A}) = q_j(A).$$

When $T(A) > 0$ (equivalently $T(\widetilde{A}) > 0$), the shares $s(\cdot) = \lambda(\cdot)/T(\cdot)$ satisfy:

1. **External invariance:** For any non–clone $i \neq j$, $s_i(\widetilde{A}) = s_i(A)$.
2. **Collective preservation:** $s_{j_a}(\widetilde{A}) + s_{j_b}(\widetilde{A}) = s_j(A)$.
3. **Proportional internal split:** If $q_{j_a}(\widetilde{A}) + q_{j_b}(\widetilde{A}) > 0$, then

$$\frac{s_{j_a}(\widetilde{A})}{s_{j_b}(\widetilde{A})} = \frac{q_{j_a}(\widetilde{A})}{q_{j_b}(\widetilde{A})}.$$

   (If $q_{j_a}(\widetilde{A}) = q_{j_b}(\widetilde{A}) = 0$, any split with sum $s_j(A)$ is allowed.)

Merging j_a, j_b back into j reverses these implications.

In other words, A3 states that If one party is replaced by two labeled subparts whose measured contributions add up to the original, then: (i) others' shares do not move merely because of relabeling; (ii) the two subparts together keep exactly the original party's share; and (iii) when both subparts are positive, their internal split tracks their measured magnitudes. Thus A3 rules out *bundling premia* and *splitting discounts* for additive units of the *same* party's deviation.

Note that A3 is both necessary and minimal. It is necessary in that imposing a weaker constraint would result several key issues. If splitting j into j_a, j_b with $q_{j_a} + q_{j_b} = q_j$ could (a) change other parties' shares, or (b) alter the *combined* share of j_a, j_b, or (c) split them non-proportionally when both positive, then j could reduce liability by mere relabeling with no change in actions. This breaks the identification "measured deviation $\rightarrow$ liability" and undermines SP comparisons, because a player could improve payoff through bookkeeping rather than behavior. A3 forbids exactly these relabeling arbitrages.

A3 is minimal in that strengthening is is untenable. Suppose one demands *own-only dependence*, i.e., $s_i = f_i(q_i)$ for all i (no reference to $\{q_k\}_{k \neq i}$). Feasibility with $T > 0$ requires $\sum_i f_i(q_i) = 1$ for all profiles. On a one–positive boundary ($q_i > 0, q_{-i} = 0$) we must have $f_i(q_i) = 1$ and $f_j(0) = 0$ for $j \neq i$ (boundary behavior). With two positives ($q_i > 0, q_j > 0$), feasibility demands $f_i(q_i) + f_j(q_j) = 1$ for *all* $q_i, q_j > 0$, which cannot be reconciled with the one-positive boundary unless each f is constant—contradicting the boundary again. The only way to restore feasibility is to normalize by the *sum* (e.g., $s_i = g(q_i)/\sum_k g(q_k)$), which necessarily reintroduces cross-dependence. Hence strict own-only dependence is incompatible with feasibility and boundary behavior. Conversely, if one strengthens A3 to require invariance under *merging distinct parties* i, j into a single label m with $q_m = q_i + q_j$ while holding others' shares fixed and the combined share of i, j constant *for every split* (q_i, q_j) with the same sum, iterating merges/splits forces the division to depend only on $Q = \sum_k q_k$, not on its distribution. That collapses party-level identifiability and violates own-measure monotonicity: increasing q_i while decreasing q_j (holding Q fixed) would be forced to leave s_i unchanged, contradicting the requirement that higher own measure should not reduce one's share.

Now, when a rule uses party–indexed inputs $q(A)$ and produces a share vector $s(A) \in \Delta^{n-1}$, it thereby induces decision regions in the observable state space. We require these regions to behave regularly and impose the following axiom.

A4 (Measurability). On $\{A \in \mathcal{A} : T(A) > 0\}$, the share map

$$s(A) = \frac{\lambda(A)}{T(A)} \in \Delta^{n-1}$$

is Borel–measurable in A (with respect to the product topology on $\mathcal{A}$).

Informally: when the total is positive, the share map must be measurable relative to the party–indexed measurements $q(A)$. Concretely, for any open set $U \subseteq \Delta^{n-1}$, the preimage $\{A : s(A) \in U\}$ is a Borel set.

This regularity is minimal but essential. If some decision region $\{A : s(A) \in U\}$ were non-Borel, there would be no way to specify or verify the rule's regions using standard (countable) operations on A (or on $q(A)$), because membership in a party's liability region would not be a measurable event. In applied terms, the rule would not admit reliable tests of compliance based on observed inputs. Moreover, if some component s_i is non-measurable, then $\lambda_i = T \cdot s_i$ is non-measurable under any probability model on $\mathcal{A}$. Consequently, standard economics predictions, such as $\mathbb{E}[\lambda_i(A)]$ or $\mathbb{E}[T(A)]$, and therefore risk and expected-liability calculations collapse.

A simple pathology makes this concrete. Let $q_1 : \mathcal{A} \to [0, 1]$ be a normalized deviation, and take a non-measurable set $V \subset [0, 1]$ (e.g., a Vitali set). Define

$$s_1(A) = \mathbf{1}\{q_1(A) \in V\}, \qquad s_j(A) = 0 \ (j \neq 1), \qquad T(A) > 0.$$

Then s_1 is non-measurable, hence $\lambda_1(A) = T(A)s_1(A)$ is non-measurable and $\mathbb{E}[\lambda_1(A)]$ does not exist. The liability region $\{A : \ s_1(A) = 1\} = \{A : \ q_1(A) \in V\}$ is non-Borel and cannot be described or checked by any finite, measurable procedure relative to q_1. A4 rules out precisely these pathologies while imposing no functional-form restrictions beyond measurability.

A4 concludes the presentation of our minimal axiomatic constraints over liability rules.

3.2 Dominance, SP constraints, and the PEB Rule

3.2.1 Dominance and Strategy Proof

The previous discussion shows that because direct activity costs are embedded in the concave benefits (b_i), strategic interaction is driven entirely by how a liability rule Λ allocates realized harm $H(A)$ across parties. We turn to *dominance* properties of liability rules: Strategy-Proofness (SP), where compliance is a dominant action regardless of others, and the stronger Obvious Strategy-Proofness (OSP), where the dominance is transparent to the decision-maker.

Definition 2 (Strategy Proof (SP) Liability Rule)**.** *A liability rule* Λ *is* strategy-proof *if, for every* i *and every deviation* $\bar{a}_i \neq a_i^*$,

$$b_i(a_i^*) - \lambda_i(a_i^*, A_{-i}) \ \geq \ b_i(\bar{a}_i) - \lambda_i(\bar{a}_i, A_{-i}) \qquad \forall\, A_{-i} \in \mathcal{A}_{-i},\ \forall\, \bar{a}_i \in \mathcal{A}_i.$$

Less formally, under an SP rule, for each player, complying at a_i^* weakly dominates any deviation—no matter what others do.

Definition 3 (Obviously Strategy Proof (OSP) Liability Rule)**.** *A liability rule* Λ *is* obviously strategy-proof *if each player's* worst *payoff from complying at* a_i^* *is at least as high as her* best *payoff obtainable by deviating.*

Hence, OSP strengthens SP by comparing worst-case compliance to best-case deviation: even the least favorable compliant outcome beats the most favorable deviation.

However, as noted above, a non-SP liability rule can nevertheless induce A^* in equilibrium. Consider the following rule:

First-fault rule. Fix a public permutation $\pi = (1, 2, \ldots, N)$. Inspect actors in order π and assign the entire harm $H(A)$ to the first player k with $a_k > a_k^*$ (set $\lambda_k(A) = H(A)$ and $\lambda_j(A) = 0$ for all $j \neq k$). If all comply ($a_i = a_i^* \ \forall i$), set $\lambda_{\pi(n)}(A) = H(A)$. This family captures simple negligence (injurer first), strict liability with contributory negligence (reverse order), and "last clear chance" (order determined by who could have last avoided the accident).

Lemma 1 (Efficiency without dominance)**.** *Under the first-fault regime, the target profile A^* is the* unique Nash equilibrium. *Exactly one actor, player $\pi(1)$, has a* strictly dominant strategy *to comply; for every other actor, a_i^* is only a* best response*, not dominant. Thus, the rule is neither SP nor OSP.*

Proof.

1. If every player complies, player $\pi(n)$'s payoff is $u_n(a_n, A^*_{-n}) = b_n(a_n) - H(a_n, A^*_{-n})$. Since A^* is socially optimal, $u_n(a_n, A^*_{-n})$ has a unique maximizer at a_n^* (by $b_n'' < 0$ and convex H with (1)).

2. Player $\pi(1)$ has a strictly dominant strategy to comply: any upward deviation makes her fully liable, irrespective of others. Moreover, it is obviously dominant: her worst (indeed only) compliant payoff $b_1(a_1^*)$ exceeds any deviating payoff $b_1(a_1) - H(A)$.

3. Players $i \geq 2$ lack dominance. If any earlier player is already at fault, a later player pays zero regardless of her action and thus gains by deviating upward; hence a_i^* is not dominant for $i \geq 2$.

4. Uniqueness: At A^*, no player benefits from deviating (by steps 1–3), so A^* is a Nash equilibrium.[2]

Therefore the first-fault family is efficient, grants strict dominance only to $\pi(1)$, and is neither SP nor OSP. □

While Efficiency without dominance is possible, if players are rationally bounded, or predict that other players are relationally bounded, and might adopt an off-equilibrium strategy as a result, they may waste resources gathering information about the actions of other players or trying to convince other players to alter their actions. Thus, even when the liability rule leads to an efficient equilibrium, SP and the more restrictive OSP are superior.

3.2.2 Global impossibility of SP & one-sided domains

The discussion of the previous subsection led some commentators to the intuition that in bilateral accidents, forms of comparative fault can have a unique compliance equilibrium in dominant strategies. We show, however, that this intuition is ungrounded. The next theorem establishes that liability rules can never be SP when the actor's actions have cross-effects.

[2]Uniqueness follows by forward induction: given that $\pi(1)$ complies (dominance), $\pi(2)$'s best response is to comply, and so on. Any profile with some $a_i > a_i^*$ assigns $H(A)$ to that i, who benefits by reducing activity, hence cannot be an equilibrium.

Theorem 1 (Impossibility of global SP under cross–effects)**.** *Fix an environment* (H, A^*) *as specified above. Suppose* H *exhibits cross–effects at* A^**: there exist* $i \neq j$ *with* $H_{ij}(A^*) \neq 0$*. Then no liability rule* Λ *can be strategy–proof on the full domain* $\mathcal{A}$*.*

Proof. Write $\lambda = \Lambda$ and $T = \sum_k \lambda_k$. Assume, for contradiction, that Λ is SP on $\mathcal{A}$. By the cross–effect assumption, choose distinct $i \neq j$, a direction for j, and $\delta > 0$ such that

$$H_{A^i}\big(a_i^*,\, a_j^* \pm \delta,\, A^*_{-\{i,j\}}\big) \;\neq\; H_{A^i}(A^*).$$

Without loss of generality, take a direction that *reduces* i's marginal harm; for concreteness, let

$$B \;:=\; (a_i^*,\, a_j^* - \delta,\, A^*_{-\{i,j\}}) \quad \text{satisfy} \quad H_{A^i}(B) \;<\; H_{A^i}(A^*).$$

Consider a small upward move $\varepsilon > 0$ by party i from B to

$$B_\varepsilon \;:=\; (a_i^* + \varepsilon,\, a_j^* - \delta,\, A^*_{-\{i,j\}}).$$

Since b_i and H are continuously differentiable and A^* is interior, first–order expansions yield

$$b_i(a_i^* + \varepsilon) - b_i(a_i^*) \;=\; b_i'(a_i^*)\,\varepsilon + o(\varepsilon), \qquad H(B_\varepsilon) - H(B) \;=\; H_{A^i}(B)\,\varepsilon + o(\varepsilon) \quad \text{as } \varepsilon \downarrow 0.$$

Fix the opponents at $B_{-i} = (a_j^* - \delta, A^*_{-\{i,j\}})$. Strategy–proofness requires a_i^* to weakly dominate every deviation, hence in particular

$$b_i(a_i^*) - \lambda_i(B) \;\geq\; b_i(a_i^* + \varepsilon) - \lambda_i(B_\varepsilon),$$

which rearranges to

$$\lambda_i(B_\varepsilon) - \lambda_i(B) \;\geq\; b_i(a_i^* + \varepsilon) - b_i(a_i^*).$$

By feasibility (A1), $T(B_\varepsilon) - T(B) \leq H(B_\varepsilon) - H(B)$ and $\lambda_i \leq T$, so

$$\lambda_i(B_\varepsilon) - \lambda_i(B) \;\leq\; T(B_\varepsilon) - T(B) \;\leq\; H(B_\varepsilon) - H(B).$$

Combining the last two displays and substituting the expansions gives

$$b_i'(a_i^*)\,\varepsilon + o(\varepsilon) \;\leq\; \lambda_i(B_\varepsilon) - \lambda_i(B) \;\leq\; H_{A^i}(B)\,\varepsilon + o(\varepsilon).$$

Divide by $\varepsilon > 0$ and let $\varepsilon \downarrow 0$ to obtain

$$b_i'(a_i^*) \;\leq\; H_{A^i}(B).$$

Using the planner first–order condition (1), $b_i'(a_i^*) = H_{A^i}(A^*)$, so

$$H_{A^i}(A^*) \;\le\; H_{A^i}(B),$$

which contradicts $H_{A^i}(B) < H_{A^i}(A^*)$ by construction of B. Therefore, no strategy–proof liability rule exists on the full domain when H has cross–effects near A^*. □

3.2.3 One-sided domains & SP constraints

Theorem 1 proves impossibility when activities have a cross-effect. When activities are independent, a rule dividing liability among several actors can be SP. For such a distribution to be effective, it should be based on the benefit that each non-compliant party gains from deviating from the standard. However, errors around the standard of care need not harm welfare symmetrically. With *complements* ($H_{ij} \ge 0$), an increase by one actor raises the *marginal* harm created by others, so overshooting is especially costly; with *substitutes* ($H_{ij} \le 0$), the reverse logic makes undershooting the dominant concern. Focusing on one-sided neighborhoods of A^* therefore directs analysis—and legal guidance—to regions where mistakes are most damaging and where dominance can be credibly targeted, while acknowledging the difficulty of achieving dominance on the full action space.

Hence, we restrict our attention now to one-sided domains around A^*:

$$\mathcal{C}^{\uparrow} := \{A \in \mathcal{A} :\; a_i \ge a_i^* \;\forall i\}(\text{upward orthant}), \quad \mathcal{C}^{\downarrow} := \{A \in \mathcal{A} :\; a_i \le a_i^* \;\forall i\}(\text{downward orthant}).$$

In the independent case no one-sided restriction is required.

Now, recall that in Definition 1 we defined q as the general, party-indexed measurement. When attention is restricted to a particular admissible domain (e.g., the upward or downward orthant), we need to specialize q to record each party's domain–relevant deviation from the standard A^*. We define below this specialized measurement as a party–indexed, nonnegative domain measurement vector $z_\diamond(A) \in \mathbb{R}^n_+$ that records each party's domain–relevant deviations from the standard A^*. On this domain we may have $z_\diamond(A) = q(A)$. Within that domain, liability shares are allowed to depend on party–indexed mesures only through $z_\diamond(A)$ (together with the chosen total $T(A)$).

Definition 4 (Domain–specific measurement vectors)**.** *For each domain $\diamond \in \{\text{nc}, \uparrow, \downarrow\}$, define a Borel-measurable mapping on that domain by either setting $z_\diamond(A) = q(A)$ or, more generally, choosing a Borel, coordinatewise nondecreasing $\phi_\diamond : \mathbb{R}^n_+ \to \mathbb{R}^n_+$ and setting $z_\diamond(A) = \phi_\diamond(q(A))$. Write*

$$z_\diamond(A) = (z_{\diamond,1}(A), \ldots, z_{\diamond,n}(A)) \in \mathbb{R}^n_+, \qquad Z_\diamond(A) := \sum_{i=1}^{n} z_{\diamond,i}(A).$$

The vector $z_\diamond$ must satisfy on its domain:

- ***Anchoring at the standard.*** *$z_{\diamond,i}(A) = 0$ whenever $a_i = a_i^*$; in particular $z_\diamond(A^*) \equiv 0$.*

- ***Nonnegativity and measurability.*** *$z_\diamond(A) \in \mathbb{R}_+^n$ for all A, and each coordinate $z_{\diamond,i}(\cdot)$ is Borel measurable in A.*

On independent (separable) domain we use the convenient label $z_{\mathrm{nc}}(A) = d(A)$.

Intuitively, $z_\diamond(A)$ records each party's domain–relevant contribution *relative to* A^*. Furthermore, small own deviations that do not change the domain-relevant position of i relative to A^* leave $z_{\diamond,i}$ unchanged locally. We assume throughout that, for each i, there exists $\epsilon_i > 0$ and a family of profiles $\{A(t)\}$ within the domain such that $z_{\diamond,i}\big(A(t)\big)$ ranges over an interval $[0, \epsilon_i)$ while $z_{\diamond,k}\big(A(t)\big)$, $k \neq i$, are held fixed. This is automatic on the independent (separable) domain, and where actions are coarser or constrained, we explicitly assume the domain still permits such local variation near A^*.

The measurement $z_\diamond(A)$ alone is of course insufficient to spell-out SP and OSP liability on one-sided domains. As we shall show, they are necessarily constrained by the following results. We begin by showing if a party's domain-relevant measure is zero, SP forbids assigning them any liability. This prevents liability gaps from being closed by arbitrarily allocating them to "bystanders."

Lemma 2 (SP–Null: zero relevant measure $\Rightarrow$ zero liability)**.** *Fix a domain $\diamond \in \{\mathrm{nc}, \uparrow, \downarrow\}$ and a liability rule that is SP on that domain. If A lies in domain $\diamond$ and $z_{\diamond,i}(A) = 0$ for some i, then $\lambda_i(A) = 0$ (equivalently, if $T(A) > 0$ then $s_i(A) = 0$).*

Proof. Suppose $z_{\diamond,i}(A) = 0$ but $\lambda_i(A) > 0$. By *Anchoring*, $z_{\diamond,i}(a_i^*, A_{-i}) = 0$. Consider player i comparing a_i (as in A) to a_i^* with A_{-i} fixed. Under SP,

$$b_i(a_i^*) - \lambda_i(a_i^*, A_{-i}) \;\geq\; b_i(a_i) - \lambda_i(a_i, A_{-i}).$$

Since $b_i' > 0$, either $b_i(a_i^*) \geq b_i(a_i)$ or (if $a_i < a_i^*$ on a downward side) the inequality is strict in favor of a_i^*. Thus the SP inequality can only hold if $\lambda_i(a_i^*, A_{-i}) \leq \lambda_i(a_i, A_{-i})$. But both profiles have the same domain–relevant measure for i (zero), and the others' measures are unchanged; by A2 (own–measure monotonicity at fixed total), i cannot face strictly higher liability at a_i^* than at a_i. Hence $\lambda_i(a_i^*, A_{-i}) \leq \lambda_i(a_i, A_{-i})$. Together with $\lambda_i(A) > 0$ this contradicts dominance of a_i^*. Therefore $\lambda_i(A) = 0$.

□

The next result deals with liability shares on the one-party boundary. We show that at any profile in the domain where exactly one party i has positive domain measurement ($z_{\diamond,i}(A) > 0$ and $z_{\diamond,j}(A) = 0$ for all $j \neq i$), Lemma 3 forces a unique split: $s_i(A) = 1$ and $s_j(A) = 0$ for $j \neq i$ (equivalently, $\lambda_i(A) = T(A)$). This eliminates any freedom to assign positive liability to zero-measure parties and fixes the shares on the one-party boundary.

Lemma 3 (SP–Boundary: a single relevant deviator bears all)**.** *Fix $\diamond$ and an SP rule on that domain. If at some A in domain $\diamond$ exactly one party i has $z_{\diamond,i}(A) > 0$ while $z_{\diamond,j}(A) = 0$ for all $j \neq i$, then $\lambda_i(A) = T(A)$ and $\lambda_j(A) = 0$ for $j \neq i$ (equivalently, if $T(A) > 0$ then $s_i(A) = 1$).*

Proof. By Lemma 2, $\lambda_j(A) = 0$ for all $j \neq i$. If $\lambda_i(A) < T(A)$, some positive liability would be assigned to a party with $z_{\diamond,j}(A) = 0$, contradicting Lemma 2. Hence $\lambda_i(A) = T(A)$. □

The next result establishes a "sufficient statistic" property: within a fixed domain, if profiles have the same measurement, then they have the same liability share.

Lemma 4 (SP–Locality: dependence only through $z_\diamond$)**.** *Fix $\diamond$ and an SP rule on that domain. If A, B lie in domain $\diamond$ with $z_\diamond(A) = z_\diamond(B)$ and $T(A) = T(B) > 0$, then $s(A) = s(B)$. Equivalently, within domain $\diamond$ party–indexed assignments may depend on party–indexed data only through $z_\diamond(\cdot)$ (and the total $T(\cdot)$).*

Proof. Assume $z_\diamond(A) = z_\diamond(B)$ and $T(A) = T(B) > 0$ but $s(A) \neq s(B)$. Then for some party i we have $s_i(A) > s_i(B)$. By A3, replace party i in the environment underlying B by two clones i_a, i_b that split $z_{\diamond,i}(B)$ in the same ratio as $s_i(A)$ would suggest, and keep all other coordinates (and the total) fixed. By clone consistency, non–clone shares are unchanged and the clones' combined share equals $s_i(B)\,T(B)$. Now consider party i facing B's others: by choosing the action profile that generates the split encoded by A (which preserves $z_\diamond$ coordinate-wise), i can move from the B–share to the higher A–share while keeping the total fixed. This strictly *reduces* i's liability at the same T, contradicting the dominance of a_i^* under SP. Hence $s(A) = s(B)$. □

Note that, coupled with the measurability axiom of the share mapping s (A4), this lemma yields the following sufficient-statistic implication: by SP–Locality, if two profiles A, B are such that $z_\diamond(A) = z_\diamond(B)$ and $T(A) = T(B) > 0$, then $s(A) = s(B)$. Combined with axiom A4, this implies that on $\{A : T(A) > 0\}$ there exists a Borel mapping $\tilde{s}$ such that $s(A) = \tilde{s}(z_\diamond(A))$ whenever $T(A) > 0$, that is liability shares only depend on the measurement $z_\diamond(A)$ and not on other information concerning A (together with the choice of the total assessment of harm $T(A)$).

Now, we turn to the more stringent case of OSP. We prove that OSP liability at the zero-measure corner eliminates any residual liability under full compliance.

Lemma 5 (OSP corner: zero total measure $\Rightarrow$ zero liability)**.** *Fix a domain $\diamond \in \{\mathrm{nc}, \uparrow, \downarrow\}$ and a rule that is OSP on that domain. If at some A we have $Z_\diamond(A) := \sum_i z_{\diamond,i}(A) = 0$, then $\lambda(A) \equiv 0$ (equivalently, $T(A) = 0$ and $s(A) \equiv 0$).*

Proof sketch. When $Z_\diamond(A) = 0$, each party's domain measurement is zero; in particular, by *Anchoring* $z_{\diamond,i}(a_i^*, A_{-i}) = 0$. Suppose, towards a contradiction, that $\lambda_\ell(A) > 0$ for some party ℓ. Under OSP, the worst payoff from compliance at a_ℓ^* must be at least the *best* payoff from deviating. But there is a continuation (staying at zero measures for everyone, which is admissible on the relevant one–sided domain) in which any deviation that preserves zero measure for ℓ yields zero liability for ℓ, while compliance would assign a strictly positive liability $\lambda_\ell(A)$. Thus worst-case compliance would be strictly below a feasible best deviation payoff, and hence a contradiction. therefore $\lambda(A) \equiv 0$. □

3.2.4 Proportional Excess-Benefit

The SP/OSP constraints imply that any implementable division must track each party's gain from non-compliance. In negligence terms, liability should scale with the cost savings from deviating from the due-care standard. We instantiate the abstract domain measurements with this legally interpretable statistic and show that the resulting **Proportional Excess–Benefit (PEB)** rule satisfies A1–A4 and the SP/OSP constraints exactly on the domains where dominance is feasible.

Excess–benefit measurements For any activity profile $A = (a_1, \dots, a_n)$, define the excess private benefit of party i:

$$d_i(A) \;:=\; \big[b_i(a_i) - b_i(a_i^*)\big]_+, \qquad D(A) \;:=\; \sum_{k=1}^{n} d_k(A).$$

Thus $d_i(A) = 0$ for downward moves below a_i^* (because $b_i' > 0$). On the independent domain we set $z_{\text{nc}}(A) = d(A)$. On one-sided domains, d serves as the domain measurement with the natural directional restriction: on $\mathcal{C}^{\uparrow}$ (complements) it measures overshooting gains; on $\mathcal{C}^{\downarrow}$ (substitutes) downward moves keep $d \equiv 0$.

If heterogeneous units/exposure warrant it, fix $\beta \in \mathbb{R}^n_{++}$ and use $\tilde{d}_i := \beta_i d_i$, $\tilde{D} := \sum_k \tilde{d}_k$. Calibration cancels inside shares and does not affect dominance arguments, and we state results for d, D.

The PEB share map

Definition 5 (Proportional Excess–Benefit (PEB))**.** *Given (H, A^*) and an admissible total $T(A) \in [0, H(A)]$,*

$$\lambda_i(A) = \begin{cases} \dfrac{d_i(A)}{D(A)}\, T(A), & \text{if } D(A) > 0, \\[2ex] w_i\, T(A), & \text{if } D(A) = 0, \end{cases} \qquad i = 1, \dots, n,$$

with $w \in \Delta^{n-1}$ the compliance weights. Feasibility ($\sum_i \lambda_i = T \le H$) is immediate. On $\{T > 0\}$, shares are $s_i = \lambda_i / T$.

By construction PEB satisfies the structural axioms and the SP constraints on any admissible domain:

- **A1 Feasibility** (nonnegativity, bounded by harm) holds by definition;
- **A2 Own-monotonicity**: holding T and others' measures fixed, larger d_i weakly increases λ_i;
- **A3 Clone-neutrality**: proportionality preserves invariance and proportional split under clone split/merge of any coordinate of d;
- **A4 Measurability**: $s(\cdot)$ is a Borel function of $d(\cdot)$ and $T(\cdot)$;

- **SP–Null, Boundary, Locality**: when $D > 0$, $s_i = 0$ iff $d_i = 0$; if only one $d_i > 0$ then $s_i = 1$; and, within a domain, shares depend on party-indexed data only through d (and T).

Lemma 6 (Excess benefit $\leq$ harm under separability)**.** *If $H(A) = \sum_{k=1}^{n} h_k(a_k)$ with each h_k increasing and convex, then for all A,*

$$D(A) \;\leq\; \sum_{k=1}^{n} \big[h_k(a_k) - h_k(a_k^*)\big] \;\leq\; H(A) - H(A^*).$$

In particular, with $H(A^*)$ normalized to 0, $D(A) \leq H(A)$.

Proof. For $a_i \geq a_i^*$, concavity of b_i and (1) give $d_i \leq b_i'(a_i^*)(a_i - a_i^*) = h_i'(a_i^*)(a_i - a_i^*) \leq h_i(a_i) - h_i(a_i^*)$. Sum over i. □

SP under independence

Theorem 2 (PEB is SP in separable environments)**.** *If H is separable and $T(A) = H(A)$, then PEB is strategy-proof on $\mathcal{A}$. More generally, SP holds for any policy $T(\cdot)$ with $T(A) \geq D(A)$ for all A.*

Proof. Fix i and A_{-i}. For any upward deviation $\bar{a}_i > a_i^*$ with $D > 0$,

$$u_i(\bar{a}_i, A_{-i}) = b_i(\bar{a}_i) - \frac{d_i}{D} T(\bar{a}_i, A_{-i}) \;\leq\; b_i(\bar{a}_i) - d_i \;=\; b_i(a_i^*).$$

If $D = 0$ at (a_i^*, A_{-i}), any upward deviation makes $d_i > 0$ and the same bound applies. Downward deviations have $d_i = 0$ and strictly reduce b_i. Thus a_i^* weakly dominates. □

One-sided SP under cross-effects When H has cross-effects, global SP is impossible (Theorem 1). On the relevant one-sided orthants, PEB restores dominance.

Theorem 3 (PEB is SP on the appropriate one-sided orthant)**.** *Let $T(A) = H(A)$. (a) If $H_{ij} \geq 0$ near A^* (complements), PEB is SP on $\mathcal{C}^{\uparrow}$. (b) If $H_{ij} \leq 0$ near A^* (substitutes), PEB is SP on $\mathcal{C}^{\downarrow}$ provided $w \equiv 0$ whenever $D = 0$.*

Proof sketch. (a) On $\mathcal{C}^{\uparrow}$, downward moves have $d_i = 0$ and reduce b_i, so are strictly worse. For upward moves, convexity and complements yield $H(A) - H(A^*) \geq \nabla H(A^*) \cdot (A - A^*) = \sum_k H_{A^k}(A^*)(a_k - a_k^*) \geq \sum_k d_k = D$, so the independence bound applies. (b) On $\mathcal{C}^{\downarrow}$, any downward move has $d_i = 0$ and—with $w \equiv 0$ at $D = 0$—$\lambda_i = 0$ while b_i strictly falls; upward moves give $d_i > 0$ and the same dominance argument. □

OSP at the zero-measure corner

Theorem 4 (OSP corner characterization)**.** *In each environment where PEB is SP (separable; complements on $\mathcal{C}^{\uparrow}$; substitutes on $\mathcal{C}^{\downarrow}$), PEB is obviously strategy-proof iff $\lambda(A) \equiv 0$ whenever $D(A) = 0$ (equivalently, $w \equiv 0$ at $D = 0$).*

Proof sketch. **If:** With $w \equiv 0$ at $D = 0$, choosing a_i^* yields the constant payoff $b_i(a_i^*)$ across continuations in the relevant domain; SP then implies worst-case compliance weakly dominates best-case deviation. **Only if:** If some $w_\ell > 0$, one can construct continuations where compliance bears $w_\ell H(\cdot)$ while a tiny upward deviation makes d_ℓ/D arbitrarily small, so the best deviation payoff exceeds the worst compliance payoff. □

Hence, at the OSP corner, $(w \equiv 0)$ is required regardless of (T). Setting $w \equiv 0$ at $D = 0$ has a natural interpretation: residual harm under full compliance is financed ex ante (e.g., first-party/public insurance), making the compliance payoff $b_i(a_i^*)$ constant and hence "obvious."

3.3 Uniqueness

This section shows that, on each admissible domain, the structural axioms (A1–A4) together with the SP/OSP constraints *uniquely* pin down a share map $(s(\cdot) = \lambda(\cdot)/T(\cdot))$. We then define the calibrated proportional sharing rule on the domain measurements and show that any SP/OSP liability rule must coincide with it (up to a calibration vector and a policy choice of the total assessment $T(\cdot)$). As a corollary, we prove that when the measurements are instantiated by excess–benefit ($z_{\mathrm{nc}} = d$), the rule specializes to PEB.

Fix a domain $\diamond \in \{\mathrm{nc}, \uparrow, \downarrow\}$, and let $z_\diamond$ be the measurable mapping as defined in Definition 4. Let Λ be any liability rule satisfying A1–A4, and let $T(A)$ be its induced total assessment.

Definition 6 (Calibrated proportional sharing on domain measurements)**.** *Given a calibration vector $\beta \in \mathbb{R}^n_{++}$, the* calibrated proportional *rule on $z_\diamond$ assigns, for $Z_\diamond(A) := \sum_k z_{\diamond,k}(A)$,*

$$\lambda_i(A) \;=\; \begin{cases} \dfrac{\beta_i\, z_{\diamond,i}(A)}{\sum_k \beta_k\, z_{\diamond,k}(A)}\; T(A), & \textit{if } Z_\diamond(A) > 0, \\[2ex] w_i(A)\, T(A), & \textit{if } Z_\diamond(A) = 0, \end{cases} \qquad i = 1, \dots, n,$$

with $w(A) \in \Delta^{n-1}$ the corner weights (feasibility holds by construction).

Theorem 5 (Uniqueness of Proportional Shares (up to calibration and Total assessment))**.** *Fix an admissible domain $\diamond \in \{\mathrm{nc}, \uparrow, \downarrow\}$ and a domain–measure map $z_\diamond$. Let Λ be a liability rule that on domain $\diamond$ satisfies A1–A4 and the SP/OSP lemmas. Then there exists $\beta \in \mathbb{R}^n_{++}$ such that, for every A in the domain with $Z_\diamond(A) = \sum_i z_{\diamond,i}(A) > 0$ and $T(A) > 0$,*

$$s_i(A) \;=\; \frac{\lambda_i(A)}{T(A)} \;=\; \frac{\beta_i\, z_{\diamond,i}(A)}{\sum_{k=1}^n \beta_k\, z_{\diamond,k}(A)} \qquad (i = 1, \dots, n).$$

If $Z_\diamond(A) = 0$, then $T(A) = 0$ and $\lambda(A) \equiv 0$ (by OSP corner*). Thus the share map on the domain is uniquely pinned down up to the calibration vector β, while the total assessment $T(\cdot) \in$*

$[0, H(\cdot)]$ *remains a feasible policy choice. On the independent domain with* $z_{\rm nc} = d$ *(excess–benefit measurements), this specializes to PEB (with optional calibration).*

Proof. Fix a domain $\diamond$ and write $z(A) = (z_1(A), \ldots, z_n(A))$ with $Z(A) = \sum_i z_i(A)$. By *SP–Locality*, for each fixed total $\tau > 0$ define a measurable share map $s^\tau : \{z \in \mathbb{R}^n_+ : Z > 0\} \to \Delta^{n-1}$ by

$$s^\tau(z) \;=\; \frac{\lambda(A)}{T(A)} \quad \text{for any } A \text{ with } z_\diamond(A) = z \text{ and } T(A) = \tau,$$

which is well-defined. It inherits the boundary properties: if $z_i = 0$ then $s^\tau_i(z) = 0$ (*SP–Null*); if exactly one coordinate $z_i > 0$ then $s^\tau_i(z) = 1$ and $s^\tau_j(z) = 0$ for $j \neq i$ (*SP–Boundary*). At $Z = 0$, *OSP corner* implies $T = 0$ and $\lambda \equiv 0$; we set $s^\tau(\mathbf{0}) = \mathbf{0}$ for notational completeness.

Fix i and hold z_{-i} fixed. Clone i into two labeled subparts with measures (x, y), $x, y \geq 0$, $x + y = t$. By A3 (external invariance and collective preservation), the combined clone share equals $s^\tau_i(t; z_{-i})$ and is independent of labels and of how (x, y) splits t. Varying the split with $x + y = t$ thus forces the map $t \mapsto s^\tau_i(t; z_{-i})$ to be additive in t over arbitrary two-way partitions. With measurability (A4), additivity and SP–Null imply linearity in own measure:

$$s^\tau_i(t; z_{-i}) = K_i(z_{-i})\, t, \qquad t \geq 0,$$

for some nonnegative measurable function $K_i(\cdot)$. Summing across parties gives, for all z with $Z > 0$,

$$1 = \sum_{k=1}^{n} s^\tau_k(z) = \frac{\sum_{k=1}^n K_k(z_{-k})\, z_k}{\sum_{k=1}^n K_k(z_{-k})\, z_k},$$

so we may write

$$s^\tau_i(z) = \frac{K_i(z_{-i})\, z_i}{\sum_{k=1}^n K_k(z_{-k})\, z_k}.$$

Now restrict to a two–party slice with indices (i, j), holding $z_{-\{i,j\}}$ fixed and letting $x := z_i > 0$, $y := z_j > 0$. Then

$$s^\tau_i(x, y) = \frac{K_i(y)\, x}{K_i(y)\, x + K_j(x)\, y}, \qquad s^\tau_j(x, y) = \frac{K_j(x)\, y}{K_i(y)\, x + K_j(x)\, y}.$$

Apply A3's proportional internal split to clones of i: for any decomposition $x = x_1 + x_2$ with $x_1, x_2 > 0$, the two clone shares satisfy $s^\tau_{i_a} : s^\tau_{i_b} = x_1 : x_2$ while j's share is unchanged. If $K_i(y)$ depended on y, the ratio of clone shares (which factors through the same denominator) would pick up spurious dependence on y, violating proportionality for fixed (x_1, x_2). Hence $K_i(y) \equiv \gamma_i$ is constant in y. By symmetry, $K_j(x) \equiv \gamma_j$ is constant in x. Therefore, for all $x, y > 0$,

$$s^\tau_i(x, y) = \frac{\gamma_i\, x}{\gamma_i\, x + \gamma_j\, y}.$$

The one–party boundary case ($z_i > 0$, $z_{-i} = 0$) forces $\gamma_i > 0$ for each i. Writing $\beta_i := \gamma_i$ and reinstating all coordinates yields the calibrated proportional form

$$s_i^{\tau}(z) = \frac{\beta_i\, z_i}{\sum_{k=1}^{n} \beta_k\, z_k}, \qquad Z > 0.$$

Since $\tau > 0$ was arbitrary, the same calibrated proportional form holds for any realized total $T(A) > 0$. When $Z = 0$, *OSP corner* gives $T(A) = 0$ and $\lambda(A) \equiv 0$. Conversely, for any $\beta \in \mathbb{R}^n_{++}$, the map $s_i(z) = \beta_i z_i / \sum_k \beta_k z_k$ (with $s(\mathbf{0}) = \mathbf{0}$), together with any feasible total $T(\cdot) \in [0, H(\cdot)]$, satisfies A1–A4 and the SP/OSP lemmas on the domain. Necessity and sufficiency follow, and the share map is uniquely pinned down up to the calibration vector $\beta \in \mathbb{R}^n_{++}$; the policy choice of $T(\cdot)$ remains free within feasibility. □

Remark (Policy and measures). The calibration β means we can multiply each party's measure by a positive constant to fix units/exposure; because shares are ratios, this rescaling does not change the shares. The total $T(\cdot)$ is free up to feasibility, and thus for SP on the separable domain it suffices that $T(A) \geq D(A)$ (as in the PEB analysis), and on one–sided domains we take $T = H$ as in the SP results above. OSP at the corner requires $w \equiv 0$ independently of T.

Corollary 1 (Specialization to PEB)**.** *On the independent (separable) domain, take $z_{\text{nc}} = d$ (the excess–benefit measurements) and write $D(A) = \sum_i d_i(A)$. Then Theorem 5 implies that for all A with $D(A) > 0$ and $T(A) > 0$, there exists $\beta \in \mathbb{R}^n_{++}$ such that*

$$s_i(A) \;=\; \frac{\lambda_i(A)}{T(A)} \;=\; \frac{\beta_i\, d_i(A)}{\sum_k \beta_k\, d_k(A)} \qquad (i = 1, \ldots, n).$$

If $D(A) = 0$, then by the OSP corner result $T(A) = 0$ and $\lambda(A) \equiv 0$. Choosing the normalization $\beta_i \equiv 1$ (or, equivalently, absorbing β into a rescaling of the units of d) yields the canonical PEB shares

$$s_i(A) = \frac{d_i(A)}{D(A)} \quad \textit{for } D(A) > 0,$$

and with any feasible total $T(\cdot)$ satisfying $T(A) \geq D(A)$ (e.g., $T = H$) we recover exactly the PEB assignments $\lambda_i(A) = \frac{d_i(A)}{D(A)}\, T(A)$.

Proof. On the separable domain set $z_{\text{nc}} = d$. By Theorem 5, for every A with $D(A) > 0$ and $T(A) > 0$ there exists $\beta \in \mathbb{R}^n_{++}$ with

$$s_i(A) = \frac{\beta_i\, z_{\text{nc},i}(A)}{\sum_k \beta_k\, z_{\text{nc},k}(A)} = \frac{\beta_i\, d_i(A)}{\sum_k \beta_k\, d_k(A)}.$$

If $D(A) = 0$, the OSP corner condition implies $T(A) = 0$ and thus $\lambda(A) \equiv 0$. To obtain the canonical PEB shares, normalize $\beta_i \equiv 1$ (or absorb β into the measurement units), giving $s_i(A) = d_i(A)/D(A)$ whenever $D(A) > 0$. Combining these shares with any feasible total $T(\cdot)$ such that $T(A) \geq D(A)$ (in particular $T = H$) yields $\lambda_i(A) = \frac{d_i(A)}{D(A)}\, T(A)$, which is precisely the PEB rule analyzed above. □

Theorem 5 and Corollary 1 show that, once admissible domain measurements $z_\diamond$ are fixed, A1–A4 hold, and the SP/OSP constraints are respected, *there is no remaining freedom* in how harm is divided among parties, apart from the harmless calibration β and the policy choice of $T(\cdot)$ (subject to the domain–specific SP conditions). In particular, on each admissible domain, the only SP–compatible division is calibrated proportional sharing on the domain measurements; and OSP additionally forces zero residual liability at the corner ($Z_\diamond = 0$). When measures are chosen as excess–benefit, the framework specializes to PEB.

4 Bilateral Accidents

In accident settings where both injurer and victim can influence the risk of harm, Brown (1973) has long shown that multiple liability regimes can achieve first-best incentives to invest in care. In particular, the classic results demonstrate that under ideal conditions (fully rational actors and correctly set legal standards of care), negligence (no liability if due care is met), strict liability with a contributory-negligence defense (injurer liable unless victim was negligent), and comparative fault (liability split by relative fault) all induce socially optimal care in equilibrium (Shavell 1980; Karapanou 2020; Landes and Posner 1987). In a pure strategy Nash equilibrium, each party takes due care and no one is found liable – an outcome that minimizes total accident costs. In essence, when one party's liability is dependent on fault, then that party, be that the injurer or the victim, will meet the due standard to avoid liability. A rational residual barer of liability will also take due care, knowing she better minimize the overall costs of accidents. As Parisi and Singh (2010) showed, a liability regime that divides the residual cost of accidents between compliant parties leads to the same equilibrium. Yet, despite this theoretical interchangeability, real-world legal systems have exhibited a strong preference for comparative fault.

Beginning in the late 20th century, contributory negligence – the old all-or-nothing rule barring recovery by a negligent victim – was steadily replaced by comparative negligence across most jurisdictions. In the United States, for example, most states follow some form of comparative fault, with only a handful of states retaining the pure contributory rule (Edelman 2007). Similar shifts occurred in other countries. Today, comparative apportionment of liability is the norm in the UK, Canada, and most of Europe (Karapanou 2020). This widespread adoption was driven in part by normative appeals to fairness – the intuition that a slightly negligent victim should not be denied all compensation – but proponents also argued that comparative fault would enhance economic efficiency in more realistic settings.

The literature has explored several explanations, each relaxing the ideal assumptions, to explain the revealed preference for a rule dividing liability between negligent parties. One argument is based on error-resilience. R. D. Cooter and Ulen (1986) argues that comparative fault better withstands court errors in determining negligence. Under contributory negligence, a small error in judging due care can radically alter liability (shifting the entire loss to one party or the other). By contrast, comparative fault smooths out these errors by apportioning damages incrementally. For example, if a court underestimates the injurer's due care level, under contributory negligence, the injurer might wrongly bear 100% of the loss. In contrast, under comparative fault, the injurer might only bear, say, 60%, reflecting a less draconian consequence of the error. This intuitive advantage has been formalized in models of evidentiary uncertainty, where random errors in measuring care levels distort incentives less under comparative sharing. However, as Bar-Gill and Ben-Shahar (2003) has shown, comparative fault does not always lead to better incentives when courts might err. Through simulations, they show that other rules can sometimes outperform comparative negligence in the presence of errors, undermining the notion that comparative fault universally minimizes error-induced inefficiency. A second line of reasoning highlights activity-level effects. A negligence rule (with or without contributory negligence) assigns residual harm to the victim, and misallocates incentives for the injurer's activity level since an injurer who meets the due care standard faces no liability for any accidents caused. Strict liability, by contrast, forces injurers to internalize accident losses (curbing their activity level) but fails to motivate victims to take care or moderate their exposure. However, a comparative fault regime is only relevant in apportioning harm between two negligent parties. Since in equilibrium both parties comply, the parties' activity level depends only on how they share the residual harm.

A third explanation involves judgment-proof injurers. If injurers have limited assets or insurance, a regime that pins 100% of damages on the injurer (as strict liability or negligence without contributory fault can in some instances) may fail to provide compensation or deterrence beyond the injurer's ability to pay. Liability rules that gradually increase the share to which insolvent injurers are liable create better incentives than all-or-nothing rules Stremitzer and Tabbach (2014). Comparative fault can provide better incentives than an all-or-nothing rule by ensuring that the victim shares in the loss when they are partly at fault. However, this rationale is somewhat limited. The positive effects of a gradual liability regime depend on the extent to which the other party is at fault. If the insolvent injurer anticipates that the victim will not be negligent, then comparative fault is essentially the same as a pure negligence regime.

While these explanations tend to provide limited support for comparative fault, the liability regime has several proponents. For instance, White's 1989 study suggests that a comparative fault regime induces less care in the real world. Low and Smith 1995 suggested that comparative negligence increases the incentive to sue (since even a partly at-fault plaintiff can recover something) and leads to higher legal costs due to arguments over apportionment.

We propose a fresh perspective: the superiority of comparative fault may lie in its game-theoretic robustness – specifically, its ability to induce optimal behavior as a dominant strategy for both parties, not merely as a Nash equilibrium. Comparative liability effectively implements a "Proportional Excess Benefit" rule: each party's share of accident losses is proportional to the benefit they obtained by deviating from the due-care standard. In practical terms, when both parties are negligent, the court compares how far each fell short of due care (e.g., the difference in investments between the optimal precaution and the actual precaution) and allocates damages in proportion to those shortfalls (Edelman 2007).In effect, comparative fault turns the bilateral care game into one where "don't be negligent" is a dominant strategy for both sides – a much stronger condition than mere Nash equilibrium. Consider a single accident with a loss of 100. Meeting the legal standard costs the injurer 18 (for example, braking or speed reduction) and the victim 12 (for example, using a crosswalk and heightened attention). When both meet the standard, accidents are rare; when one or both fall short, accidents are more frequent. The table below lists care costs and, separately, how accident risk translates into expected harm.

Accident risk and expected harm (loss = 100)

Who takes due care?	Accident probability	Expected harm
Both	5%	5
Exactly one	25%	25
Neither	50%	50

Under comparative fault, if both parties are negligent, the court splits the loss by relative shortfalls from due care. Here, that means 60 to the injurer and 40 to the victim (reflecting the 18 versus 12 shortfalls). If only one party is negligent, that party bears the full loss. If no one is negligent, there is no liability.

The injurer's choice does not depend on the victim. If the victim meets the standard, the injurer can either pay 18 to comply and owe nothing, or shirk and face an expected outlay of 25 as the sole negligent party. Paying 18 is cheaper than risking 25. If the victim falls short, the injurer can again pay 18 and owe nothing (because the victim is negligent), or shirk and expect to pay 30 (which is 60% of the expected harm 50). Paying 18 is cheaper than 30. Thus, no matter what the victim does, the injurer is better off meeting the standard.

The victim faces a similar incentive structure. If the injurer meets the standard, the victim can also meet the standard, pay her 12 in care, and face the small residual expected harm of 5 that no one is liable for, for a total of 17. If she shirks when the injurer complies, she is the sole negligent party and faces an expected outlay of 25. If the injurer falls short, the victim can either pay 12 to comply and be fully compensated (so no expected harm out of pocket), or shirk and expect to pay 20 (which is 40% of the expected harm 50). Twelve is cheaper than twenty. Thus, no matter what the injurer does, the victim is better off meeting the standard. In this simplified example, the liability rule is also OSP - the worst outcome for both parties when complying (18 for the injurer and 17 for the victim) is better than the best outcome when shirking (25 for the injurer and 20 for the victim). This is not always the case. Consider the following variation to the example: **Accident risk and expected harm (loss = 100)**

Who takes due care?	Accident probability	Expected harm
Both	55%	55
Exactly one	75%	75
Neither	100%	100

In this case, the injurer still has an obvious dominant strategy to invest in care. By meeting the standard, the Injurer pays 18, which is lower than shirking and paying 75, if the victim invested in care, or 60 if she did not. The victim has a dominant strategy to take due care, but it is no longer obvious. When meeting the standard, the worst outcome for the victim is paying 67 in care and residual liability. This is more than the best outcome when the victim is negligent - if both parties are negligent, the victim incurs only 40 in accident costs.

This mechanism-design insight helps explain the appeal of comparative fault. The rule essentially implements a strategy-proof allocation of accident costs in the bilateral case: no party can gain by "gaming" the system or unilaterally adjusting their care level away from the social optimum. In reality, of course, humans may not perceive the incentive structure with perfect clarity, but the comparative fault rule makes the choice to take care *simple and credible.* Neither party has to second-guess whether the other will be careful – each knows that, regardless of the other's choice, they are better off being careful themselves.

The analysis can also help us determine who should be the residual cost-bearer. Suppose a social planner is concerned that injurers and victims might err and not play a dominant strategy. In that case, comparative fault can be refined to be not only SP but OSP by eliminating residual liability for both parties. One practical way to do this is to ensure that when both parties exercise due care, any remaining accident losses are fully covered by insurance or a compensation fund. In other words, if both comply with the legal standard, neither will have to bear the loss from a bad-luck accident personally. Such a system might involve, say, each party carrying first-party insurance that pays out if an accident occurs despite both parties' due care (akin to a no-fault coverage triggered only in the no-negligence scenario). Thus, from each individual's perspective, taking due care strictly dominates in an obvious way: it removes not only liability risk but even the risk of uncompensated loss due to the other's fault or no one's fault. If fault-contingent insurance is impractical, residual liability should be assigned to the more rational party, which is more likely to adopt a non-obvious dominant strategy.

5 Multiple Tortfeasors

Many of the most difficult tort problems involve several actors, each imposing risk on outsiders: factories discharging into a watershed, drivers who collectively create danger by staging an illegal street race, or manufacturers that produce an identical hazardous product. The doctrinal menu for such cases—joint and several liability, equal-split or market-share apportionment, comparative negligence across multiple defendants, and causation-based schemes—aims to spread losses fairly and to deter harm. But when liability schemes ignore the extent of each party's deviation from the standard, an injurer's best response may depend on what others do. In consequence, care can be inefficiently low because each injurer internalizes only a fraction of the harm (Shavell 1987; Dillbary 2013; Dillbary 2016), or inefficiently high as defensive care when precautions act as (strategic) complements or when the Hand test is applied with limited information about others' behavior (Guttel, Procaccia, and Winter 2021). Both effects are well-documented in the literature on multi-party accidents and group causation doctrines (e.g., concerted action, concurrent causes, or alternative liability). For multiple risk creators, apportionment rules frequently fail to make "take care" an obviously best choice for every actor, regardless of others' behavior. Cooter and Porat 2007 suggested a simple fix—make each tortfeasor pay the *entire* excessive harm—but the aggregate payout then exceeds realized harm, violating the budget constraint of liability rules.

The earlier analysis implies a practical OSP liability rule for settings in which victims cannot influence the externality (pollution, drag racing harm to bystanders): (i) when at least one injurer is non-compliant, divide the realized harm among all non-compliant injurers in proportion to each one's excess benefit from deviating (the costs saved or gains realized by not meeting the legal standard); (ii) when all injurers comply, assign any residual harm to the victim(s). This Proportional Excess-Benefit (PEB) sharing respects the budget constraint—total liability never exceeds harm and is never negative—and, crucially, makes compliance an *obviously* dominant strategy for every injurer: the worst-case outlay from complying (the compliance cost itself) is strictly less than the best-case outlay from any deviation (a positive share of harm), regardless of others' behavior. The metric is also implementable: it relies on shortfalls from the prescribed standard (abatement expenditures not made; safety steps skipped), not opaque percentage fault scores (Edelman 2007). The following example illustrates how the PEB liability rule may solve the incentive dilution problem for multiple tortfeasors.

Imagine three drivers compete in an illegal drag race in the streets of the city, endangering pedestrians. Since the drivers participate in an unlawful activity, they are all jointly and severally liable for accidental harm, under the concerted-action doctrine, to ensure that no victim remains uncompensated when the direct injurer is judgment-proof or unidentified. There remains the question of how to divide the harm between participants in the race, whether they are sued together or if some pay and then sue the others in a claim for contribution. To examine the incentives, assume that each driver can reduce accidental harm by 30, but the cost of care for each is different - 5 for the first driver, 15 for the second and 20 for the third. The default, "equal-split" rule dilutes incentives. By investing in care, each driver reduces the overall risk by 30, and because they split the cost of accidents, each internalizes only 10 in saved accident costs. Thus, the first driver will pay 5 to reduce liability by 10, but the other two drivers will not invest in care.

Sharing the costs according to PEB solves the incentive problem. When all drivers are negligent, they create an expected cost of 90. The third driver, having underinvested 20 out of the 40 they underinvested in care in total, is held liable for half the harm - 45, making investment in care preferable for being negligent. Similarly, the first driver also prefers to invest 5 in care to avoid liability of 11.25 ($\frac{90}{8}$). In fact, no matter what anyone else is doing, each driver's best outcome is to invest in care and not face liability. The PEB framework extends naturally to other multi-party tort contexts where traditional apportionment schemes fail to create proper incentives. In product liability cases involving multiple manufacturers of identical harmful products, courts could allocate damages based on each manufacturer's saved safety investments rather than market share. In mass tort scenarios with multiple corporate defendants, liability could be proportioned by the compliance costs each defendant avoided rather than their relative "causal contribution"—a notoriously difficult concept to operationalize. The key practical advantage of PEB is its reliance on observable, quantifiable measures: the actual costs of safety equipment not purchased, training programs not implemented, or quality controls not adopted. This contrasts favorably with traditional comparative fault assessments that require courts to make subjective judgments about relative blameworthiness. While implementation would require careful documentation of industry safety standards and their costs, this information is often already available through regulatory compliance data, industry best practices, and expert testimony about reasonable precautionary measures.

6 Conclusion

This Article has explored when liability rules can make compliance with legal standards a dominant strategy, and when that dominance is obvious to decision-makers. We showed that while traditional fault-based regimes can achieve efficient equilibria, they fail to provide the robust incentive structure that dominance offers. The Proportional Excess-Benefit rule emerges as a practical solution that makes compliance dominant in independent externality settings and on relevant one-sided domains when activities interact.

However, our analysis confronts a fundamental challenge inherent to accident law: accidents occur only with some probability. This stochastic nature means that no liability rule can be truly obviously strategy-proof in the strictest sense. When actors choose their care levels, they face uncertainty about whether an accident will occur at all. A negligent actor who escapes liability simply because no accident materialized experiences the best possible outcome—full private benefits with zero liability. This reality undermines the transparent dominance that defines OSP, since the worst outcome from compliance (paying care costs plus expected residual liability) may exceed the best outcome from deviation (full private benefits when no accident occurs).

This limitation suggests that truly obvious mechanisms in accident settings may require actors to internalize the marginal risk they create, not just the conditional liability given an accident. Insurance schemes that make actors pay premiums proportional to the risk they generate, such as those analyzed by Ben-Shahar (2023), could restore the obvious dominance property by ensuring that risk-creators face costs regardless of whether accidents materialize. Under such schemes, the worst outcome from compliance becomes more favorable relative to the best outcome from deviation, potentially achieving genuine OSP. This insight reveals an important trade-off in liability design: more restrictive equilibrium concepts require more expensive implementation. The historical evolution of tort law illustrates this pattern. The nineteenth-century shift from strict liability to negligence reflected courts' willingness to gather information about defendants' behavior to achieve more nuanced incentives. The twentieth-century move from all-or-nothing rules (negligence with contributory negligence defenses) to comparative fault regimes demanded even more information—courts now had to measure and compare the degree of fault across multiple parties rather than simply determining whether a threshold was met.

As data collection becomes cheaper and more sophisticated, we should expect liability systems to adopt increasingly granulated compensation mechanisms. These systems will require more information to implement, but will create more obviously optimal incentives for care among multiple parties. The framework developed here guides this evolution, showing how proportional sharing based on measurable deviations from standards can achieve robust incentive properties that simpler rules cannot match.

The analysis also clarifies why comparative fault has become the dominant approach in modern tort systems. Beyond its intuitive fairness appeal, comparative fault implements a form of proportional excess-benefit sharing that makes compliance a dominant strategy for both injurers and victims. This game-theoretic robustness may explain its widespread adoption better than traditional efficiency arguments that assume perfect rationality and information. Looking forward, the principles identified here extend beyond traditional accident settings to any context where multiple actors create joint externalities—from environmental regulation to financial systemic risk. As legal systems grapple with increasingly complex multi-party harms, the tools developed in this Article offer a principled approach to designing liability rules that guide behavior even when actors cannot perfectly predict others' choices or when courts cannot perfectly observe all relevant actions.

References

Bar-Gill, Oren and Omri Ben-Shahar (2003). "The uneasy case for comparative negligence". In: *American Law and Economics Review* 5.2, pp. 433–469.

Ben-Shahar, Omri (2023). "Privacy Protection, At What Cost? Exploring the Regulatory Resistance to Data Technology in Auto Insurance". In: *Journal of Legal Analysis* 15.1, pp. 129–157.

Brown, John Prather (1973). "Toward an economic theory of liability". In: *The Journal of Legal Studies* 2.2, pp. 323–349.

Cooter, Robert and Ariel Porat (2007). "Total liability for excessive harm". In: *The Journal of Legal Studies* 36.1, pp. 63–80.

Cooter, Robert D. and Thomas S. Ulen (1986). "An economic case for comparative negligence". In: *NYUL rev.* 61, p. 1067.

Dillbary, J Shahar (2013). “Tortfest”. In: *U. Chi. L. Rev.* 80, p. 953.

— (2016). “Causation actually”. In: *Ga. L. Rev.* 51, p. 1.

Edelman, Paul H (2007). “What Are We Comparing in Comparative Negligence”. In: *Wash. UL Rev.* 85, p. 73.

Guttel, Ehud, Yuval Procaccia, and Eyal Winter (2021). “Shared liability and excessive care”. In: *The Journal of Law, Economics, and Organization* 37.2, pp. 358–391.

Haddock, David and Christopher Curran (1985). “An economic theory of comparative negligence”. In: *The Journal of Legal Studies* 14.1, pp. 49–72.

Karapanou, Vaia (2020). “Strict Liability Versus Negligence”. In: *Encyclopedia of Law and Economics.* Ed. by Alain Marciano and Giovanni Battista Ramello. New York, NY: Springer New York, pp. 1–8.

Landes, William M and Richard A Posner (1987). *The economic structure of tort law.* Harvard University Press.

Li, Shengwu (2017). “Obviously Strategy-Proof Mechanisms”. In: *American Economic Review* 107.11, pp. 3259–3287.

Low, Stuart and Janet Kiholm Smith (1995). “Decisions to retain attorneys and file lawsuits: An examination of the comparative negligence rule in accident law”. In: *The Journal of Legal Studies* 24.2, pp. 535–557.

Miceli, Thomas J. (1997). *Economics of the Law: Torts, Contracts, Property, Litigation.* None. OUP Catalogue. Oxford University Press.

Parisi, Francesco and Ram Singh (2010). “The efficiency of comparative causation”. In: *Review of Law & Economics* 6.2, pp. 219–245.

Shavell, Steven (1980). “Strict Liability versus Negligence”. In: *The Journal of Legal Studies* 9.1, pp. 1–25. (Visited on 08/17/2025).

— (1987). *Economic analysis of accident law.* Harvard University Press.

Stremitzer, Alexander and Avraham D Tabbach (2014). “The robustness case for proportional liability”. In: *The BE Journal of Theoretical Economics* 14.1, pp. 371–395.

Van der Linden, Martin (2018). “Bounded rationality and the choice of jury selection procedures”. In: *The Journal of Law and Eoternomics* 61.4, pp. 711–738.

White, Michelle J (1989). “An empirical test of the comparative and contributory negligence rules in accident law”. In: *The RAND Journal of Economics*, pp. 308–330.