Divide-and-Conquer against Coordinated Plaintiffs

Ehud Guttel, Marina Halac, Ilan Kremer,
Yuval Procaccia, and Eyal Winter*

May 19, 2026

Abstract

Divide-and-conquer is often viewed as ineffective against coordinated opponents. We show that fee-shifting rules—which can shift litigation costs to plaintiffs based on rejected offers—overturn this intuition in settlement bargaining. A defendant faces multiple plaintiffs whose claims are worth more when more plaintiffs sue, and plaintiffs can coordinate through side contracts. Even so, and even when fee shifting would have no effect against a single plaintiff, it gives the defendant a divide-and-conquer advantage, lowering settlement costs. We characterize optimal mechanisms and show that, when litigation costs are high, the defendant does as well as if side contracts were unavailable.

*Guttel: The Hebrew University (email: ehud.guttel@mail.huji.ac.il). Halac: Yale University and CEPR (email: marina.halac@yale.edu). Kremer: University of Warwick and The Hebrew University (email: ikremer@mail.huji.ac.il). Procaccia: Bar-Ilan University (email: yuval.procaccia@biu.ac.il). Winter: Lancaster University and The Hebrew University (email: eyal.winter@mail.huji.ac.il). Acknowledgments to be added. Hongcheng Li provided excellent research assistance.

1 Introduction

Divide-and-conquer is a long-standing strategic principle. It applies when a principal confronts multiple agents who benefit from acting together. By contracting dynamically or discriminating across agents, the principal can prevent them from fully capturing those gains (e.g., Segal, 2003). Yet divide-and-conquer strategies are often viewed as ineffective when used against agents who can coordinate with one another. Following Coasean logic (Coase, 1960), coordination should allow them to internalize how each agent's decision affects the others' payoffs and thus bargain as a single party.

In this paper, we show that this intuition breaks down in common legal settings. In our environment, the principal is a defendant and the agents are plaintiffs. The plaintiffs' claims exhibit complementarities: the value of each claim increases with the number of plaintiffs who sue, for example because related claims are assessed together and damages depend on the scale of proven harm. Before deciding whether to litigate, plaintiffs receive settlement offers from the defendant. Plaintiffs can coordinate through binding side contracts, so their litigation decisions maximize their joint payoff. Even so, we show that divide-and-conquer remains effective. In fact, under certain conditions, the defendant's optimal settlement offers extract as much surplus as when side contracts are not available to the plaintiffs.

Consistent with the Coasean benchmark, plaintiffs in our setting face no direct contracting costs. Instead, the key institutional feature is a *fee-shifting rule*. Legal systems often use such rules to encourage settlement and reduce litigation costs and docket congestion. While terminology and details vary across jurisdictions—as we describe in Section 2—the basic structure is similar: a litigant who rejects a settlement offer and then obtains a judgment no more favorable than the offer must reimburse the opposing party's litigation costs. Similar principles also apply in international arbitration, where rejecting an offer may trigger an adverse costs order if the final award is not more favorable.

A simple example illustrates the main idea. Consider a defendant and two plaintiffs, A and B. Each plaintiff recovers 150 at trial if both sue and 100 if only one sues. Suppose first that no fee-shifting rule applies. If plaintiffs cannot use side contracts, the defendant can use a divide-and-conquer strategy to settle for a total cost strictly below 300. Indeed, she can use a dynamic mechanism that offers 100 to A, followed by an offer of 100 to B if A accepts and 150 if A rejects. By backward induction, both plaintiffs are willing to settle for 100 each, so the defendant extracts the gains from coordination.[1] However, side contracts remove the advantage of these strategies: A would compensate B to reject after his own rejection whenever the defendant's total offers sum to less than 300. Thus, when plaintiffs

[1]The defendant can guarantee that this is the unique induced outcome by slightly increasing all offers.

can use side contracts, the defendant must pay their full joint trial award of 300.

Now suppose a fee-shifting rule applies: a plaintiff who rejects a settlement offer at least as high as his trial award must pay the defendant's litigation costs. Suppose these costs are 80 per litigant, and consider the following variant of the mechanism above. The defendant first offers 100 to each plaintiff and, if both reject, returns sequentially with offers of 150, first to A and then to B only if A rejects. If both plaintiffs accept the initial offers, they receive 200 in total. If both reject all offers, the court awards 150 to each, but each plaintiff must pay the defendant 80 in litigation costs, so the plaintiffs' joint payoff is $300-160<200$. If both reject the initial offers of 100 but one later accepts 150, or if exactly one rejects 100, then the remaining litigant is awarded 100 and must pay the defendant 80 in litigation costs, so the plaintiffs are again left with less than 200 in total. Therefore, under this mechanism, accepting both initial offers maximizes the plaintiffs' joint payoff. Despite their ability to use side contracts, the defendant extracts the gains from coordination and induces them to settle for 100 each. As we explain below, the defendant achieves the same outcome even when litigation costs are lower, as low as 50 per litigant in this example.

We emphasize that if the plaintiffs were a single legal entity, the fee-shifting rule would have no effect in the setting just described. The same would be true in our main model, which purposely assumes away any uncertainty. The fee-shifting rule matters in this environment precisely because plaintiffs receive separate offers and must rely on side contracts to coordinate. Intuitively, it makes rejection histories matter plaintiff by plaintiff; side contracts can redistribute payoffs across plaintiffs, but they cannot eliminate the cost consequences of separate rejected offers. Thus, even in the absence of direct transaction costs, fee shifting introduces a friction that undermines the Coasean logic.

Fee-shifting rules have recently drawn public attention following Hugh Grant's 2023 lawsuit against News Group Newspapers (NGN), publisher of *The Sun*. Grant's case was one of several related actions alleging privacy violations that were heard together before the same judge. He ultimately settled with NGN, citing concerns about fee shifting: "[T]he rules around civil litigation mean that if I proceed to trial and the court awards me damages that are even a penny less than the settlement offer, I would have to pay the legal costs of both sides. [...] Rupert Murdoch's lawyers are very expensive. So even if every allegation is proven in court, I would still be liable for something approaching £10m in costs."[2]

[2]See "Hugh Grant settles court claim against Sun publisher for 'enormous sum'," *The Guardian*, April 17, 2024. The fee-shifting rule would have applied to Grant's case even if the damages were equal to NGN's settlement offer. At the time of Grant's settlement, Prince Harry and 41 other claimants remained scheduled to go to trial, and one of their lawyers argued that settlement offers made proceeding impossible because of the risk of adverse costs; see "Hugh Grant says he took 'enormous sum' to settle suit alleging illegal snooping by The Sun tabloid," *Associated Press*, April 17, 2024.

We study fee-shifting rules when a defendant faces multiple plaintiffs who can collaborate. In Grant's case, plaintiffs shared substantial legal work and certain litigation expenses across their claims. Similar situations arise when polluters are sued by affected property owners, employers are sued by workers for employment-related harms, or large service providers (e.g., hospitals or investment firms) are sued by consumers.[3] Although plaintiffs sue individually, they often coordinate by sharing information, pooling resources, and aligning legal strategies. We analyze the defendant's optimal settlement mechanism in such environments and show that fee shifting can give defendants an important advantage.

Our model consists of a liable defendant and a set of plaintiffs. Plaintiffs litigate individually but may contract with each other. The value of each claim increases with the number of plaintiffs who sue, so plaintiffs' litigation decisions are strategic complements. The defendant commits to a dynamic mechanism of settlement offers with a finite number of rounds. In each round, she may make offers to the unsettled plaintiffs, who then decide whether to accept or reject. Plaintiffs who reject all offers proceed to trial, together with any others who do the same, and the defendant incurs a litigation cost for each plaintiff who litigates. Under the fee-shifting rule, this cost is shifted to a litigating plaintiff if and only if he has previously rejected an offer at least as high as his trial award.

Because there is no private information and litigation is socially costly, the defendant can be taken, without loss, to use a mechanism that induces full settlement. Her problem is then to minimize total settlement payments subject to the constraint that no coordinated profile of acceptances and rejections yields the plaintiffs a higher joint payoff.

Section 4 studies two benchmarks. Without fee shifting, side contracts are fully effective: the plaintiffs can act as a single party, reject any profile of offers that gives them less than the full joint trial award, and therefore extract that payoff in settlement. Without side contracts, by contrast, the defendant can use a sequential divide-and-conquer mechanism to induce each plaintiff to settle for the claim value he would obtain by litigating alone. These benchmarks isolate the separate roles of fee shifting and plaintiff coordination.

Section 5 contains our main results. We begin by describing a class of coordinated litigation threats available to plaintiffs. For any given, possibly empty, subset of plaintiffs, consider the coordinated strategy in which those plaintiffs commit to litigate, while each other plaintiff agrees to settle only if offered at least the trial award he would obtain by litigating alongside that subset. This guarantees the plaintiffs a minimum joint payoff regardless of the defendant's mechanism. The strongest threat in this class therefore yields a lower bound on the defendant's optimal settlement cost. In the two-plaintiff example above, a strongest litigation threat involves no plaintiff committing to litigation when the per-plaintiff litigation

[3] A closely related application is eminent domain; see the concluding remarks.

cost is at least 50, and both plaintiffs committing when the cost is below 50.

We then show that this lower bound on the defendant's optimal cost is tight, and characterize optimal phased mechanisms that achieve it. In phase 1, the defendant makes a common offer to all plaintiffs, equal to the per-plaintiff amount implied by their strongest litigation threat. If not all plaintiffs accept, the mechanism then proceeds through later phases, which are divide-and-conquer phases: in each phase, the defendant approaches plaintiffs who have not yet accepted sequentially, stopping once enough plaintiffs have settled and otherwise moving to the next phase if all remaining plaintiffs reject. Across these later phases, offers increase with the relevant trial values, ensuring that any coordinated deviation triggers fee shifting for the litigating plaintiffs.

The two-plaintiff example above again provides an illustration. When the per-plaintiff litigation cost is at least 50, the defendant can settle for 200 in total by first offering 100 to each plaintiff and, if both plaintiffs reject, returning sequentially with offers of 150. This ensures that on any path ending in trial, every litigating plaintiff has previously rejected an offer at least as high as his trial award and therefore bears fee-shifting costs. When the per-plaintiff litigation cost is below 50, an optimal phased mechanism still improves on the no-fee-shifting benchmark, allowing the defendant to settle for 300 less the litigation costs she would incur if both plaintiffs litigated.

This characterization yields two key implications. First, fee shifting makes divide-and-conquer effective: relative to the benchmark with side contracting and no fee shifting, it strictly lowers the defendant's optimal settlement cost, no matter how low litigation costs are. Second, when litigation costs are sufficiently high, the plaintiffs' strongest threat involves no litigation at all, and consequently fee shifting eliminates the plaintiffs' gains from coordination: the defendant attains the same settlement cost as in the benchmark without side contracts.[4]

Section 6 studies two alternative bargaining protocols. First, suppose the defendant is restricted to bilateral offers, so she must make a single round of simultaneous offers. This removes her ability to condition later offers on earlier rejections and therefore raises her optimal cost relative to the baseline. Even so, fee shifting continues to make divide-and-conquer effective: the defendant's optimal cost remains strictly below the no-fee-shifting benchmark, and under sufficiently high litigation costs she attains the same cost as in the corresponding bilateral-offers benchmark without side contracts.

Second, we consider mutual offers, allowing plaintiffs to make counteroffers during settle-

[4]We derive a sufficient litigation-cost threshold for this full-surplus-extraction result. If the claim-value function is weakly convex, it is enough that the defendant's per-plaintiff litigation cost be at least the difference between the per-plaintiff award when all plaintiffs sue and when only one sues.

ment bargaining. In this environment, the fee-shifting rule depends on the last offers made by both sides before litigation. Counteroffers partly restore plaintiffs' bargaining power because they allow them to place offers on the record that protect them against adverse fee shifting at trial. This raises the defendant's optimal cost relative to the baseline; in fact, it restores the same lower bound as under bilateral offers. Yet fee shifting still makes divide-and-conquer effective, and under sufficiently high litigation costs the defendant attains that lower bound.[5]

Section 7 considers three variants of our model. In the first, plaintiffs may be heterogeneous in their claim values and in the defendant's litigation costs attributable to each plaintiff. In the second, plaintiffs bear their own litigation costs and may shift these costs to the defendant. In the third, the per-plaintiff trial award is fixed, but the defendant's liability is uncertain, and the probability that litigating plaintiffs prevail in court depends on the number of plaintiffs who sue. We show that the divide-and-conquer logic extends to these environments, while also clarifying where our baseline results must be qualified.

Our results identify a novel divide-and-conquer effect of fee shifting. By exposing plaintiffs to adverse cost consequences when they reject separate offers, fee-shifting rules weaken their ability to coordinate against the defendant. These rules may have benefits that we do not study here, such as discouraging socially costly litigation in environments with private information or other frictions that may lead cases to trial.[6] But whatever the tradeoffs are in single-plaintiff settings, our point is that they change when a defendant faces multiple plaintiffs: fee shifting then affects not only whether cases settle but also how much of the gains from coordination the defendant can capture.

As noted, fee-shifting principles are common across legal systems as well as in international arbitration. More broadly, the mechanism we identify is not unique to legal settings. Similar effects may arise whenever agents who reject a principal's offer face additional costs if the eventual outcome is no more favorable than the rejected terms. These costs need not take the form of formal fee shifting or even entail a transfer of costs from the principal to the agents. They may instead arise from reputational, political, or other consequences of declining an offer and later doing no better. They may also be psychological, such as regret after rejecting an offer and then failing to improve on it. Whenever such costs are present, they can give rise to similar divide-and-conquer effects, weakening agents' ability to coordinate against the principal even when they can side-contract without direct costs.

[5] We derive sufficient litigation-cost thresholds for the defendant to achieve the lower-bound cost under bilateral and mutual offers. The threshold is lower in the latter case.

[6] Prior work examines how fee-shifting rules affect settlement rates in single-plaintiff settings. We discuss this literature below.

Literature. Our paper relates to three strands of literature. The first is the literature on contracting with externalities and divide-and-conquer, pioneered by Segal (1999, 2003). A main strand of this literature studies how a principal can ensure a desired outcome in multi-agent settings with strategic complementarities; see, among others, Winter (2004), Bernstein and Winter (2012), Halac, Kremer and Winter (2020), and Halac, Lipnowski and Rappoport (2026).[7] Che and Spier (2008) is especially close to our setting: it studies settlement bargaining with multiple plaintiffs and shows how a defendant can exploit plaintiffs through divide-and-conquer strategies.[8] These papers all study agents who act separately, so the principal's problem is shaped by the strategic uncertainty this creates. We instead consider plaintiffs who coordinate through binding side contracts, and show that fee shifting can nevertheless restore a divide-and-conquer advantage for the defendant.

Second, our paper relates to the literature on fee shifting in settlement bargaining. Starting from Spier (1994), this literature studies how fee-shifting and other offer-of-settlement rules affect bargaining positions, settlement terms, and the incidence of trial; see also Miller (1986), Chung (1996), Bebchuk and Chang (1999), Farmer and Pecorino (2000), Chen and Wang (2007), and, more broadly, Klement and Neeman (2005).[9] This work studies single-plaintiff environments, where fee shifting matters by changing an individual plaintiff's incentive to accept or reject a settlement offer. Our contribution is to show that when a defendant faces multiple coordinated plaintiffs, fee shifting has an additional effect: by exposing plaintiffs to separate adverse cost consequences, it weakens their ability to coordinate in bargaining.

Third, our paper relates to the literature on plaintiff coordination and collective negotiations. Che (2002) studies collective negotiation in multi-plaintiff litigation. Spier (2002, 2003), Hua and Spier (2005), Daughety and Reinganum (2011), and Choi and Spier (2022) study how related claims shape bargaining, aggregation, commitment, and litigation incentives across plaintiffs. Our main departure is to study how coordinated plaintiffs interact with fee-shifting rules and dynamic settlement offers. In our setting, plaintiffs can use side contracts, but their coordination remains imperfect because fee shifting operates separately across claims.

[7] For a broader discussion of recent work on contracting for coordination, see Halac (2025).

[8] Related political-economy applications include Dal Bó (2007) and Iaryczower and Oliveros (2017), which study how actors can influence collective outcomes by exploiting externalities among individual decisions.

[9] For surveys of litigation and settlement, see Hay and Spier (1998) and Spier (2007).

2 Fee shifting in practice

Conditions and variations. Offer-based fee-shifting rules (known by names such as the "offer-of-judgment" rule, "Calderbank offers," and "offers without prejudice save as to costs") apply across many jurisdictions. In the United States, the federal example is Rule 68 of the Federal Rules of Civil Procedure, and many states have enacted analogous rules that apply, and often expand, Rule 68's mechanism. In the United Kingdom, the analogue is Part 36 of the Civil Procedure Rules. Similar rules exist in Australia, Canada, Hong Kong, Singapore, and civil-law jurisdictions such as Italy and Spain. Related principles also arise in international arbitration. We provide supporting references in Appendix A.

The formal conditions are straightforward and broadly similar across jurisdictions. Under Rule 68, for example, a defendant must serve the plaintiff with a settlement offer at least 14 days before trial. In multi-plaintiff lawsuits, the defendant may extend separate offers to different plaintiffs and may make one plaintiff's offer conditional on the other plaintiffs' decisions to accept or reject. If a plaintiff rejects and later obtains a judgment less than or equal to what was offered, the plaintiff loses the right to recover post-offer costs and must pay the defendant's post-offer costs.[10]

The main variations are along three dimensions. First, some jurisdictions limit the rule to offers made by defendants, as under Rule 68, whereas others apply the mechanism symmetrically to both plaintiffs and defendants, as under Part 36. Second, in some jurisdictions the rule is triggered only if the rejected offer exceeds the ultimate judgment by a specified threshold. Third, jurisdictions differ on what counts as recoverable costs, including whether attorneys' fees, enhanced interest, or double costs are available.

Real cases underscore the significance of these rules. In recent years, U.S. courts applying Rule 68 or state-law analogs have awarded millions of dollars in costs in ordinary civil actions.[11] In some cases, these awards approach, or even exceed, the amounts claimed in the suit.[12] Courts and arbitral tribunals in other jurisdictions have reached comparable results.[13]

Concerns over strategic use. Courts have long recognized that fee-shifting rules exert pressure on the litigant receiving a settlement offer. But that pressure has generally been understood as part of the rule's design. As the U.S. Supreme Court observed, Rule 68

[10] Information concerning the rejected settlement offer is disclosed to the judge only after judgment on the merits has been rendered and is considered solely for the purpose of determining costs.

[11] For example, trial courts in Georgia and Nevada entered cost-shifting awards of about $14 million and $11 million, respectively. See Atlanta Women's Specialists, LLC v. Trabue, 377 Ga. App. 755 (2025); Ana v. Republic Silver State Disposal, Inc., 2022 Nev. Dist. LEXIS 555.

[12] See, e.g., Bosnali v. Michaud, 2020 ONCA 7.

[13] See, e.g., OMV Petrom SA v. Glencore International AG [2017] EWCA Civ 195; Pye v. Di Trapani, 2025 ONCA 355.

requires plaintiffs to "think very hard" about whether continued litigation is worthwhile.[14]

In multiple-plaintiff settings, however, the strategic use of the rule has generated concern. Specifically, courts have considered whether defendants may use Rule 68 offers to "pick off" named plaintiffs in class actions by offering complete relief and thereby mooting the action before certification. The issue reached the Supreme Court twice. In the first case, the Court appeared to endorse the strategy; in the second, it limited defendants' ability to use Rule 68 offers to deprive the entire class of its claims.[15]

This line of cases suggests that defendants have tried to use the rule not only to facilitate settlement but also to strategically reduce their overall exposure. It also highlights that the rule becomes especially consequential when multiple plaintiffs are involved. This broader strategic issue is the focus of our analysis.

3 Model

Setup. Consider a set $N = \{1, \dots, n\}$ of $n \geq 2$ plaintiffs who may take a defendant to trial. The value of each plaintiff's claim depends on how many plaintiffs sue. If $s \in \{1, \dots, n\}$ plaintiffs sue, the defendant must compensate each of them for damages of $x(s)$, where the function $x : \{1, \dots, n\} \to \mathbb{R}_+$ is strictly increasing, and we let $x(0) := 0$. The plaintiffs' litigation decisions are thus strategic complements.

Before plaintiffs decide whether to sue, the defendant can make settlement offers. We consider a dynamic setting with a finite but arbitrarily large number R of rounds. The defendant commits to a public mechanism $T = (t(r))_{r=1}^{R}$ that specifies offers $t(r) = (t_1(r), \dots, t_n(r)) \in \mathbb{R}_+^n$ in each round $r = 1, \dots, R$ as a function of the history of past offers and acceptance decisions up to that round, where not making an offer to plaintiff i in round r is equivalent to setting $t_i(r) = 0$. Denote by $\mathcal{T}$ the class of such mechanisms. In each round r, each plaintiff who receives a strictly positive offer chooses whether to accept or reject it, with rejected offers expiring at the end of the round.[16] Without loss, we set $t_i(r') = 0$ for all $r' = r + 1, \dots, R$ if plaintiff i accepts in round r, and we also set $t(r') = (0, \dots, 0)$ for all $r' = r + 1, \dots, R$ if $t(r) = (0, \dots, 0)$.

The plaintiffs can coordinate their settlement decisions through binding side contracts. They remain separate legal claimants and cannot commit ex ante, before the defendant chooses the mechanism, to a binding acceptance rule. Once the defendant has proposed the mechanism, however, the plaintiffs can redistribute payoffs among themselves, for example

[14] Marek v. Chesny, 473 U.S. 1, 11 (1985).

[15] Genesis HealthCare Corp. v. Symczyk, 569 U.S. 66 (2013); Campbell-Ewald Co. v. Gomez, 577 U.S. 153, 165–66 (2016).

[16] If $t_i(r) = 0$, plaintiff i is understood to receive no live offer in that round.

by compensating one another for rejecting offers or by agreeing how to allocate settlement proceeds. Thus, given a mechanism T, plaintiffs choose an acceptance profile to maximize the sum of their payoffs. Let $\alpha \in \{1, \ldots, R, R+1\}^n =: \mathcal{A}$ record the round at which each plaintiff accepts an offer. For example, if all plaintiffs accept in the first round, then $\alpha = (1, \ldots, 1)$. The decision $\alpha_i = R+1$ means plaintiff i rejects all offers; in this case, the plaintiff may either abandon the claim for zero or proceed to trial.

In the case of a lawsuit, the defendant faces litigation costs of $k > 0$ per litigating plaintiff.[17] As discussed in the Introduction, the allocation of these costs is determined by a fee-shifting rule based on settlement offers. Specifically, suppose that given the defendant's mechanism T, a set $S \subseteq N$ of plaintiffs sue, with $s := |S|$ and $i \in S$. Then the defendant pays the cost of k corresponding to i unless she has made a settlement offer to i at least as high as the awarded damages of $x(s)$, in which case i must bear this cost. To simplify the exposition, we assume throughout that $k < x(1)$. This means that any plaintiff who rejects all offers and remains unsettled (i.e., any $i \in N$ with $\alpha_i = R+1$) strictly prefers continuing to trial rather than abandoning the claim.[18]

Given $T \in \mathcal{T}$ and $\alpha \in \mathcal{A}$, denote the set of litigating plaintiffs by $S(\alpha) := \{i \in N : \alpha_i = R+1\}$, with $s(\alpha) := |S(\alpha)|$, and let $(t^{T,\alpha}(r))_{r=1}^{R}$ denote the sequence of offer vectors realized under T along the history induced by α. Writing $t_i^{T,\alpha}(r)$ for the ith component of $t^{T,\alpha}(r)$, define

$$z(T,\alpha) := \left|\left\{i \in S(\alpha) : \max\{t_i^{T,\alpha}(1), \ldots, t_i^{T,\alpha}(R)\} \geq x(s(\alpha))\right\}\right|$$

to be the number of litigating plaintiffs responsible for their litigation costs. The sum of payoffs to the $s(\alpha)$ litigating plaintiffs is thus $s(\alpha)x(s(\alpha)) - z(T,\alpha)k$.[19] The defendant's total cost is

$$\sum_{i \in N \setminus S(\alpha)} t_i^{T,\alpha}(\alpha_i) + s(\alpha)x(s(\alpha)) + (s(\alpha) - z(T,\alpha))k.$$

In particular, the defendant's total cost exceeds the plaintiffs' joint payoff by exactly $s(\alpha)k$, i.e. by the total litigation costs incurred given α.

[17] In the model, k should be interpreted as the recoverable portion of the defendant's litigation costs attributable to a litigating plaintiff.

[18] All statements below about litigation costs are made under the maintained assumption $k < x(1)$. Since the litigation-cost thresholds we derive in our results depend on differences in claim values rather than their absolute levels, one can always add a common constant to $x(1), \ldots, x(n)$ so that this inequality holds. Thus this assumption rules out abandonment without narrowing our substantive conclusions. Allowing $k \geq x(1)$ can only weakly lower the defendant's optimal settlement cost in our model.

[19] Because the fee-shifting rule is applied only after the court determines the damages award, $x(\cdot)$ is independent of whether plaintiffs bear the litigation cost k.

Defendant's problem. The defendant chooses a mechanism of settlement offers in order to minimize her total cost. As just described, this cost consists of any payments made to the plaintiffs—both during settlement and in court—plus any litigation costs the defendant bears. Given the mechanism proposed by the defendant, the plaintiffs choose an acceptance profile that maximizes their joint payoff. We adopt the tie-breaking assumption that plaintiffs accept when indifferent, and among payoff-maximizing response profiles, they select one with the lowest number of rejections of strictly positive offers.

Since going to court is socially costly, efficiency requires that no plaintiff sue. Moreover, since there is no private information in our model, the defendant minimizes her cost by proposing a mechanism that induces this efficient outcome. In fact, we can show that it is without loss of optimality to focus on mechanisms that induce each plaintiff to accept his first strictly positive offer.

Lemma 1. *Without loss of optimality, the defendant proposes a mechanism* $T = (t(r))_{r=1}^{R}$ *that induces each plaintiff* $i \in N$ *to settle by accepting his first strictly positive offer.*

Denote the class of such mechanisms by $\hat{\mathcal{T}} \subseteq \mathcal{T}$. Fix a mechanism $T \in \hat{\mathcal{T}}$, and let $\hat{\alpha} \in \mathcal{A}$ denote the acceptance profile in which each plaintiff accepts in the first round in which he receives a strictly positive offer.[20] Along the history induced by $\hat{\alpha}$, both the defendant's cost and the plaintiffs' joint payoff are equal to $\sum_{i \in N} t_i^{T,\hat{\alpha}}(\hat{\alpha}_i)$, and the following incentive compatibility constraint must hold (recall $x(0) := 0$):

$$\sum_{i \in N} t_i^{T,\hat{\alpha}}(\hat{\alpha}_i) \geq \sum_{i \in N \setminus S(\alpha)} t_i^{T,\alpha}(\alpha_i) + s(\alpha)x(s(\alpha)) - z(T,\alpha)k \quad \forall \alpha \in \mathcal{A}. \tag{IC}$$

That is, the sum of settlement payments when each plaintiff accepts his first strictly positive offer is no smaller than the plaintiffs' joint payoff from choosing any other acceptance profile.

Under our tie-breaking assumption, the defendant's problem can thus be written as follows:

$$\min_{T \in \hat{\mathcal{T}}} \; W(T) = \sum_{i \in N} t_i^{T,\hat{\alpha}}(\hat{\alpha}_i)$$
$$\text{subject to (IC)}.$$

We denote an optimal mechanism by T^* and the defendant's optimal cost by $W^* := W(T^*)$. Observe that the defendant can guarantee a cost no larger than $nx(n)$ by offering the highest claim value of $x(n)$ to each of the n plaintiffs. At the same time, the plaintiffs

[20] Since $T \in \hat{\mathcal{T}}$, each plaintiff settles, so $\hat{\alpha}_i \in \{1, \ldots, R\}$ for all $i \in N$.

can guarantee themselves a joint payoff no smaller than $nx(1)$ by having each plaintiff reject any offer strictly below $x(1)$ and accept otherwise. The following lemma records these observations. Let $\overline{W} := nx(n)$ and $\underline{W} := nx(1)$.

Lemma 2. *The defendant's optimal cost satisfies* $\underline{W} \leq W^* \leq \overline{W}$.

The gap $\overline{W} - \underline{W} = n(x(n) - x(1))$ is the coordination surplus in our model. This gap arises because we consider multiple plaintiffs whose claims exhibit complementarities. If the defendant faced a single plaintiff—or, more generally, if claim values were independent of the number of litigating plaintiffs—then $\overline{W} = \underline{W}$. In that case, the fee-shifting rule would have no effect on the defendant's optimal cost.[21]

Corollary 1. *With a single plaintiff, the defendant's optimal cost would be independent of fee shifting.*

The rest of the paper asks how much of the coordination surplus $\overline{W} - \underline{W}$ the defendant can capture in our setting with multiple plaintiffs, and how this depends on the fee-shifting rule. We say that divide-and-conquer is effective if the defendant can extract some of this surplus, i.e. if $W^* < \overline{W}$.

4 Benchmarks

There are two central ingredients in our model. The first is the defendant's ability to shift litigation costs to the plaintiffs as a function of her settlement offers. The second is the plaintiffs' ability to coordinate their settlement decisions using side contracts. To clarify the role of each ingredient, we next examine what would occur in the absence of either.

4.1 No fee-shifting rule

Suppose there is no fee-shifting rule, so the defendant is responsible for the per-plaintiff litigation cost of k regardless of the offers she makes during settlement negotiations.

As a warm-up, consider a single-plaintiff setting as discussed in the previous section, and take the plaintiff's claim value to be $nx(n)$. It is straightforward that the plaintiff would reject any offer strictly below $nx(n)$, and as a result, the defendant would optimally settle for that amount.

The same outcome obtains with n plaintiffs whose total award from litigating together is $nx(n)$ and who can coordinate via side contracts. Without fee shifting, side contracts

[21] This conclusion would not apply in a setting where the level of damages is uncertain; see Spier (1994) and Section 7.3.

are sufficient for the plaintiffs to act as a single plaintiff: they can agree not to accept any settlement path that gives them strictly less than $nx(n)$ in total, and thus extract $nx(n)$ in settlement.[22]

Proposition 1. *Suppose there is no fee-shifting rule. Then the defendant's optimal cost is* $W^* = \overline{W}$.

Thus, absent fee shifting, divide-and-conquer is not effective when plaintiffs can contract with one another: the defendant cannot extract any of the coordination surplus $\overline{W} - \underline{W}$.

4.2 No side contracts

Suppose that the plaintiffs cannot sign side contracts. In this case, any mechanism T proposed by the defendant induces a dynamic game between the plaintiffs: in each round $r = 1, \ldots, R$, they simultaneously make accept/reject decisions to maximize their own individual payoffs, after observing the history of settlement offers and all plaintiffs' past decisions. We study the subgame-perfect Nash equilibria of the induced game.

In principle, the induced game may admit multiple equilibria. We show that, under our tie-breaking assumption that a plaintiff accepts when indifferent, the defendant can achieve the lowest possible cost regardless of equilibrium selection.

Proposition 2. *Suppose that the plaintiffs cannot sign side contracts. Then the defendant's optimal cost is* $W^* = \underline{W}$.

The idea is simple. Suppose the defendant approaches the plaintiffs one by one according to some order. She offers $x(1)$ to the first plaintiff, and continues offering $x(1)$ to the next plaintiff as long as all previous offers have been accepted. If some plaintiff rejects, the defendant switches to offering $x(n)$ to each remaining plaintiff. By backward induction (and our tie-breaking assumption), this mechanism induces all plaintiffs to accept.[23] Introducing fee shifting does not affect this logic.

Remark 1. *The result in Proposition 2 holds regardless of whether a fee-shifting rule is in place.*

[22] Observe that the plaintiffs' agreement can be reached after the defendant proposes the mechanism; no ex ante commitment to an acceptance rule is needed. At the same time, the result does rely on costless coordination: if coordination were costly, the defendant could reduce settlement payments accordingly.

[23] This result relates to the literature on contracting with externalities (Segal, 2003; Che and Spier, 2008). Unlike that literature, Proposition 2 allows the defendant to use a dynamic mechanism and focuses on the subgame-perfect Nash equilibria of the induced game. Things are different when the defendant is constrained to using static mechanisms. See Section 6.1.

Thus, when side contracts are unavailable, the defendant can use divide-and-conquer to attain $\underline{W}$ independently of fee shifting, thereby extracting the entire coordination surplus $\overline{W} - \underline{W}$.

5 Main results

Our model combines a fee-shifting rule with binding side contracts among the plaintiffs. Relative to the benchmark cases in Section 4, we show that fee shifting makes divide-and-conquer effective even against coordinated plaintiffs. For sufficiently high litigation costs, it even allows the defendant to capture the entire coordination surplus $\overline{W} - \underline{W}$.

We begin by deriving a lower bound on the defendant's optimal cost, tighter than that in Lemma 2. We exploit the fact that, via side contracting, the plaintiffs can use coordinated *litigation threats* to secure a minimum joint payoff regardless of the defendant's mechanism. By the incentive constraint (IC), any mechanism that induces full settlement must therefore deliver the plaintiffs a joint settlement payoff at least as large as this guaranteed value.

For any set $S \subseteq N$ with $|S| = s \in \{0, \dots, n-1\}$, consider the following coordinated strategy by the plaintiffs. Each plaintiff in S rejects all offers and proceeds to trial. Each plaintiff in $N \setminus S$ accepts the first offer at least $x(s+1)$ (if such an offer is ever made) and otherwise rejects all offers. Note that depending on the defendant's offers, this coordinated litigation threat can lead to more than s plaintiffs ultimately litigating.

If exactly the s plaintiffs in S litigate, then the $n-s$ non-litigants each settle for at least $x(s+1)$. The s litigants each receive $x(s)$ at trial, and, in the worst case, each of them bears litigation cost k under fee shifting. Thus, in this case the plaintiffs' joint payoff is at least

$$V_s := (n-s)x(s+1) + s(x(s) - k).$$

If instead more than s plaintiffs litigate, then each additional litigant $i \notin S$ must have rejected only offers strictly below $x(s+1)$. Since the trial award is at least $x(s+1)$ in this case, such an additional litigant cannot be made responsible for litigation costs under fee shifting. Moreover, each plaintiff in $N \setminus S$ obtains at least $x(s+1)$ whether he settles (by construction) or litigates (since the trial award is at least $x(s+1)$). Therefore, along any history induced by this coordinated litigation threat, at most the s plaintiffs in S can bear litigation costs, and the plaintiffs can guarantee themselves a joint payoff of at least V_s, regardless of the defendant's mechanism.

A similar logic applies when the set S of plaintiffs who reject all offers and proceed to trial is the whole set N. In this case, the plaintiffs can guarantee themselves a joint payoff

of at least

$$V_n := n(x(n) - k).$$

The plaintiffs can coordinate on the strongest litigation threat that guarantees them the highest payoff regardless of the defendant's mechanism. This yields a lower bound on the defendant's optimal cost. Let $\overline{V} := \max_{s\in\{0,\dots,n\}} V_s$.

Lemma 3. *The defendant's optimal cost satisfies* $W^* \geq \overline{V}$.

Note that $V_0 = \underline{W}$, which is the lower bound in Lemma 2.

To illustrate, consider an example similar to that described in the Introduction but with $n = 3$ plaintiffs and claim values $(x(1), x(2), x(3)) = (100, 125, 150)$. If all three plaintiffs reject all offers and sue, then, in the worst case, each bears cost k, so they can secure $V_3 = 3(150 - k)$. If instead two plaintiffs commit to litigate while the remaining plaintiff insists on settling for at least 150, then again in the worst case the two litigants bear cost k, so the plaintiffs can secure $V_2 = 150+2(125-k)$. If one plaintiff commits to litigate while the other two insist on settling for at least 125, then they can secure $V_1 = 2(125)+100-k$. Finally, by insisting on receiving at least 100 each, they can secure $V_0 = 3(100)$. In this example, a strongest litigation threat yields $\overline{V} = V_0 = 300$ when $k \geq 50$ and $\overline{V} = V_3 = 450 - 3k$ when $k \leq 50$.[24]

The next result shows that the lower bound in Lemma 3 is tight: the defendant can guarantee a cost as low as $\overline{V}$.

For any $\nu \geq 0$, let $\mathcal{T}(\nu) \subseteq \mathcal{T}$ denote the class of phased mechanisms defined as follows. In phase 1, the defendant makes simultaneous offers of ν to all n plaintiffs. If all accept, the mechanism ends. Otherwise, for each phase $q \in \{2, \dots, n\}$, the defendant approaches unsettled plaintiffs in some fixed order and makes each of them an offer of $x(q)$, continuing in this way only while more than $q-1$ plaintiffs remain unsettled. If at the start of phase q only $q-1$ plaintiffs remain unsettled, or as soon as this number falls to $q-1$, the mechanism ends and these plaintiffs proceed to trial. If instead all remaining plaintiffs reject $x(q)$, the mechanism moves to phase $q+1$ (and if $q = n$, the remaining plaintiffs proceed to trial).

Let $\nu^* := \overline{V}/n$.

Proposition 3. *The defendant's optimal cost is* $W^* = \overline{V}$. *Moreover, this cost is achieved by a phased mechanism* $T^* \in \mathcal{T}(\nu^*)$.

The phased mechanism in Proposition 3 has a natural two-part structure. Phase 1 is a common offer whose total value is pinned down by the plaintiffs' strongest litigation threat.

[24] As we show below, when x is weakly convex, a strongest litigation threat always involves either $s = 0$ or $s = n$ plaintiffs committing to litigation, so $\overline{V} = \max\{V_0, V_n\} = \max\{\underline{W}, \overline{W} - nk\}$ in this case.

The continuation phases $q \geq 2$ are divide-and-conquer phases: if phase 1 does not settle the case, the defendant approaches unsettled plaintiffs sequentially and uses fee shifting to make continued coordination against the phase-1 offer unattractive.

To see this, note that the mechanism is designed so that every terminal history gives the plaintiffs a joint payoff no larger than one of the litigation-threat values V_s used to define $\overline{V}$. If the mechanism ends in phase $q \in \{2, \ldots, n\}$ with $q-1$ plaintiffs unsettled, then exactly $q-1$ plaintiffs litigate. When $q = 2$, the litigant rejected the phase-1 offer ν^*, and since $\nu^* = \overline{V}/n \geq V_0/n = x(1)$, fee shifting makes him bear cost k at trial. When $q \in \{3, \ldots, n\}$, phase $q-1$ must have ended with all remaining plaintiffs rejecting $x(q-1)$, so each litigant has rejected an offer at least as large as his trial award and therefore bears cost k. Moreover, any plaintiff who settles receives at most $\max\{\nu^*, x(q)\}$. Hence, in any such terminal history, the plaintiffs' joint payoff is at most

$$(n - q + 1)\max\{\nu^*, x(q)\} + (q-1)(x(q-1) - k).$$

If $x(q) \geq \nu^*$, this expression equals V_{q-1}. If $x(q) < \nu^*$, then it is strictly below $n\nu^*$, since $x(q-1) < x(q) < \nu^*$ and $k > 0$. Finally, if all plaintiffs reject all offers and proceed to trial, the plaintiffs' joint payoff is V_n. Since $n\nu^* = \overline{V} \geq V_s$ for all s, no acceptance profile can yield the plaintiffs a joint payoff strictly above $\overline{V}$, while accepting all phase-1 offers yields exactly $\overline{V}$. Given our tie-breaking assumption, the plaintiffs accept in phase 1 and the defendant settles at total cost $\overline{V}$.

To illustrate, return to the 3-plaintiff example with $(x(1), x(2), x(3)) = (100, 125, 150)$. As shown above, $\overline{V} = 300$ when $k \geq 50$ and $\overline{V} = 450 - 3k$ when $k \leq 50$. Hence $\nu^* = \overline{V}/3$ equals 100 in the high-litigation-cost region and equals $150 - k$ in the low-litigation-cost region. The mechanism in $\mathcal{T}(\nu^*)$ begins by offering ν^* to each plaintiff in phase 1. If at least two plaintiffs accept, the defendant makes no further settlement offers and any remaining plaintiff proceeds to trial. Otherwise, she next approaches plaintiffs who have not yet accepted and offers 125 sequentially, stopping as soon as two plaintiffs have settled (either for ν^* or for 125), or once one plaintiff has settled and the other two have rejected 125. If all three have rejected 125, she approaches the plaintiffs sequentially with an offer of 150, and stops as soon as one plaintiff has settled, or once all have rejected 150.

Figure 1 summarizes the numerical example. In this case, $\underline{W} = 300$ and $\overline{W} = 450$, while Proposition 3 yields $W^* = \overline{V} = \max\{300, 450 - 3k\}$. Thus W^* declines linearly with the litigation cost when $k < 50$, and it is flat once $k \geq 50$. Note that W^* is strictly below $\overline{W}$ for every $k > 0$, and it coincides with $\underline{W}$ once k is sufficiently high. These observations extend beyond the example, as we discuss next.

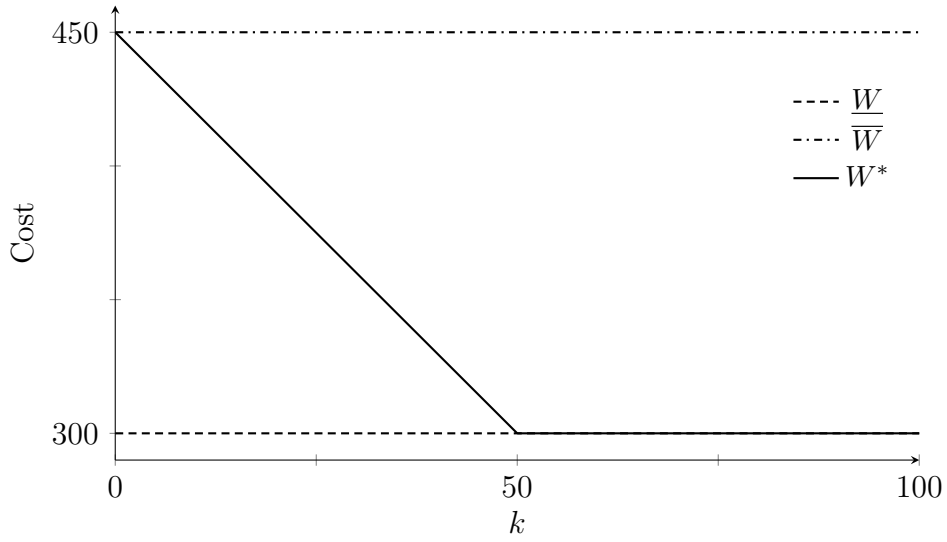


Figure 1: Illustration of the defendant's optimal cost in the 3-plaintiff example. See the text for details.

Two corollaries follow from Proposition 3. First, fee shifting makes divide-and-conquer effective. Indeed, we can verify that each litigation-threat value V_s is strictly below the benchmark cost level $\overline{W} = nx(n)$. Hence $\overline{V} = \max_s V_s < \overline{W}$, and Proposition 3 implies $W^* = \overline{V} < \overline{W}$. Combined with Proposition 1, this shows that fee shifting strictly lowers the defendant's optimal cost relative to the benchmark with side contracting and no fee shifting. This conclusion holds for every $k > 0$, so even arbitrarily low litigation costs suffice to make divide-and-conquer effective.

Corollary 2. *Fee shifting makes divide-and-conquer effective: $W^* < \overline{W}$. In particular, relative to the benchmark with side contracting and no fee shifting, it strictly lowers the defendant's optimal cost and therefore the plaintiffs' joint payoff.*

Second, when litigation costs are sufficiently high, the plaintiffs' strongest litigation threat involves no plaintiff committing to litigation, so $\overline{V}$ collapses to $V_0 = \underline{W}$. Intuitively, any litigation threat that relies on some plaintiffs actually suing exposes those litigants to fee shifting; when k is large, such threats are less attractive than simply insisting that each plaintiff receive at least $x(1)$.

For each $q \in \{2, \ldots, n\}$, define

$$\overline{\kappa}_q := n\Big(\frac{x(q) - x(1)}{q-1}\Big) - \big(x(q) - x(q-1)\big),$$

and let

$$\overline{k} := \max\Big\{x(n) - x(1), \max_{q \in \{2,\ldots,n\}} \overline{\kappa}_q\Big\}.$$

We can verify that $\overline{V} = V_0 = \underline{W}$ whenever $k \geq \overline{k}$. Combined with Proposition 2, this shows that when k is high enough, fee shifting eliminates any benefit of side contracting for the plaintiffs and allows the defendant to capture the entire coordination surplus $\overline{W} - \underline{W}$.

If x is weakly convex, one may take $\overline{k} = x(n) - x(1)$. In that case, a strongest litigation threat always involves either $s = 0$ or $s = n$ plaintiffs committing to litigation, so $\overline{V} = \max\{V_0, V_n\} = \max\{\underline{W}, \overline{W} - nk\}$, and therefore $\overline{V} = \underline{W}$ whenever $nx(1) \geq nx(n) - nk$.

Corollary 3. *If $k \geq \overline{k}$, then the defendant's optimal cost satisfies $W^* = \underline{W}$. In particular, fee shifting allows the defendant to capture the entire surplus $\overline{W} - \underline{W}$, and the plaintiffs' joint payoff is the same as when side contracting is unavailable. If x is weakly convex, one may take $\overline{k} = x(n) - x(1)$.*

6 Bargaining protocols

In our baseline model, the defendant can commit to a dynamic mechanism with multiple rounds of offers, and the plaintiffs can only accept or reject in each round. This section considers two alternative bargaining protocols. The first restricts the defendant to a single round of simultaneous (and hence bilateral) offers. The second allows plaintiffs to make counteroffers. While either modification can raise the defendant's optimal cost relative to the baseline, fee shifting continues to make divide-and-conquer effective in each case.

6.1 Bilateral offers

Our baseline model allows the defendant to condition offers in each round on the history of acceptances and rejections. We now impose the one-round restriction $R = 1$. A mechanism is therefore an offer vector $t = (t_i)_{i \in N}$, and each plaintiff $i \in N$ can either accept or reject his offer t_i. Denote the defendant's optimal cost under such bilateral offers by W_b^*.

As in Section 4.1, the defendant's optimal cost in the absence of a fee-shifting rule would be $\overline{W} = nx(n)$. We show that, even with $R = 1$, fee shifting allows the defendant to reduce her cost strictly below this benchmark level. She can do so by offering one plaintiff the full amount $x(n)$, and offering each remaining plaintiff a lower amount that is nonetheless high enough to make rejection unattractive under fee shifting. Specifically, consider the one-round offer vector $\hat{t}$ with $\hat{t}_1 = x(n)$ and

$$\hat{t}_i = \max\left\{x(n) - \frac{k}{n-1},\, x(n-1)\right\}$$

for all $i \in \{2, \ldots, n\}$. If all plaintiffs accept, the defendant pays strictly less than $\overline{W}$. At the same time, the offers are chosen so that any other acceptance profile leaves the plaintiffs

with a joint payoff no larger than that from accepting. Thus $\hat{t}$ induces full settlement, implying $W_b^* < \overline{W}$. That is, fee shifting makes divide-and-conquer effective even under the bilateral-offers restriction.

To characterize how low the defendant's cost can be under bilateral offers, let

$$\underline{W}_b := \sum_{i=1}^{n} x(i).$$

The quantity $\underline{W}_b$ arises naturally as a lower bound on W_b^*. Consider any offer vector that induces full settlement, and order the offers from lowest to highest. To induce full settlement, the defendant cannot let the sth-lowest offer fall strictly below $x(s)$: otherwise the s plaintiffs receiving the lowest offers could all reject and litigate together, and each would obtain $x(s)$ at trial without bearing litigation costs under fee shifting. Summing these constraints over $s = 1, \ldots, n$ yields $\underline{W}_b$.

Remark 2. *A complementary way to motivate the lower bound $\underline{W}_b$ is to consider the bilateral-offers environment when side contracts are not available, but plaintiffs coordinate on selecting the equilibrium that maximizes their joint payoff (equivalently, the defendant is concerned with her worst-case equilibrium). Results from the contracting-with-externalities literature (see, e.g. Segal, 2003) imply that the defendant's optimal worst-case offer vector then yields a total cost of $\underline{W}_b$. Since allowing binding side contracts can only expand the set of joint payoffs the plaintiffs can secure, the defendant cannot achieve a strictly lower cost when side contracting is available, and $\underline{W}_b$ remains a lower bound on W_b^*.*

The next proposition summarizes these observations and shows that the lower bound $\underline{W}_b$ is attainable when litigation costs are sufficiently high.

Proposition 4. *Suppose only bilateral offers are feasible. The defendant's optimal cost satisfies*

$$\underline{W}_b \leq W_b^* < \overline{W}.$$

Moreover, there exists $\overline{k}_b$ such that, if the litigation cost is $k \geq \overline{k}_b$, then the defendant's optimal cost is $W_b^* = \underline{W}_b$.

As noted in Remark 2, $\underline{W}_b$ is the defendant's optimal cost when plaintiffs cannot use side contracts but can coordinate on their preferred equilibrium. Since Proposition 4 shows that this level is attained for high litigation costs, the following corollary is immediate.

Corollary 4. *Suppose only bilateral offers are feasible. Under high enough litigation costs, fee shifting allows the defendant to capture the entire surplus $\overline{W} - \underline{W}_b$, and the plaintiffs' joint payoff is the same as in their preferred equilibrium when side contracting is unavailable.*

To see why the lower bound $\underline{W}_b$ is attainable when litigation costs are high, consider the bilateral-offers mechanism that assigns the offers $\{x(1), \ldots, x(n)\}$ one-to-one across plaintiffs (e.g., $t_i = x(i)$ for each $i \in \{1, \ldots, n\}$). If all plaintiffs accept, the defendant's cost is $\underline{W}_b$. If instead exactly $s \in \{1, \ldots, n\}$ plaintiffs reject, then the plaintiffs maximize their joint payoff by having the s lowest offers be rejected. In that case the trial award is $x(s)$ and at least one litigant has rejected an offer of $x(s)$, so fee shifting makes at least one litigant bear cost k. A sufficient condition ensuring that accepting all offers maximizes the plaintiffs' joint payoff is therefore

$$k \geq \overline{W} - \underline{W}_b = \sum_{i \in N} (x(n) - x(i)).$$

Remark 3. *The threshold $\overline{k}_b$ in Proposition 4 can be chosen as*

$$\overline{k}_b := \sum_{i \in N} (x(n) - x(i)).$$

In the 3-plaintiff example of Section 5, with claim values $(x(1), x(2), x(3)) = (100, 125, 150)$, we have $\overline{W} = 450$ and $\underline{W}_b = 100 + 125 + 150 = 375$. Proposition 4 then implies $375 = \underline{W}_b \leq W_b^* < \overline{W} = 450$. Moreover, the litigation-cost threshold can be chosen as $\overline{k}_b = 75$, so that $W_b^* = \underline{W}_b = 375$ whenever $k \geq 75$.

6.2 Mutual offers

Our baseline model assumes only the defendant can make settlement offers to the plaintiffs. We now consider a bargaining setting with mutual offers. In each round, the timing is as follows: first, the defendant makes offers to any subset of plaintiffs who have not yet settled; second, each plaintiff who receives an offer chooses whether to accept or reject; third, each plaintiff who has not yet settled may make an offer to the defendant; finally, the defendant chooses whether to accept or reject each offer she receives.

As in our baseline, the defendant commits to a dynamic mechanism—here specifying both her offers and her acceptance decisions following plaintiffs' offers as a function of the history—and plaintiffs can coordinate through binding side contracts. We extend the fee-shifting rule to account for offers made by both sides as follows. If a plaintiff litigates and is awarded damages of $x(s)$, then he bears the defendant's litigation cost k if and only if he rejected some defendant offer at least $x(s)$ and, after that rejection, did not make any counteroffer weakly below $x(s)$.[25] Given this, we may assume without loss that if a plaintiff rejects an offer and makes a counteroffer in the same round, the counteroffer is strictly higher

[25] We do not consider plaintiffs' own litigation costs here, but see Section 7.2 for an analysis.

than the rejected offer.[26] To rule out delayed counteroffers that would otherwise undermine the fee-shifting logic, we also assume that if a plaintiff rejects an offer and does not make a counteroffer weakly below that offer in the same round, then his ability to make further offers lapses until either the defendant makes him a new offer or another plaintiff settles.[27]

We denote the defendant's optimal cost under mutual offers by W_m^*. The next proposition shows that allowing plaintiffs to make counteroffers has an effect similar to the one-round restriction studied in Section 6.1: fee shifting continues to make divide-and-conquer effective, in the sense that the defendant can still reduce her cost strictly below $\overline{W}$, but the relevant lower bound is again $\underline{W}_b$ rather than $\underline{W}$. Intuitively, counteroffers allow plaintiffs to place on the record an offer that prevents adverse fee shifting at trial.[28] When plaintiffs coordinate, they can use this option to force the defendant to "buy out" claims at the relevant trial values, much as in the bilateral-offers environment.

Proposition 5. *Suppose mutual offers are feasible. The defendant's optimal cost satisfies*

$$\underline{W}_b \leq W_m^* < \overline{W}.$$

Moreover, there exists $\overline{k}_m$ such that, if the litigation cost is $k \geq \overline{k}_m$, then the defendant's optimal cost is $W_m^* = \underline{W}_b$.

To see why $\underline{W}_b$ is attainable for high litigation costs, consider a mechanism that builds on the logic of Section 6.1. While u plaintiffs remain unsettled, the defendant approaches them sequentially and offers $x(u)$ until one plaintiff accepts, in which case the mechanism moves on to offers of $x(u-1)$; if all u plaintiffs reject, they proceed to trial. The defendant also commits to accept any counteroffer from the plaintiff she is addressing that is weakly below $x(u)$, and to reject all higher counteroffers and all offers from other plaintiffs. On the intended path, exactly one plaintiff settles at each value $x(n), x(n-1), \ldots, x(1)$, so the defendant's total cost is $\underline{W}_b$. If instead the mechanism ends with s plaintiffs at trial, then each must have rejected an offer of $x(s)$ without making a counteroffer weakly below $x(s)$, so fee shifting makes each litigant bear cost k. We derive a threshold $\overline{k}_m$ such that, whenever $k \geq \overline{k}_m$, any such terminal history involving litigation yields the plaintiffs a joint payoff no

[26] The defendant can always commit to accept any counteroffer weakly below her last offer (which weakly reduces her cost), so a counteroffer weakly below a rejected offer would be accepted and is weakly dominated for the plaintiffs by accepting the rejected offer.

[27] This assumption prevents plaintiffs from waiting until the terminal round to place a counteroffer on the record. Our results are unchanged if, instead, any delayed counteroffer following a rejected defendant offer resets the clock, giving the defendant a fresh R-round continuation from that history.

[28] Here counteroffers matter only through the fee-shifting rule. Because the defendant commits to a mechanism specifying both her offers and her responses to plaintiffs' counteroffers after every history, allowing counteroffers would not otherwise change her optimal cost.

greater than $\underline{W}_b$.

Remark 4. *The threshold $\overline{k}_m$ in Proposition 5 can be chosen as*

$$\overline{k}_m := \max_{s\in\{1,\dots,n\}} \frac{1}{s}\sum_{i=1}^{s}(x(s)-x(i)).$$

In particular, $\overline{k}_m \le \overline{k}_b/2 < \overline{k}_b$, where $\overline{k}_b$ is the threshold in Remark 3. If x is weakly convex, the maximum is attained at $s=n$, so $\overline{k}_m = \overline{k}_b/n$.

In the 3-plaintiff example with $(x(1),x(2),x(3)) = (100,125,150)$, Proposition 5 implies $375 = \underline{W}_b \le W_m^* < \overline{W} = 450$. Moreover, the litigation-cost threshold can be chosen as $\overline{k}_m = 25$, so that $W_m^* = \underline{W}_b = 375$ whenever $k \ge 25$.

7 Discussion

This section explores three variants of our model. We first allow for heterogeneous plaintiffs, then introduce plaintiffs' own litigation costs, and finally consider a setting with uncertain liability. These variants show that the divide-and-conquer logic behind our main results applies more broadly, but also that its force depends on the nature of the trial environment.

7.1 Heterogeneous plaintiffs

Our baseline model considers symmetric plaintiffs. We now allow plaintiffs to differ in both their claim values and the defendant's litigation costs. Specifically, suppose that if $s \in \{1,\dots,n\}$ plaintiffs litigate, plaintiff i's trial award is $x_i(s)$, where each function $x_i : \{1,\dots,n\} \to \mathbb{R}_+$ is strictly increasing.[29] The defendant's litigation cost attributable to plaintiff i is $k_i > 0$; as in our baseline model, $k_i < x_i(1)$ for all $i \in N$. The fee-shifting rule applies plaintiff by plaintiff: if s plaintiffs litigate, litigating plaintiff i bears cost k_i if and only if he previously rejected an offer weakly above $x_i(s)$.

The no-fee-shifting and no-side-contract benchmarks extend readily to this heterogeneous-plaintiff setting. The defendant's optimal settlement costs in those benchmarks are, respectively,

$$\overline{W}_h := \sum_{i\in N} x_i(n), \qquad \underline{W}_h := \sum_{i\in N} x_i(1).$$

Under fee shifting and side contracting, we can define plaintiffs' litigation threats analo-

[29] We maintain the assumption that plaintiff i's trial award is independent of which other plaintiffs litigate; only the total number of litigating plaintiffs matters.

gous to those in Section 5. For $S \subsetneq N$, let

$$V(S) := \sum_{i \in N \setminus S} x_i(|S|+1) + \sum_{i \in S}(x_i(|S|) - k_i),$$

and for $S = N$, let

$$V(N) := \sum_{i \in N}(x_i(n) - k_i).$$

Thus $V(S)$ is the joint payoff plaintiffs can guarantee when the plaintiffs in S commit to litigate, while each plaintiff outside S accepts only offers at least as high as the award he would obtain if he joined them. The strongest such litigation threat, $\overline{V} := \max_{S \subseteq N} V(S)$, yields a lower bound on the defendant's optimal cost. The next proposition shows that this lower bound remains tight. We denote by W_h^* the defendant's optimal cost in this heterogeneous-plaintiff setting.

Proposition 6. *Consider the heterogeneous-plaintiff setting. The defendant's optimal cost is $W_h^* = \overline{V}$. Moreover, this cost is achieved by a modified phased mechanism.*

The modified phased mechanism replaces the common first-phase offer in the phased mechanisms of Section 5 with personalized offers. For each nonempty set of unsettled plaintiffs $U \subseteq N$, and each $S \subsetneq U$, define

$$V_U(S) := \sum_{i \in U \setminus S} x_i(|S|+1) + \sum_{i \in S}(x_i(|S|) - k_i),$$

and

$$V_U(U) := \sum_{i \in U}(x_i(|U|) - k_i).$$

Letting $\overline{V}_U := \max_{S \subseteq U} V_U(S)$ and $\overline{V}_\emptyset := 0$, we thus have $V(S) = V_N(S)$ and $\overline{V} = \overline{V}_N$. Starting with $U = N$, the defendant selects some plaintiff $i \in U$ and offers him $\overline{V}_U - \overline{V}_{U \setminus \{i\}}$. If plaintiff i accepts, he settles, the set of unsettled plaintiffs becomes $U \setminus \{i\}$, and the defendant repeats the same step. If plaintiff i rejects, the mechanism proceeds to divide-and-conquer phases analogous to those in Section 5, as described in the proof of Proposition 6.

The proposition implies that our findings are robust to heterogeneity. Since $\overline{V} < \overline{W}_h$, fee shifting continues to make divide-and-conquer effective, strictly lowering the defendant's optimal cost relative to the benchmark with side contracting and no fee shifting. Moreover, for high enough litigation costs, the defendant attains the no-side-contract benchmark cost

$\underline{W}_h$. This holds if

$$\sum_{i\in N\setminus S}(x_i(|S|+1)-x_i(1))+\sum_{i\in S}(x_i(|S|)-x_i(1))\leq\sum_{i\in S}k_i$$

for every $S\subsetneq N$, and

$$\sum_{i\in N}(x_i(n)-x_i(1))\leq\sum_{i\in N}k_i,$$

so that $V(S)\leq V(\emptyset)=\underline{W}_h$ for every $S\subseteq N$.

7.2 Plaintiffs' litigation costs

Consider again the symmetric baseline model, but suppose that, in addition to the defendant's litigation cost $k>0$, each plaintiff who litigates incurs a litigation cost $c>0$. Analogous to the baseline assumption $k<x(1)$, assume $c+k<x(1)$, so that any plaintiff who rejects all settlement offers strictly prefers pursuing the case in court rather than abandoning it.

Without a fee-shifting rule, the defendant's optimal cost is $\overline{W}_c:=\overline{W}-nc=n(x(n)-c)$, regardless of the bargaining protocol. The reason is simple: plaintiffs can always secure $\overline{W}_c$ by rejecting all offers and litigating together, while the defendant can induce full settlement at that same cost by offering $x(n)-c$ to each plaintiff and, under mutual offers, committing to reject any counteroffer strictly above that level.

Now consider the defendant's optimal cost in the presence of fee shifting. Under the baseline and bilateral-offers protocols, fee shifting weakens any coordinated response by the plaintiffs that involves litigation and therefore cannot hurt the defendant. Proposition 7 below shows that fee shifting can give the defendant a divide-and-conquer advantage relative to the benchmark without fee shifting: her optimal cost is weakly below $\overline{W}_c$ for all c, and strictly below this level when c is sufficiently small.

The mutual-offers protocol is more subtle. In that environment, fee shifting extends symmetrically to plaintiffs' litigation costs by applying the same last-offer logic to both sides. Thus, if s plaintiffs litigate and plaintiff i is awarded damages of $x(s)$, then the defendant bears plaintiff i's litigation cost c if and only if she rejected some offer by plaintiff i weakly below $x(s)$ and, after that rejection, did not make any counteroffer to i weakly above $x(s)$. Because plaintiffs can then place offers on the record that protect them against bearing c at trial, two-sided fee shifting can either benefit or harm the defendant relative to the benchmark without fee shifting. Proposition 7 shows that it benefits the defendant whenever her cost k is sufficiently high and plaintiffs' costs c are sufficiently low.

Let W_c^*, $W_{b,c}^*$, and $W_{m,c}^*$ denote the defendant's optimal costs in the baseline, bilateral-

offers, and mutual-offers protocols, respectively, when each litigating plaintiff incurs litigation cost c.

Proposition 7. *Consider the setting with plaintiffs' litigation costs c. If only offers from the defendant are feasible, then the defendant's optimal cost satisfies*

$$W_c^* \leq \overline{W}_c \quad \textit{and} \quad W_{b,c}^* \leq \overline{W}_c,$$

with strict inequality in both cases for c small enough.

If mutual offers are feasible and the fee-shifting rule is two-sided, then the defendant's optimal cost can lie above or below $\overline{W}_c$, depending on parameters. If $k \geq \overline{k}_m$, then $W_{m,c}^* = \underline{W}_b$ for all c, and therefore $W_{m,c}^* < \overline{W}_c$ for c small enough.

To understand the mutual-offers part, recall that when $c = 0$, the plaintiffs can guarantee themselves a joint payoff of at least $\underline{W}_b$ by using counteroffers to prevent the defendant from shifting her litigation cost k to them. This lower bound remains valid when $c > 0$ and fee shifting is two-sided: plaintiffs can also protect themselves against bearing c at trial. At the same time, when $k \geq \overline{k}_m$, the defendant can still attain this lower bound, just as in Proposition 5. Hence her optimal cost remains equal to $\underline{W}_b$ for all c. Since the no-fee-shifting benchmark cost is $\overline{W}_c = \overline{W} - nc$, it follows that two-sided fee shifting helps the defendant for all $c < (\overline{W} - \underline{W}_b)/n$.[30]

7.3 Uncertain liability

Our analysis has assumed that plaintiffs' awards at trial are a deterministic function of the number of plaintiffs who litigate. We now return to the baseline bargaining protocol and consider a variant of the model in which trial outcomes are uncertain. Specifically, suppose plaintiffs are awarded damages of $\overline{x} > 0$ each if they win in court, independent of the number of plaintiffs who sue. However, the probability that plaintiffs win depends on how many sue: if $s \in \{1, \ldots, n\}$ plaintiffs litigate, then all of them win with probability $p(s)$ and all lose otherwise. The function $p : \{1, \ldots, n\} \to [0, 1]$ is strictly increasing, so the plaintiffs' litigation decisions are strategic complements.

The fee-shifting rule is applied based on the realized award. If the plaintiffs lose, the realized award is 0, so any litigating plaintiff bears the defendant's litigation cost k (recall that not making an offer is equivalent to offering 0, and offers are nonnegative). If the plaintiffs win, the realized award is $\overline{x}$, so a litigating plaintiff bears cost k if and only if he

[30] If, in addition, $\overline{k}_m \leq k < x(1) - (\overline{W} - \underline{W}_b)/n$, then two-sided fee shifting hurts the defendant for all $c \in ((\overline{W} - \underline{W}_b)/n, x(1) - k)$.

previously rejected a defendant offer weakly above $\overline{x}$. Thus, for a litigating plaintiff there are two relevant expected payoffs. If he has not rejected any offer weakly above $\overline{x}$, then his expected payoff when s plaintiffs litigate is

$$\mu(s) := \overline{x}p(s) - k(1 - p(s)).$$

If instead he has rejected some offer weakly above $\overline{x}$, then his expected payoff becomes $\overline{x}p(s) - k = \mu(s) - p(s)k$. Since p is strictly increasing, $\mu(s)$ is strictly increasing in s. Analogous to our assumptions in the baseline model, we assume $k < \overline{x}p(1)$, so that any plaintiff who rejects all settlement offers strictly prefers pursuing the case in court rather than abandoning it. In particular, this implies $\mu(1) > 0$.

Observe that, unlike in our baseline model with deterministic awards, here the fee-shifting rule yields a strict benefit to the defendant even in the case of a single plaintiff. With one plaintiff, the defendant can settle for a total cost of $\overline{x}p(1) - k(1 - p(1)) = \mu(1)$, which is strictly lower than the plaintiff's expected award $\overline{x}p(1)$. Thus the relevant question is whether the presence of multiple plaintiffs generates an additional divide-and-conquer advantage beyond this one-plaintiff effect. Proposition 8 below shows that when there are $n \geq 2$ plaintiffs, fee shifting can indeed give the defendant an additional advantage, in the sense that her optimal total cost can be strictly lower than $n\mu(n)$. On the other hand, unlike in our baseline, the defendant cannot extract the full surplus from plaintiffs' coordination: her cost is always strictly greater than $n\mu(1)$.

Let W^*_μ denote the defendant's optimal cost in the uncertain-liability setting. Define $\overline{W}_\mu := n\mu(n)$ and $\underline{W}_\mu := n\mu(1)$.

Proposition 8. *Consider the uncertain-liability setting. The defendant's optimal cost satisfies*

$$\underline{W}_\mu < W^*_\mu \leq \overline{W}_\mu.$$

Moreover, fixing $\overline{x}$, k, and $(p(1), \ldots, p(n-1))$, there exists $\varepsilon > 0$ such that if $p(n) \geq 1 - \varepsilon$, then $W^*_\mu < \overline{W}_\mu$.

The defendant can guarantee a cost no greater than $\overline{W}_\mu$ by offering $\mu(n)$ to each plaintiff. Under any acceptance profile, each settling plaintiff receives $\mu(n)$, while each litigant's expected payoff is at most $\mu(n)$. Hence, the plaintiffs' joint payoff is at most $n\mu(n) = \overline{W}_\mu$, and by tie-breaking they accept.

The proposition further shows that fee shifting can make divide-and-conquer effective when $p(n)$ is close to 1. As in our main model, fee shifting gives the defendant a way to weaken the plaintiffs' litigation threat, but this now requires her to make an offer as high

as $\overline{x}$ to at least one plaintiff. Specifically, by doing so the defendant can ensure that if all n plaintiffs litigate and win, at least one litigant bears the cost k. This reduces the plaintiffs' expected payoff from going to trial, and when $p(n)$ is sufficiently high, this reduction more than compensates for the extra cost of making a guaranteed payment of $\overline{x}$ to one plaintiff.

Proposition 8 also shows that the defendant's optimal cost is bounded below by $\underline{W}_\mu$; however, in contrast with our baseline findings, she cannot achieve this lower bound, no matter how high litigation costs are. The reason is that the plaintiffs can guarantee themselves strictly more than $\underline{W}_\mu$ by having each plaintiff accept any offer weakly above $\overline{x}$ and reject all strictly lower offers. Under this strategy, if at least one plaintiff settles, some plaintiff receives at least $\overline{x} > \mu(1)$, while each litigant obtains at least $\mu(1)$ because no litigating plaintiff has rejected an offer weakly above $\overline{x}$. If instead all plaintiffs litigate, then each obtains $\mu(n) > \mu(1)$. Thus, regardless of the defendant's mechanism, the plaintiffs can secure strictly more than $\underline{W}_\mu = n\mu(1)$, and therefore the defendant's optimal cost is always strictly above this level.

8 Concluding remarks

This paper has studied settlement bargaining between a defendant and multiple plaintiffs whose claims exhibit complementarities. Plaintiffs can use binding side contracts, so their litigation decisions maximize their joint payoff. Yet litigation costs are allocated according to fee-shifting rules that apply plaintiff by plaintiff as a function of previously rejected settlement offers. We showed that this feature overturns the usual Coasean intuition: a divide-and-conquer strategy is effective even against coordinated plaintiffs. In fact, when the defendant's litigation costs are sufficiently high, divide-and-conquer can allow the defendant to do as well as if side contracts were unavailable to the plaintiffs. We examined these effects under different bargaining protocols and trial environments.

While we have focused on fee-shifting rules, other legal settings may give rise to similar effects. A closely related application is eminent domain. When a public authority seeks to assemble land for a road, railway, or redevelopment project, it may negotiate separately with multiple owners who can nevertheless coordinate by sharing appraisals, expert testimony, and legal work. An owner's payoff from resisting may be higher when more other owners also resist, because collective resistance can strengthen the case for higher compensation. In certain jurisdictions, rejected offers affect the cost consequences of condemnation proceedings: an owner who rejects and later fails to obtain a sufficiently better award may bear additional litigation costs or lose the right to recover related expenses. These offer-contingent cost consequences can undermine coordination among owners.

More broadly, the mechanism we study is not unique to legal settings. It does not depend on formal fee shifting, or even on any transfer of costs from a principal to agents. What matters instead is that a principal bargains separately with multiple agents who can coordinate, each agent's payoff from resisting is higher when more others also resist, and rejecting an offer may expose an agent to costs if the eventual outcome is no more favorable than the rejected terms. In our setting, those costs take the form of fee shifting, but similar costs may arise from reputational, political, or organizational consequences of declining an offer and later doing no better. Such costs may also be psychological, such as regret after rejecting an offer and then failing to improve on it. Side contracts can compensate agents for bearing these costs, and thereby redistribute their burden across agents, but they cannot eliminate them once they are triggered by separate rejected offers.

One setting in which this logic seems natural is labor bargaining. The agents in this context may be union leaders or local worker representatives bargaining on behalf of distinct worker groups that can coordinate their bargaining positions. A group's payoff from resisting may be higher when more other groups also resist, as coordinated resistance can limit the employer's ability to isolate any one group or shift concessions across groups. Ashenfelter and Johnson (1969) stress that union leaders who bring back an agreement below members' expectations face internal dissension and ratification problems. In such an environment, rejecting an employer's offer and later securing no more favorable terms may impose a political cost on the leader, thereby weakening the gains from coordination across worker groups.

A similar logic may also apply to political bargaining, for example in legislative bargaining or international negotiations. Support is often assembled through separate negotiations with multiple actors who benefit from coordinating their responses. When those actors are accountable to constituents or other audiences, publicly rejecting a compromise and later accepting no more favorable terms may entail audience costs in the spirit of Fearon (1994). Such costs could give rise to divide-and-conquer effects similar to those of fee shifting, and thereby undermine coordination.

References

Ashenfelter, Orley and George E. Johnson, "Bargaining Theory, Trade Unions, and Industrial Strike Activity," *American Economic Review*, 1969, *59* (1), 35–49.

Bebchuk, Lucian Arye and Howard F. Chang, "The Effect of Offer-of-Settlement Rules on the Terms of Settlement," *Journal of Legal Studies*, 1999, *28* (2), 489–513.

Bernstein, Shai and Eyal Winter, "Contracting with Heterogeneous Externalities," *American Economic Journal: Microeconomics*, 2012, *4* (2), 50–76.

Che, Yeon-Koo, "The Economics of Collective Negotiation in Pretrial Bargaining," *International Economic Review*, 2002, *43* (2), 549–576.

_ **and Kathryn E. Spier**, "Exploiting Plaintiffs through Settlement: Divide and Conquer," *Journal of Institutional and Theoretical Economics*, 2008, *164* (1), 4–23.

Chen, Kong-Pin and Jue-Shyan Wang, "Fee-Shifting Rules in Litigation with Contingency Fees," *Journal of Law, Economics, and Organization*, 2007, *23* (3), 519–546.

Choi, Albert H. and Kathryn E. Spier, "Class Actions and Private Antitrust Litigation," *American Economic Journal: Microeconomics*, 2022, *14* (3), 131–163.

Chung, Tai-Yeong, "Settlement of Litigation under Rule 68: An Economic Analysis," *Journal of Legal Studies*, 1996, *25* (1), 261–286.

Coase, R. H., "The Problem of Social Cost," *Journal of Law and Economics*, 1960, *3*, 1–44.

Dal Bó, Ernesto, "Bribing Voters," *American Journal of Political Science*, 2007, *51* (4), 789–803.

Daughety, Andrew F. and Jennifer F. Reinganum, "A Dynamic Model of Lawsuit Joinder and Settlement," *RAND Journal of Economics*, 2011, *42* (3), 471–494.

Farmer, Amy and Paul Pecorino, "Conditional Cost Shifting and the Incidence of Trial: Pretrial Bargaining in the Face of a Rule 68 Offer," *American Law and Economics Review*, 2000, *2* (2), 318–341.

Fearon, James D., "Domestic Political Audiences and the Escalation of International Disputes," *American Political Science Review*, 1994, *88* (3), 577–592.

Halac, Marina, "Contracting for Coordination," *Journal of the European Economic Association*, 2025, *23* (3), 815–844.

_ **, Elliot Lipnowski, and Daniel Rappoport**, "Pricing for Coordination," 2026. Working paper.

_ **, Ilan Kremer, and Eyal Winter**, "Raising Capital from Heterogeneous Investors," *American Economic Review*, 2020, *110* (3), 889–921.

Hay, Bruce L. and Kathryn E. Spier, "Settlement of Litigation," in Peter Newman, ed., *The New Palgrave Dictionary of Economics and the Law*, Vol. 3, New York: Stockton Press, 1998, pp. 442–451.

Hua, Xinyu and Kathryn E. Spier, "Information and Externalities in Sequential Litigation," *Journal of Institutional and Theoretical Economics*, 2005, *161* (2), 215–232.

Iaryczower, Matias and Santiago Oliveros, "Competing for Loyalty: The Dynamics of Rallying Support," *American Economic Review*, 2017, *107* (10), 2990–3005.

Klement, Alon and Zvika Neeman, "Against Compromise: A Mechanism Design Approach," *The Journal of Law, Economics, and Organization*, 2005, *21* (2), 285–314.

Miller, Geoffrey P., "An Economic Analysis of Rule 68," *Journal of Legal Studies*, 1986, *15* (1), 93–125.

Segal, Ilya, "Contracting with Externalities," *Quarterly Journal of Economics*, 1999, *114* (2), 337–388.

_ , "Coordination and Discrimination in Contracting with Externalities: Divide and Conquer?," *Journal of Economic Theory*, 2003, *113* (2), 147–181.

Spier, Kathryn E., "Pretrial Bargaining and the Design of Fee-Shifting Rules," *RAND Journal of Economics*, 1994, *25* (2), 197–214.

_ , "Settlement with Multiple Plaintiffs: The Role of Insolvency," *Journal of Law, Economics, and Organization*, 2002, *18* (2), 295–323.

_ , "The Use of "Most-Favored-Nation" Clauses in Settlement of Litigation," *RAND Journal of Economics*, 2003, *34* (1), 78–95.

_ , "Litigation," in A. Mitchell Polinsky and Steven Shavell, eds., *Handbook of Law and Economics*, Vol. 1, Amsterdam: Elsevier, 2007, chapter 4, pp. 259–342.

Winter, Eyal, "Incentives and Discrimination," *American Economic Review*, 2004, *94* (3), 764–773.

A Supporting references on fee-shifting rules

U.S. law

United States Courts. *Federal Rules of Civil Procedure*. Rule 68: "Offer of Judgment."

California. *Code of Civil Procedure* § 998 (enacted 1971).

Florida. *Statutes* § 768.79 (enacted 1986, amended 2022).

New Jersey. *Court Rules*, R. 4:58-3 (rev. 2024).

Texas. *Civil Practice and Remedies Code* ch. 42 (added 2003).

Common-law jurisdictions

Australia

New South Wales. *Uniform Civil Procedure Rules 2005* (NSW), r. 20.26.

Beazley, Margaret. "Offers of Compromise: Following, Bending and Breaking the Rules." Speech presented at the NSW Law Society, Sydney, August 5, 2015. Supreme Court of New South Wales.

Canada

Federal Courts Rules, SOR/98-106, r. 420.

Sgayias, David, Meg N. Kinnear, Donald J. Rennie, and Brian J. Saunders. *Federal Court Practice, 2004*. Toronto: Carswell, 2003, 873–874.

Ontario. *Rules of Civil Procedure*, R.R.O. 1990, Reg. 194, r. 49.

Akazaki, R. Lee. "Peering over the Plaintiff's and Defendants' Parapets: Redesigning Ontario Rule 49 to Encourage Meaningful Offers of Settlement." *Advocates' Quarterly* 30, no. 3 (2005): 285–296.

England

Ministry of Justice (UK). *Civil Procedure Rules*, Pt. 36, "Offers to Settle," § 36.17.

Ward, Daniel L. "New Carrots and Sticks: Proposals for Reform of CPR Part 36," *The Modern Law Review* 70, no. 2 (2007): 278–293, at 280–283.

Hong Kong

Rules of the High Court (Cap. 4A), O. 22.

Meggitt, Gary. "Civil Justice Reform in Hong Kong — Its Progress and Its Future." *Hong Kong Law Journal* 38 (2008): 89–128, at 111–114.

Singapore

Rules of Court (Cap. 322, R 5, 2014 Rev. Ed.), O. 22A.

Chen, Siyuan, and Eunice Hui Han Chua. "Recent Developments in the Offer to Settle Regime in Singapore — Accepting an Offer to Settle before Damages Are Assessed and the Contra Proferentem Rule: *Ong & Ong Pte Ltd. v. Fairview Development Pte Ltd* [2014] 2

SLR 1285." *Singapore Academy of Law Journal* 26, no. 2 (2014): 687–699, at 688–690.

Civil-law jurisdictions

Italy

Codice di procedura civile [Code of Civil Procedure], art. 91, *Gazzetta Ufficiale della Repubblica Italiana*.

De Luca, Alessandra. "Italy: A Tale of Successful Resistance?" In *Cost and Fee Allocation in Civil Procedure*, edited by Mathias Reimann, 185–194. Berlin: Springer, 2012.

Spain

Ley 1/2000, de 7 de enero, de Enjuiciamiento Civil [Civil Procedure Act], art. 245(5), *Boletín Oficial del Estado*, no. 7 (8 Jan. 2000), as amended by *Ley Orgánica 1/2025*, *Boletín Oficial del Estado*, no. 3 (3 Jan. 2025).

Arbitration and international tribunals

Born, Gary B. *International Commercial Arbitration*. 3rd ed. Alphen aan den Rijn: Kluwer Law International, 2021.

International Chamber of Commerce. *Decisions on Costs in International Arbitration: ICC Commission Report*. Paris: International Chamber of Commerce, 2015.

New York City Bar Association. *Awards of Costs in International Arbitration in New York: Current Approaches and Best Practices*. New York: New York City Bar Association, 2024.

B Proofs for Section 3 and Section 4

B.1 Proof of Lemma 1

We proceed in four steps.

Step 1: We claim that it is without loss to consider mechanisms that induce full settlement (i.e., no plaintiff goes to court along the induced history). Fix a mechanism $T \in \mathcal{T}$, and let h be the history selected by the plaintiffs' (coordinated) response under our tie-breaking rule. Suppose that h ends with a set $S \subseteq N$ of plaintiffs going to court, with $s := |S| \geq 1$.

Construct a new mechanism that coincides with the original one except that, following history h, it includes a final round (appending one additional last round if needed, which is without loss since R is arbitrarily large) in which the defendant offers $t_i = x(s)$ to each

plaintiff $i \in S$. If all plaintiffs in S accept, their joint payoff is $sx(s)$. Consider instead any response in this final round in which a nonempty subset $L \subseteq S$ rejects and goes to court, with $\ell := |L| \geq 1$. Then the litigating plaintiffs in L are awarded damages of $x(\ell)$. Moreover, since each $i \in L$ has received a settlement offer of $x(s) \geq x(\ell)$, the fee-shifting rule implies that each such plaintiff bears his litigation cost k. Hence, under this response, the joint payoff to the plaintiffs in S is

$$(s - \ell)x(s) + \ell(x(\ell) - k) < sx(s),$$

where the strict inequality uses that x is strictly increasing and $k > 0$. Therefore, all plaintiffs $i \in S$ accept in the added final round.

The proposed modification changes the mechanism only after history h and leaves payoffs following any history that diverges from h unchanged. Since it (weakly) increases the plaintiffs' payoff conditional on reaching h, it does not affect the induced play prior to the terminal stage. At the same time, the modification weakly reduces the defendant's total cost, since it replaces litigation by settlement and thus weakly saves on litigation costs. It follows that it is without loss to restrict attention to mechanisms that induce full settlement.

Step 2: We claim that it is without loss to consider mechanisms that admit an all-accept history in which every plaintiff receives (and accepts) a strictly positive offer. Consider a mechanism satisfying Step 1. Let h_{yes} be a history in which every strictly positive offer is accepted. It may be that along h_{yes} some set $S \subseteq N$ of plaintiffs never receives a strictly positive offer and would therefore go to court. Let $s := |S|$. Modify the mechanism only by adding a final round following h_{yes} (appending one additional last round if needed) in which the defendant offers $t_i = x(s)$ to each $i \in S$. By the same argument as in Step 1, all plaintiffs in S accept, and the defendant's total cost weakly declines. Hence, without loss, there exists an all-accept history h_{yes} in which every plaintiff receives (and accepts) a strictly positive offer.

Step 3: We claim that it is without loss to consider mechanisms that make at most one strictly positive offer per round (after any history). Suppose that at some history and round r the mechanism makes $\ell_r > 1$ strictly positive offers, to plaintiffs in a set $L_r \subseteq N$, and let $\rho_r \in \{\text{yes}, \text{no}\}^{\ell_r}$ denote their responses.

Replace round r by ℓ_r consecutive rounds, each containing one of the original strictly positive offers (and zeros, i.e. no offers, to all other plaintiffs), and specify that these ℓ_r offers do not depend on responses within this block of ℓ_r rounds (so the same response vector ρ_r can be implemented in either representation). After these ℓ_r rounds, the mechanism continues

exactly as the original mechanism would after round r given the response vector ρ_r. Then for every response vector ρ_r, the induced payments (and hence plaintiffs' total payoffs) are the same as under the original mechanism, and the defendant's cost is unchanged. Iterating this argument across rounds, it is without loss to consider mechanisms that make at most one strictly positive offer in each round.

Step 4: We claim that it is without loss to focus on mechanisms whose induced path is h_{yes}. Consider a mechanism satisfying Steps 1-3, and let h_{yes} be an all-accept history in which every plaintiff receives (and accepts) a strictly positive offer. Let h be the history selected by the plaintiffs' coordinated response under the mechanism. If $h = h_{\text{yes}}$, then no plaintiff rejects a strictly positive offer along the induced path and we are done.

Otherwise, h involves at least one rejection. Since the plaintiffs maximize joint payoff, their joint payoff under h is strictly larger than under h_{yes}; if it were equal, tie-breaking would select h_{yes}. Let $\Delta > 0$ denote this difference. Modify the mechanism only by increasing the accepted offer in the last round along h_{yes} by Δ, and after a rejection of this increased offer let the mechanism continue as it would have after a rejection of the original offer. By Step 3, there is at most one strictly positive offer in that round, so this change raises the plaintiffs' joint payoff along h_{yes} by exactly Δ. Histories that diverge from h_{yes} before that last round are unchanged. Histories that follow h_{yes} until that round and then reject the increased offer do not yield a higher payoff than before: rejecting the higher offer creates no additional settlement payment, and if litigation occurs, it can only weakly increase the costs borne by plaintiffs under fee shifting. Hence no history yields the plaintiffs a joint payoff strictly above the original joint payoff under h, while h_{yes} now yields that joint payoff. Therefore h_{yes} becomes payoff-maximizing, and by our tie-breaking assumption, the plaintiffs select h_{yes}.

By construction, along the induced history, no plaintiff rejects a strictly positive offer and every plaintiff receives a strictly positive offer. Therefore, each plaintiff settles by accepting his first strictly positive offer, as claimed. Since there is no litigation along either h_{yes} or the originally induced history h, the defendant's total cost equals the plaintiffs' joint payoff on these histories; since the modification in this step makes the plaintiffs' joint payoff along h_{yes} equal to that along h, the defendant's total cost is unchanged.

B.2 Proof of Lemma 2

Fix any mechanism $T \in \mathcal{T}$. We show that the plaintiffs can secure a joint payoff of at least $\underline{W}$. Since the defendant's total cost always weakly exceeds the plaintiffs' joint payoff (the difference is the total litigation costs), it follows that the defendant cannot achieve a cost strictly below $\underline{W}$, yielding the lower bound in the lemma.

Consider the following coordinated strategy for the plaintiffs: in each round, accept any offer at least $x(1)$ and reject all strictly lower offers. Let α denote the acceptance profile that results under this strategy given T, and let $S(\alpha)$ be the set of plaintiffs who go to court. By construction, each plaintiff $i \notin S(\alpha)$ receives (and accepts) an offer of at least $x(1)$. Moreover, if $i \in S(\alpha)$, then along the history induced by α all offers made to i are strictly below $x(1)$. Since x is increasing, $x(s(\alpha)) \geq x(1)$ whenever $s(\alpha) \geq 1$, and therefore $\max\{t_i^{T,\alpha}(1), \ldots, t_i^{T,\alpha}(R)\} < x(s(\alpha))$ for each $i \in S(\alpha)$. This implies $z(T, \alpha) = 0$. Hence, the plaintiffs' joint payoff under this strategy is

$$\sum_{i \in N \setminus S(\alpha)} t_i^{T,\alpha}(\alpha_i) + s(\alpha)x(s(\alpha)) \geq (n - s(\alpha))x(1) + s(\alpha)x(1) = nx(1) = \underline{W}.$$

Therefore, regardless of the defendant's mechanism, the plaintiffs can secure a joint payoff of at least $\underline{W}$, and thus the defendant's optimal cost must satisfy $W^* \geq \underline{W}$.

To prove the upper bound, consider a mechanism that offers $x(n)$ to each plaintiff (e.g., in round 1). Fix any plaintiffs' response, and let S be the set of plaintiffs who litigate, with $s := |S|$. For each $i \in S$, the defendant pays damages of $x(s)$, and bears no litigation cost for i under the fee-shifting rule since $x(n) \geq x(s)$. For each $i \notin S$, the defendant pays $x(n)$ upon settlement. Thus, the defendant's total cost under this mechanism is

$$(n - s)x(n) + sx(s) \leq (n - s)x(n) + sx(n) = nx(n) = \overline{W},$$

where the inequality uses that x is increasing. It follows that $W^* \leq \overline{W}$.

B.3 Proof of Proposition 1

When there is no fee-shifting rule, the plaintiffs' litigation payoffs do not depend on the defendant's past settlement offers. Hence, regardless of the mechanism proposed by the defendant, since the plaintiffs can coordinate their acceptance decisions, they can secure a total payoff of $\overline{W}$ by rejecting all offers and going to court. This implies $W^* \geq \overline{W}$.

For the reverse inequality, consider a mechanism that offers $x(n)$ to each plaintiff (e.g., in round 1). If all plaintiffs accept, the defendant pays $\overline{W}$. Under any other acceptance profile, suppose s plaintiffs litigate. Each settling plaintiff receives $x(n)$, while each litigant receives $x(s) \leq x(n)$. Hence the plaintiffs' joint payoff under any acceptance profile is weakly below $\overline{W}$. By our tie-breaking assumption, the plaintiffs accept, so the defendant incurs cost $\overline{W}$. Therefore $W^* \leq \overline{W}$.

B.4 Proof of Proposition 2

Consider the following mechanism T with $R = n$ rounds. In round r, the defendant makes an offer only to plaintiff $i = r$. Plaintiff r is offered $x(1)$ if all offers in rounds $\{1, \ldots, r-1\}$ were accepted, and is offered $x(n)$ otherwise.

We claim that in the induced game, given our tie-breaking assumption under indifference, every plaintiff accepts in every subgame, and this outcome is the unique subgame-perfect equilibrium outcome.

Fix a round r and consider any history through round $r-1$. Suppose first that some previous offer was rejected. Then plaintiff r is offered $x(n)$. Rejecting cannot yield a strictly higher payoff for plaintiff r than accepting $x(n)$ (since $x(n)$ is at least as large as plaintiff r's litigation payoff in any such history), and therefore plaintiff r accepts.

Suppose next that all previous offers were accepted. Then plaintiff r is offered $x(1)$. If plaintiff r rejects, the mechanism switches to offering $x(n)$ to each remaining plaintiff, who (by the argument above) accept. Thus, if plaintiff r rejects in this history, he becomes the only plaintiff to proceed to court, and so obtains a litigation payoff of at most $x(1)$. Since rejecting cannot yield a strictly higher payoff for plaintiff r than accepting $x(1)$, the plaintiff accepts.

Backward induction then implies that all plaintiffs accept along every history, so the defendant incurs cost $\underline{W}$ under the proposed mechanism. This implies $W^* \leq \underline{W}$.

For the reverse inequality, fix any mechanism and any equilibrium. No plaintiff can accept an offer strictly below $x(1)$: by rejecting every offer strictly below $x(1)$ and accepting any later offer weakly above $x(1)$, he guarantees himself at least $x(1)$, since if s plaintiffs litigate then his trial award is $x(s) \geq x(1)$, and fee shifting cannot make him bear the defendant's litigation cost after rejecting only offers strictly below $x(1)$. Thus each plaintiff's equilibrium payoff is at least $x(1)$, so the plaintiffs' joint payoff is at least $\underline{W}$. Since the defendant's total cost weakly exceeds the plaintiffs' joint payoff, it follows that $W^* \geq \underline{W}$.

Both the upper-bound construction and the lower-bound argument apply regardless of whether a fee-shifting rule is in place. This establishes Remark 1.

C Proofs for Section 5

C.1 Proof of Lemma 3

Fix any mechanism $T \in \hat{\mathcal{T}}$, and let $\hat{\alpha}$ be the acceptance profile under which each plaintiff accepts his first strictly positive offer. Then

$$W(T) = \sum_{i \in N} t_i^{T,\hat{\alpha}}(\hat{\alpha}_i),$$

and the incentive constraint (IC) implies that for every $\alpha \in \mathcal{A}$,

$$W(T) \geq \sum_{i \in N \setminus S(\alpha)} t_i^{T,\alpha}(\alpha_i) + s(\alpha)x(s(\alpha)) - z(T,\alpha)k.$$

Fix $s \in \{0, \ldots, n-1\}$, and let $S \subset N$ satisfy $|S| = s$. Define an acceptance profile $\alpha^s \in \mathcal{A}$ as follows. For each $i \in S$, set $\alpha_i^s := R+1$. For each $i \in N \setminus S$, let α_i^s be the first round r along the history induced by α^s such that $t_i^{T,\alpha^s}(r) \geq x(s+1)$, if such a round exists, and set $\alpha_i^s := R+1$ otherwise. Let $\ell := s(\alpha^s)$ be the realized number of litigants.

Under α^s, each plaintiff in $N \setminus S$ obtains at least $x(s+1)$: either he settles for an offer weakly above $x(s+1)$ by construction, or he litigates, in which case $\ell \geq s+1$ and his trial award is $x(\ell) \geq x(s+1)$. Each plaintiff in S litigates and receives trial award $x(\ell) \geq x(s)$ since $\ell \geq s$. Moreover, if $\ell \geq s+1$, then any additional litigant $i \in S(\alpha^s) \setminus S$ has rejected only offers strictly below $x(s+1) \leq x(\ell)$, so such a litigant cannot be made responsible for litigation costs under fee shifting. Hence $z(T,\alpha^s) \leq s$. It follows that the plaintiffs' joint payoff under α^s is at least

$$(n-s)x(s+1) + s(x(s) - k) = V_s.$$

Applying (IC) to α^s yields $W(T) \geq V_s$.

Now consider $\alpha^n := (R+1, \ldots, R+1)$, under which all n plaintiffs litigate. Since $z(T,\alpha^n) \leq n$, their joint payoff is at least $n(x(n) - k) = V_n$, so (IC) implies $W(T) \geq V_n$.

Since $W(T) \geq V_s$ for all $s \in \{0, \ldots, n\}$, we have $W(T) \geq \max_{s \in \{0,\ldots,n\}} V_s = \overline{V}$. Since any optimal mechanism can be taken to belong to $\hat{\mathcal{T}}$, it follows that $W^* \geq \overline{V}$.

C.2 Proof of Proposition 3

By Lemma 3, $W^* \geq \overline{V}$. Hence it suffices to show that a phased mechanism $T \in \mathcal{T}(\nu^*)$ induces full settlement at total cost $\overline{V}$. In particular, we will show that any such mechanism

belongs to $\hat{\mathcal{T}}$.

Let $\nu^* := \overline{V}/n$, and fix any mechanism $T \in \mathcal{T}(\nu^*)$. If all plaintiffs accept in phase 1, each receives ν^* and the defendant's total cost is $n\nu^* = \overline{V}$. We claim that under T, the plaintiffs' joint payoff from any other acceptance/rejection profile is weakly below $n\nu^*$. Since plaintiffs coordinate to maximize their joint payoff, and given our tie-breaking rule, this implies that they accept all phase-1 offers and the defendant attains total cost $\overline{V}$.

First note that $\nu^* \geq x(1)$, since $\overline{V} \geq V_0 = nx(1)$.

If not all plaintiffs accept in phase 1, then play ends in one of the following scenarios:

Scenario 1: All plaintiffs reject all offers and go to trial. Then phase n is reached and each litigant has rejected an offer of $x(n)$, so by the fee-shifting rule each litigant bears cost k. Thus the plaintiffs' joint payoff is $n(x(n) - k) = V_n \leq \overline{V} = n\nu^*$.

Scenario 2: The mechanism ends in some phase $q \in \{2, \ldots, n\}$. Then exactly $n - q + 1$ plaintiffs have accepted offers and the remaining $q - 1$ plaintiffs litigate. Each settling plaintiff accepts an offer of at most $\max\{\nu^*, x(q)\}$, so the total settlement payment is at most $(n - q + 1)\max\{\nu^*, x(q)\}$.

If $q = 2$, then the single litigant rejected the phase-1 offer $\nu^* \geq x(1)$, so he bears cost k at trial. If $q \geq 3$, then phase $q - 1$ must have ended with all remaining plaintiffs rejecting $x(q-1)$, so each litigant rejected an offer of $x(q-1)$ and therefore bears cost k at trial. In all cases, the plaintiffs' joint payoff is no larger than

$$(n - q + 1)\max\{\nu^*, x(q)\} + (q - 1)(x(q - 1) - k).$$

If $x(q) \geq \nu^*$, then this expression equals

$$(n - q + 1)x(q) + (q - 1)(x(q - 1) - k) = V_{q-1} \leq \overline{V} = n\nu^*.$$

If instead $x(q) < \nu^*$, then $x(q-1) < x(q) < \nu^*$ and $k > 0$, so $x(q-1) - k < \nu^*$ and therefore

$$(n - q + 1)\max\{\nu^*, x(q)\} + (q - 1)(x(q - 1) - k) < (n - q + 1)\nu^* + (q - 1)\nu^* = n\nu^*.$$

We have shown that accepting all phase-1 offers yields joint payoff $n\nu^*$ and no other response yields strictly more. By our tie-breaking assumption, the plaintiffs accept in phase 1, so the defendant attains total cost $n\nu^* = \overline{V}$ under T. Thus $W^* \leq \overline{V}$. Together with $W^* \geq \overline{V}$, we conclude that $W^* = \overline{V}$.

C.3 Proof of Corollary 2

By Proposition 3, $W^* = \overline{V} = \max_{s\in\{0,\dots,n\}} V_s$. We show that $V_s < \overline{W}$ for all $s \in \{0,\dots,n\}$, which implies $\overline{V} < \overline{W}$ and hence $W^* < \overline{W}$.

If $s = 0$, then $V_0 = nx(1) < nx(n) = \overline{W}$ since x is strictly increasing. If $s \in \{1,\dots,n-1\}$, then $x(s+1) \le x(n)$ and $x(s) \le x(n)$, so

$$V_s = (n-s)x(s+1) + s(x(s) - k) \le (n-s)x(n) + s(x(n) - k) = \overline{W} - sk < \overline{W},$$

since $k > 0$. Finally, if $s = n$, then $V_n = n(x(n) - k) = \overline{W} - nk < \overline{W}$.

C.4 Proof of Corollary 3

Note that $V_0 = nx(1) = \underline{W}$. Also, $V_n \le V_0$ is equivalent to $k \ge x(n) - x(1)$.

For each $q \in \{2,\dots,n\}$, we have

$$V_{q-1} = (n-q+1)x(q) + (q-1)(x(q-1) - k).$$

Hence $V_{q-1} \le V_0$ is equivalent to $k \ge \overline{\kappa}_q$, where recall that

$$\overline{\kappa}_q = n\Big(\frac{x(q) - x(1)}{q-1}\Big) - (x(q) - x(q-1)).$$

Thus, if $k \ge \overline{k} = \max\Big\{x(n) - x(1), \max_{q\in\{2,\dots,n\}} \overline{\kappa}_q\Big\}$, then $V_s \le V_0$ for all $s \in \{0,\dots,n\}$, so $\overline{V} = \max_s V_s = V_0 = \underline{W}$. By Proposition 3, $W^* = \overline{V} = \underline{W}$.

Now suppose x is weakly convex, i.e., $\Delta_j := x(j+1) - x(j)$ is weakly increasing in j. Fix $s \in \{1,\dots,n-1\}$. We claim that

$$V_s \le \frac{n-s}{n} V_0 + \frac{s}{n} V_n.$$

Since $V_s = (n-s)x(s+1) + s(x(s) - k)$, this inequality is equivalent to

$$(n-s)(x(s+1) - x(1)) \le s(x(n) - x(s)).$$

Using $x(s+1) - x(1) = \Delta_1 + \cdots + \Delta_s$ and $x(n) - x(s) = \Delta_s + \cdots + \Delta_{n-1}$, weak convexity implies

$$\frac{x(s+1) - x(1)}{s} = \frac{\Delta_1 + \cdots + \Delta_s}{s} \le \Delta_s \le \frac{\Delta_s + \cdots + \Delta_{n-1}}{n-s} = \frac{x(n) - x(s)}{n-s},$$

which proves the claim. Therefore, $V_s \leq \max\{V_0, V_n\}$ for all $s \in \{1, \ldots, n-1\}$, and hence $\overline{V} = \max\{V_0, V_n\} = \max\{\underline{W}, \overline{W} - nk\}$. It follows that $\overline{V} = \underline{W}$ whenever $V_0 \geq V_n$, i.e. whenever $k \geq x(n) - x(1)$. Thus, under weak convexity one may take $\overline{k} = x(n) - x(1)$.

D Proofs for Section 6

D.1 Proof of Proposition 4

Under bilateral offers ($R = 1$), a mechanism is an offer vector $t = (t_i)_{i \in N}$. We proceed in three steps.

Step 1: We prove the lower bound $\underline{W}_b \leq W_b^*$. Fix any bilateral offer vector t, and let α be the acceptance profile selected by the plaintiffs under t. Let $s := s(\alpha)$ be the number of litigating plaintiffs.

If $s = n$, then the defendant's total cost is at least $nx(n) \geq \underline{W}_b$, since x is increasing. Thus suppose $s \leq n-1$. Let $a_{(1)} \leq \cdots \leq a_{(n-s)}$ denote the offers accepted under α, ordered from lowest to highest. We claim that $a_{(j)} \geq x(s+j)$ for all $j \in \{1, \ldots, n-s\}$. Suppose instead that $a_{(j)} < x(s+j)$ for some $j \in \{1, \ldots, n-s\}$. Consider the deviation in which the j plaintiffs who accepted the lowest offers now reject as well, while all other plaintiffs keep the same decisions as under α. Then exactly $s+j$ plaintiffs litigate, so the trial award is $x(s+j)$. Each of the j new litigants has rejected an offer strictly below $x(s+j)$, so fee shifting cannot make any of them responsible for litigation costs. Moreover, each of the original s litigants now receives a strictly higher trial award than under α, and the higher trial award can only weakly reduce the set of litigants who bear litigation costs under fee shifting. Hence this deviation strictly raises the plaintiffs' joint payoff, contradicting that α is payoff-maximizing.

Therefore $a_{(j)} \geq x(s+j)$ for all $j \in \{1, \ldots, n-s\}$. It follows that the total settlement payment under α is at least

$$\sum_{j=1}^{n-s} a_{(j)} \geq \sum_{j=1}^{n-s} x(s+j) = \sum_{\ell=s+1}^{n} x(\ell).$$

Since the defendant's total cost is weakly above the sum of settlement payments plus trial damages, we have

$$W(t) \geq \sum_{\ell=s+1}^{n} x(\ell) + sx(s).$$

Because x is increasing, $sx(s) \geq \sum_{\ell=1}^{s} x(\ell)$, and therefore

$$W(t) \geq \sum_{\ell=s+1}^{n} x(\ell) + \sum_{\ell=1}^{s} x(\ell) = \sum_{\ell=1}^{n} x(\ell) = \underline{W}_b.$$

Since this holds for every offer vector t, we conclude that $W_b^* \geq \underline{W}_b$.

Step 2: We prove the strict inequality $W_b^* < \overline{W}$. Consider the offer vector $\hat{t}$ with $\hat{t}_1 = x(n)$ and

$$\hat{t}_i = \max\Big\{x(n) - \frac{k}{n-1},\, x(n-1)\Big\} =: \hat{t}^*$$

for all $i \in \{2, \ldots, n\}$. If all plaintiffs accept, the defendant pays

$$\hat{W} := x(n) + (n-1)\hat{t}^*.$$

Since $k > 0$ and x is strictly increasing, we have $\hat{t}^* < x(n)$ and therefore $\hat{W} < nx(n) = \overline{W}$.

We show that the plaintiffs' joint payoff from any other acceptance profile is weakly below $\hat{W}$. There are three cases.

Case 1: Plaintiff 1 accepts and exactly $s \in \{1, \ldots, n-1\}$ of the other plaintiffs reject. Then s plaintiffs litigate and each receives $x(s)$. Since $x(s) \leq x(n-1) \leq \hat{t}^*$, each litigant has rejected an offer at least $x(s)$ and therefore bears cost k. Thus the plaintiffs' joint payoff is

$$x(n) + (n-1-s)\hat{t}^* + s(x(s) - k) = \hat{W} - s(\hat{t}^* - x(s) + k) < \hat{W}.$$

Case 2: Plaintiff 1 rejects and the total number of rejecting plaintiffs is $s \in \{1, \ldots, n-1\}$. Then s plaintiffs litigate and each receives $x(s)$. Plaintiff 1 has rejected $x(n) \geq x(s)$, and every other litigant has rejected $\hat{t}^* \geq x(n-1) \geq x(s)$, so every litigant bears cost k. Hence the plaintiffs' joint payoff is

$$(n-s)\hat{t}^* + s(x(s) - k) = \hat{W} - (x(n) - x(s) + k) - (s-1)(\hat{t}^* - x(s) + k) < \hat{W}.$$

Case 3: All n plaintiffs reject. Then each plaintiff receives $x(n)$ at trial. Plaintiff 1 has rejected an offer of $x(n)$ and therefore bears cost k. Since $\hat{t}^* < x(n)$, no plaintiff $i \geq 2$ has rejected an offer weakly above $x(n)$, so fee shifting does not shift costs to them. Thus the plaintiffs' joint payoff is $nx(n) - k$, and

$$nx(n) - k = x(n) + (n-1)\Big(x(n) - \frac{k}{n-1}\Big) \leq x(n) + (n-1)\hat{t}^* = \hat{W},$$

where the inequality uses $\hat{t}^* \geq x(n) - k/(n-1)$.

Therefore, any acceptance profile with some rejection yields joint payoff weakly below $\hat{W}$, while accepting all offers yields exactly $\hat{W}$. By tie-breaking, the plaintiffs accept, so $\hat{t}$ induces full settlement at total cost $\hat{W} < \overline{W}$. Hence $W_b^* < \overline{W}$.

Step 3: We prove that high litigation costs imply $W_b^* = \underline{W}_b$. Consider any bilateral offer vector that assigns the values $\{x(1), \ldots, x(n)\}$ one-to-one across plaintiffs (e.g., $t_i = x(i)$). If all plaintiffs accept, the defendant pays $\underline{W}_b$.

If some plaintiffs reject, then the plaintiffs maximize their joint payoff by having the plaintiffs who received the lowest offers reject. If exactly $s \in \{1, \ldots, n\}$ plaintiffs reject, then the trial award is $x(s)$ and at least one litigant has rejected an offer of $x(s)$, so fee shifting makes at least one litigant bear cost k. Moreover, no plaintiff can ever receive more than $x(n)$. Hence, under any acceptance profile with at least one rejection, the plaintiffs' joint payoff is at most $nx(n) - k$.

Thus accepting all offers maximizes the plaintiffs' joint payoff whenever $\underline{W}_b \geq nx(n) - k$, which is equivalent to

$$k \geq nx(n) - \sum_{i=1}^{n} x(i) = \sum_{i=1}^{n} (x(n) - x(i)) = \overline{k}_b.$$

For such k, tie-breaking implies that all plaintiffs accept, so the defendant attains cost $\underline{W}_b$. Together with Step 1, this argument yields $W_b^* = \underline{W}_b$ for all $k \geq \overline{k}_b$. This also establishes Remark 3.

D.2 Proof of Corollary 4

In the absence of fee shifting, plaintiffs can guarantee $\overline{W} = nx(n)$ by rejecting all offers and litigating, so the defendant's optimal cost is at least $\overline{W}$. Conversely, offering $x(n)$ to each plaintiff in a single round yields full settlement at total cost $\overline{W}$. Hence the defendant's optimal cost without fee shifting is $\overline{W}$.

Under fee shifting and for $k \geq \overline{k}_b$, Proposition 4 gives $W_b^* = \underline{W}_b$. Therefore, under high enough litigation costs, fee shifting allows the defendant to appropriate the full surplus $\overline{W} - \underline{W}_b$, and plaintiffs obtain no benefit relative to their preferred equilibrium under no side contracting (see Remark 2).

D.3 Proof of Proposition 5

We proceed in three steps.

Step 1: We prove the lower bound $\underline{W}_b \leq W_m^*$. Fix any defendant mechanism in the mutual-offers environment. Consider the following coordinated strategy for the plaintiffs. At the start of any round, let u be the number of plaintiffs who have not yet settled. Each unsettled plaintiff accepts any defendant offer weakly above $x(u)$ and rejects any strictly lower offer. If he does not accept an offer in that round and is able to make a counteroffer, he makes a counteroffer of $x(u)$.

Let s be the number of plaintiffs who ultimately litigate under this strategy. We claim that the plaintiffs' joint payoff is at least $\underline{W}_b$.

First, consider the settlement payments. Order the $n - s$ settlements chronologically, breaking ties arbitrarily if multiple settlements occur in the same round. Before the jth settlement, at least $n - j + 1$ plaintiffs remain unsettled, so by the strategy the settling plaintiff receives at least $x(n - j + 1)$. Hence total settlement payments are at least

$$\sum_{j=1}^{n-s} x(n - j + 1) = \sum_{\ell=s+1}^{n} x(\ell).$$

Second, consider litigation payoffs. If $s = 0$, then all plaintiffs settle and the total payoff is at least $\underline{W}_b$ by the previous bound. Suppose $s \geq 1$. We claim that no litigating plaintiff can be made responsible for the defendant's litigation cost k under the fee-shifting rule.

Fix a litigant i. If i never rejected a defendant offer weakly above the trial award $x(s)$, then fee shifting cannot apply to i. Otherwise, let r_s be the first round in which exactly s plaintiffs remain unsettled. Any rejection by i of a defendant offer weakly above $x(s)$ must occur before round r_s, since with exactly s unsettled plaintiffs the strategy accepts any defendant offer weakly above $x(s)$. We claim that in round r_s, plaintiff i makes a counteroffer of $x(s)$. If the defendant makes an offer to i in round r_s, then since i does not accept, that offer must be strictly below $x(s)$; by the coordinated strategy, i then counteroffers $x(s)$ in the same round. If instead the defendant makes no offer to i in round r_s, then i can still make an offer: his last rejection must have occurred when more than s plaintiffs were unsettled, and because round r_s is the first round with exactly s unsettled plaintiffs, some other plaintiff must have settled after that rejection, which restores i's ability to make offers under the protocol in Section 6.2. Hence, after any rejection of a defendant offer weakly above $x(s)$, plaintiff i later makes a counteroffer weakly below $x(s)$. Therefore the fee-shifting rule does not make i responsible for the defendant's litigation cost.

It follows that the litigants' joint trial payoff is $sx(s)$. Since x is increasing, we have

$sx(s) \geq \sum_{\ell=1}^{s} x(\ell)$. Therefore the plaintiffs' total payoff is at least

$$\sum_{\ell=s+1}^{n} x(\ell) + sx(s) \geq \sum_{\ell=s+1}^{n} x(\ell) + \sum_{\ell=1}^{s} x(\ell) = \underline{W}_b.$$

Thus the plaintiffs can guarantee themselves a joint payoff of at least $\underline{W}_b$ regardless of the defendant's mechanism. Any mechanism that induces full settlement must therefore deliver a joint settlement payoff at least $\underline{W}_b$, implying $W_m^* \geq \underline{W}_b$.

Step 2: We prove the strict inequality $W_m^* < \overline{W}$. Consider the following one-round mutual-offers mechanism. The defendant makes offers $\hat{t}_1 = x(n)$ and

$$\hat{t}_i = \max\left\{x(n) - \frac{k}{n-1},\, x(n-1)\right\} =: \hat{t}^*$$

for all $i \in \{2, \ldots, n\}$. Additionally, she commits to accept any plaintiff counteroffer weakly below the offer she made to that plaintiff, and to reject any strictly higher counteroffer.

Under this defendant's acceptance rule, any counteroffer that is accepted yields the plaintiff no more than accepting the defendant's offer, and any rejected counteroffer is strictly above the rejected offer. In particular, rejected counteroffers cannot prevent fee shifting at trial after rejecting an offer weakly above the trial award. Hence, for the purpose of maximizing the plaintiffs' joint payoff, we may restrict attention to accept/reject behavior without counteroffers, and the payoff comparison is identical to the bilateral-offers case.

Let $\hat{W} := x(n) + (n-1)\hat{t}^*$. Since $\hat{t}^* < x(n)$, we have $\hat{W} < nx(n) = \overline{W}$. A case analysis as in the proof of Proposition 4 shows that the plaintiffs' joint payoff from any acceptance profile with some litigation is weakly below $\hat{W}$, while accepting all offers yields exactly $\hat{W}$. By tie-breaking, the plaintiffs accept all offers, and the defendant induces full settlement at cost $\hat{W} < \overline{W}$. Therefore $W_m^* < \overline{W}$.

Step 3: We prove that high litigation costs imply $W_m^* = \underline{W}_b$. Fix $k \geq \overline{k}_m$, recalling that

$$\overline{k}_m = \max_{s \in \{1,\ldots,n\}} \frac{1}{s} \sum_{i=1}^{s} (x(s) - x(i)).$$

Consider the following phased mutual-offers mechanism. While u plaintiffs remain unsettled, the defendant approaches them sequentially and offers $x(u)$ until one plaintiff accepts, in which case that plaintiff settles and the mechanism moves to the phase with $u-1$ unsettled plaintiffs. In every round in which the defendant offers $x(u)$ to a plaintiff, she also commits to accept any counteroffer from that plaintiff weakly below $x(u)$ and to reject any strictly

higher counteroffer, and to reject any offers from plaintiffs to whom she did not make an offer in that round. If the defendant has offered $x(u)$ to each of the u unsettled plaintiffs (in some order) and all have rejected, then the defendant makes no further offers and rejects any future offers.

If all plaintiffs accept along the intended path, exactly one plaintiff settles at each value $x(n), x(n-1), \ldots, x(1)$, so the defendant's total cost is $\underline{W}_b$.

Now consider any other response profile, and let $s \geq 1$ be the number of litigating plaintiffs in the terminal history. Then $n - s$ plaintiffs settled in earlier phases, each at a payment of at most the phase value, so total settlement payments are at most $\sum_{\ell=s+1}^{n} x(\ell)$.

Moreover, in the final phase the s litigants all rejected an offer of $x(s)$. Since the defendant accepts any counteroffer weakly below $x(s)$ from a plaintiff to whom she offers $x(s)$ in that round, any litigant who proceeds to trial must have rejected that offer without making such a counteroffer. Also, since no plaintiff settles after these rejections and the defendant makes no further offers, by our assumptions in Section 6.2, no litigant could have made further offers after rejecting $x(s)$. Hence, by the mutual-offers fee-shifting rule, each litigant bears cost k and the litigants' joint payoff is at most $s(x(s) - k)$. Therefore the plaintiffs' joint payoff is no larger than

$$\sum_{\ell=s+1}^{n} x(\ell) + s(x(s) - k).$$

This is weakly below $\underline{W}_b = \sum_{i=1}^{n} x(i)$ if and only if

$$\sum_{i=1}^{s} x(i) \geq s(x(s) - k) \iff k \geq \frac{1}{s} \sum_{i=1}^{s} (x(s) - x(i)),$$

which holds for all s when $k \geq \overline{k}_m$.

Thus, for $k \geq \overline{k}_m$, no deviation yields the plaintiffs a joint payoff strictly above $\underline{W}_b$, while accepting along the intended path yields exactly $\underline{W}_b$. By tie-breaking, the plaintiffs accept along the intended path, and the defendant attains cost $\underline{W}_b$. Combined with Step 1, this implies $W_m^* = \underline{W}_b$ for all $k \geq \overline{k}_m$.

D.4 Proof of Remark 4

We proved the first claim of the remark in Step 3 of the proof of Proposition 5. To prove the remaining claims, first note that for any $s \in \{1, \ldots, n\}$, since x is increasing,

$$\frac{1}{s} \sum_{i=1}^{s} (x(s) - x(i)) \leq \frac{1}{s} \sum_{i=1}^{s} (x(n) - x(i)) \leq \frac{1}{s} \sum_{i=1}^{n} (x(n) - x(i)) = \frac{\overline{k}_b}{s}.$$

Since the $s = 1$ term equals 0 and $1/s \leq 1/2$ for all $s \geq 2$, it follows that

$$\overline{k}_m \leq \frac{\overline{k}_b}{2} < \overline{k}_b.$$

Now suppose x is weakly convex, i.e., $\Delta_j := x(j+1) - x(j)$ is weakly increasing in j. For $s \geq 2$,

$$\frac{1}{s}\sum_{i=1}^{s}(x(s) - x(i)) = \frac{1}{s}\sum_{j=1}^{s-1} j\Delta_j =: \psi(s).$$

For $s \in \{2, \ldots, n-1\}$,

$$\psi(s+1) - \psi(s) = \frac{s^2\Delta_s - \sum_{j=1}^{s-1} j\Delta_j}{s(s+1)} \geq 0,$$

since $\Delta_j \leq \Delta_s$ for all $j \leq s-1$ implies

$$\sum_{j=1}^{s-1} j\Delta_j \leq \Delta_s \sum_{j=1}^{s-1} j = \Delta_s \frac{s(s-1)}{2} \leq s^2\Delta_s.$$

Thus ψ is weakly increasing and the maximum defining $\overline{k}_m$ is attained at $s = n$, so

$$\overline{k}_m = \psi(n) = \frac{1}{n}\sum_{i=1}^{n}(x(n) - x(i)) = \frac{\overline{k}_b}{n}.$$

E Proofs for Section 7

E.1 Proof of Proposition 6

The proof of Lemma 1 extends directly to the heterogeneous setting. We therefore restrict attention to mechanisms $T \in \hat{\mathcal{T}}$ that induce each plaintiff to accept his first strictly positive offer. Fix any such mechanism T. Let

$$W_h(T) = \sum_{i \in N} t_i^{T,\hat{\alpha}}(\hat{\alpha}_i),$$

where $\hat{\alpha}$ is the acceptance profile under which each plaintiff accepts his first strictly positive offer. For any $\alpha \in \mathcal{A}$, define

$$Z_h(T, \alpha) := \{i \in S(\alpha) : \max\{t_i^{T,\alpha}(1), \ldots, t_i^{T,\alpha}(R)\} \geq x_i(s(\alpha))\}.$$

The heterogeneous analogue of (IC) is

$$W_h(T) \geq \sum_{i \in N \setminus S(\alpha)} t_i^{T,\alpha}(\alpha_i) + \sum_{i \in S(\alpha)} x_i(s(\alpha)) - \sum_{i \in Z_h(T,\alpha)} k_i. \tag{1}$$

We now proceed in two steps.

Step 1: We prove $W_h^* \geq \overline{V}$. Fix $S \subsetneq N$. Define an acceptance profile $\alpha^S \in \mathcal{A}$ as follows. For each $i \in S$, set $\alpha_i^S := R+1$. For each $i \in N \setminus S$, let α_i^S be the first round r along the history induced by α^S such that $t_i^{T,\alpha^S}(r) \geq x_i(|S|+1)$, if such a round exists, and set $\alpha_i^S := R+1$ otherwise. Let $\ell := s(\alpha^S)$ be the realized number of litigants.

Under α^S, each plaintiff $i \in N \setminus S$ obtains at least $x_i(|S|+1)$: either he settles for an offer weakly above $x_i(|S|+1)$ by construction, or he litigates, in which case $\ell \geq |S|+1$ and his trial award is $x_i(\ell) \geq x_i(|S|+1)$. Each plaintiff $i \in S$ litigates and receives trial award $x_i(\ell) \geq x_i(|S|)$ since $\ell \geq |S|$. Moreover, if $\ell \geq |S|+1$, then any additional litigant $i \in S(\alpha^S) \setminus S$ has rejected only offers strictly below $x_i(|S|+1) \leq x_i(\ell)$, so such a litigant cannot be made responsible for litigation costs under fee shifting. Hence $Z_h(T, \alpha^S) \subseteq S$. It follows that the plaintiffs' joint payoff under α^S is at least

$$\sum_{i \in N \setminus S} x_i(|S|+1) + \sum_{i \in S} (x_i(|S|) - k_i) = V(S).$$

Applying (1) to α^S yields $W_h(T) \geq V(S)$.

Now consider $\alpha^N := (R+1, \ldots, R+1)$, under which all plaintiffs litigate. Since $Z_h(T, \alpha^N) \subseteq N$, their joint payoff is at least

$$\sum_{i \in N} (x_i(n) - k_i) = V(N),$$

so (1) implies $W_h(T) \geq V(N)$.

Since $W_h(T) \geq V(S)$ for all $S \subseteq N$, we have $W_h(T) \geq \max_{S \subseteq N} V(S) = \overline{V}$. Since any optimal mechanism can be taken to belong to $\hat{\mathcal{T}}$, it follows that $W_h^* \geq \overline{V}$.

Step 2: We prove $W_h^* \leq \overline{V}$. Consider the following modified phased mechanism. Starting with $U = N$, the defendant selects some plaintiff $i \in U$ and offers him $\overline{V}_U - \overline{V}_{U \setminus \{i\}}$. If plaintiff i accepts, he settles, the set of unsettled plaintiffs becomes $U \setminus \{i\}$, and the defendant repeats the same step. If plaintiff i rejects and $|U| = 1$, he proceeds to trial. If plaintiff i rejects and $|U| \geq 2$, the mechanism proceeds to phases $q \in \{2, \ldots, |U|\}$ for the current set U. In each phase q, the defendant approaches the unsettled plaintiffs in some order and offers plaintiff j the amount $x_j(q)$. If $q = 2$ and plaintiff i rejected $\overline{V}_U - \overline{V}_{U \setminus \{i\}}$, he is approached last. If

at the start of phase q only $q-1$ plaintiffs in U remain unsettled, or as soon as this number falls to $q-1$, the mechanism ends and these plaintiffs proceed to trial. If this does not occur by the end of the phase, the mechanism moves to phase $q+1$, and if $q = |U|$, all remaining plaintiffs proceed to trial.

First note that for each nonempty set $U \subseteq N$ and each plaintiff $i \in U$, the personalized offer $\overline{V}_U - \overline{V}_{U\setminus\{i\}}$ is weakly above his stand-alone claim value, $x_i(1)$. If $U = \{i\}$, then $\overline{V}_U = x_i(1)$ and $\overline{V}_\emptyset = 0$. Suppose instead that $U \setminus \{i\}$ is nonempty, and let $S \subseteq U \setminus \{i\}$ maximize $V_{U\setminus\{i\}}(S)$. If $S \subsetneq U \setminus \{i\}$, then

$$V_U(S) = V_{U\setminus\{i\}}(S) + x_i(|S|+1).$$

If $S = U \setminus \{i\}$, then

$$V_U(S) = V_{U\setminus\{i\}}(S) + x_i(|U|).$$

In both cases, $\overline{V}_U \geq \overline{V}_{U\setminus\{i\}} + x_i(1)$.

If all personalized offers are accepted, total settlement payments equal

$$\sum(\overline{V}_U - \overline{V}_{U\setminus\{i\}}) = \overline{V}_N - \overline{V}_\emptyset = \overline{V}.$$

We claim that under this mechanism, the plaintiffs' joint payoff from any other acceptance profile is weakly below $\overline{V}$. Consider any profile in which some personalized offer is rejected, and let U be the set of unsettled plaintiffs at the first such rejection. The personalized offers accepted before this rejection sum to $\overline{V}_N - \overline{V}_U$. It thus suffices to show that the continuation following this rejection gives the plaintiffs in U a joint payoff no greater than $\overline{V}_U$.

Let i be the plaintiff who rejects the personalized offer at U. If $|U| = 1$, plaintiff i proceeds to trial after rejecting an offer weakly above $x_i(1)$, so his payoff is

$$x_i(1) - k_i = V_U(U) \leq \overline{V}_U.$$

Suppose then that $|U| \geq 2$. Let $L \subseteq U$ be the set of plaintiffs in U who litigate at the end of the continuation, and let $\ell := |L|$. By construction, $\ell \geq 1$. If $1 \leq \ell < |U|$, the mechanism ends in phase $\ell+1$. Each settling plaintiff $j \in U \setminus L$ accepts an offer no greater than $x_j(\ell+1)$. Each litigating plaintiff $j \in L$ bears cost k_j: if $\ell = 1$, then the sole litigant is either plaintiff i, who rejected the personalized offer weakly above $x_i(1)$, or a plaintiff who rejected an offer of $x_j(2)$ in phase 2; if $\ell \geq 2$, then every plaintiff in L rejected an offer of

$x_j(\ell)$ in phase ℓ. Hence the plaintiffs in U obtain joint payoff no larger than

$$\sum_{j\in U\setminus L} x_j(\ell+1) + \sum_{j\in L}(x_j(\ell) - k_j) = V_U(L) \le \overline{V}_U.$$

If instead $\ell = |U|$, all plaintiffs in U litigate after rejecting offers of $x_j(|U|)$, so their joint payoff is

$$\sum_{j\in U}(x_j(|U|) - k_j) = V_U(U) \le \overline{V}_U.$$

Thus, after the first rejected personalized offer, the plaintiffs in U can obtain at most $\overline{V}_U$. Combining this with the personalized offers accepted before that rejection, the plaintiffs' joint payoff is at most

$$(\overline{V}_N - \overline{V}_U) + \overline{V}_U = \overline{V}_N = \overline{V}.$$

We have shown that accepting all personalized offers yields joint payoff $\overline{V}$, and no other response yields strictly more. By our tie-breaking assumption, the plaintiffs accept all personalized offers. Since these offers are strictly positive, the mechanism belongs to $\hat{\mathcal{T}}$, and the defendant attains total cost $\overline{V}$. Thus $W_h^* \le \overline{V}$. Together with $W_h^* \ge \overline{V}$, we conclude that $W_h^* = \overline{V}$.

E.2 Proof of Proposition 7

We prove three claims.

Claim 1. Under the baseline protocol, $W_c^* \le \overline{W}_c$, with strict inequality for c small enough.

Proof of Claim 1. Consider the one-round mechanism that offers $x(n) - c$ to each plaintiff. If all plaintiffs accept, the defendant pays $n(x(n) - c) = \overline{W}_c$. Under any other acceptance profile, suppose s plaintiffs litigate. Each settling plaintiff receives $x(n) - c$, while each litigant's net payoff is at most $x(s) - c \le x(n) - c$ (and is strictly lower if the litigant bears the defendant's litigation cost k under fee shifting). Hence the plaintiffs' joint payoff under any acceptance profile is weakly below $\overline{W}_c$, while accepting all offers yields exactly $\overline{W}_c$. By tie-breaking, the plaintiffs accept, so $W_c^* \le \overline{W}_c$.

To prove strict inequality for sufficiently small c, let T^* be an optimal mechanism in the baseline protocol when $c = 0$, with total cost W^*. Since introducing plaintiffs' litigation costs can only weakly reduce the plaintiffs' payoff from any terminal history, the same mechanism T^* also induces full settlement when $c > 0$, at the same total cost W^*. By Corollary 2, $W^* < \overline{W}$. Therefore, if $0 < c < (\overline{W} - W^*)/n$, then $W_c^* \le W^* < \overline{W} - nc = \overline{W}_c$. □

Claim 2. Under bilateral offers, $W_{b,c}^* \le \overline{W}_c$, with strict inequality for c small enough.

Proof of Claim 2. The same one-round offer vector used in Claim 1 is feasible under bilateral offers. Thus, by the same argument, $W^*_{b,c} \leq \overline{W}_c$.

To prove strict inequality for sufficiently small c, let t^* be an optimal bilateral-offers vector when $c = 0$, with total cost W^*_b. Since plaintiffs' litigation costs can only weakly reduce the plaintiffs' payoff from any acceptance profile with litigation, the same vector t^* also induces full settlement when $c > 0$, at the same total cost W^*_b. By Proposition 4, $W^*_b < \overline{W}$. Therefore, if $0 < c < (\overline{W} - W^*_b)/n$, then $W^*_{b,c} \leq W^*_b < \overline{W} - nc = \overline{W}_c$. □

Claim 3. Consider mutual offers, and suppose $k \geq \overline{k}_m$. Then $W^*_{m,c} = \underline{W}_b$ for all c and therefore $W^*_{m,c} < \overline{W}_c$ for all $c < (\overline{W} - \underline{W}_b)/n = \overline{k}_b/n$. If, in addition, $k < x(1) - \overline{k}_b/n$, then $W^*_{m,c} > \overline{W}_c$ for all $c \in (\overline{k}_b/n, x(1) - k)$.

Proof of Claim 3. We first prove that $W^*_{m,c} \leq \underline{W}_b$. By Proposition 5, given $k \geq \overline{k}_m$, there exists a mutual-offers mechanism that induces full settlement at total cost $\underline{W}_b$ when $c = 0$. Under the same mechanism, introducing plaintiffs' litigation costs cannot increase the plaintiffs' joint payoff in any terminal history: each litigating plaintiff either bears cost c or has that cost shifted to the defendant under the two-sided fee-shifting rule, in which case his net payoff is the same as when $c = 0$. Thus, the same mechanism still induces full settlement when $c > 0$, and therefore $W^*_{m,c} \leq \underline{W}_b$.

Next we prove that $W^*_{m,c} \geq \underline{W}_b$. Fix any defendant mechanism in the mutual-offers environment. Consider the coordinated strategy from Step 1 of the proof of Proposition 5: at the start of any round, if u plaintiffs have not yet settled, then each unsettled plaintiff accepts any defendant offer weakly above $x(u)$ and rejects any strictly lower offer; if he does not accept an offer in that round and is able to make a counteroffer, he offers $x(u)$. Let s be the number of plaintiffs who ultimately litigate under this strategy. As shown in Step 1 of the proof of Proposition 5, total settlement payments are at least

$$\sum_{j=s+1}^{n} x(j),$$

and no litigating plaintiff can be made responsible for the defendant's litigation cost k under the mutual-offers fee-shifting rule.

We claim that no litigating plaintiff bears his own litigation cost c either. Fix a litigant i, and let r_s be the first round in which exactly s plaintiffs remain unsettled. We claim that in round r_s, plaintiff i makes a counteroffer of $x(s)$. If the defendant makes an offer to i in round r_s, then since i does not accept, that offer must be strictly below $x(s)$; by the coordinated strategy, i then counteroffers $x(s)$ in the same round. If instead the defendant

makes no offer to i in round r_s, then i can still make an offer: either he has never rejected a defendant offer, or his last rejection occurred when more than s plaintiffs were unsettled. In the latter case, because round r_s is the first round with exactly s unsettled plaintiffs, some other plaintiff must have settled after that rejection, which restores i's ability to make offers. In either case, i counteroffers $x(s)$ in round r_s.

Thus every litigant i makes a counteroffer of $x(s)$ in the first round with exactly s plaintiffs remaining unsettled. Since i ultimately litigates, that counteroffer must be rejected by the defendant, and any later offer by the defendant to i must be strictly below $x(s)$, for otherwise i would accept. Hence, by the two-sided fee-shifting rule, the defendant bears that litigant's cost c. Therefore no litigating plaintiff bears litigation cost c.

It follows that the litigants' joint trial payoff is $sx(s)$. Since x is increasing, $sx(s) \geq \sum_{j=1}^{s} x(j)$. Therefore the plaintiffs' total payoff under this coordinated strategy is at least

$$\sum_{j=s+1}^{n} x(j) + sx(s) \geq \sum_{j=s+1}^{n} x(j) + \sum_{j=1}^{s} x(j) = \underline{W}_b.$$

Thus the plaintiffs can guarantee themselves a joint payoff of at least $\underline{W}_b$ regardless of the defendant's mechanism. Any mechanism that induces full settlement must therefore deliver a joint settlement payoff at least $\underline{W}_b$, implying $W^*_{m,c} \geq \underline{W}_b$.

Combining the two inequalities yields $W^*_{m,c} = \underline{W}_b$. Finally, let $\bar{c} := (\overline{W} - \underline{W}_b)/n = \bar{k}_b/n$, so that $\underline{W}_b = \overline{W} - n\bar{c} = \overline{W}_{\bar{c}}$. We obtain $W^*_{m,c} < \overline{W}_c$ for all $0 < c < \bar{c}$. Moreover, under $k < x(1) - \bar{k}_b/n$, we have $\bar{c} \in (0, x(1) - k)$, and thus $W^*_{m,c} > \overline{W}_c$ for all $\bar{c} < c < x(1) - k$. □

E.3 Proof of Proposition 8

We proceed in three steps.

Step 1: We prove the strict inequality $\underline{W}_\mu < W^*_\mu$. Consider the following coordinated strategy for the plaintiffs: accept any defendant offer weakly above $\overline{x}$ and reject all strictly lower offers. Let s be the (endogenous) number of litigating plaintiffs under this strategy.

If $s = 0$, then every plaintiff settles for at least $\overline{x}$, so the plaintiffs' joint payoff is at least $n\overline{x}$. Since $p(1) < 1$ (because p is strictly increasing and $n \geq 2$) and $k > 0$, we have $\mu(1) < \overline{x}$, and therefore $n\overline{x} > n\mu(1) = \underline{W}_\mu$.

If $s \geq 1$, then under this strategy no litigating plaintiff has rejected an offer weakly above $\overline{x}$, so each litigant's expected payoff is $\mu(s) \geq \mu(1)$. If $s = n$, the joint litigation payoff is $n\mu(n) > n\mu(1) = \underline{W}_\mu$ since μ is strictly increasing. If instead $s \leq n - 1$, then at least one plaintiff settles for at least $\overline{x}$ while the remaining plaintiffs each obtain at least $\mu(1)$ in

expectation, so the plaintiffs' joint payoff is at least

$$\overline{x} + (n-1)\mu(1) > \underline{W}_\mu,$$

where the strict inequality uses $\overline{x} > \mu(1)$.

Thus, regardless of the defendant's mechanism, the plaintiffs can guarantee themselves a joint payoff strictly above $\underline{W}_\mu$. To induce settlement, the defendant must therefore give the plaintiffs a joint settlement payoff strictly above $\underline{W}_\mu$, implying $W^*_\mu > \underline{W}_\mu$.

Step 2: We prove the upper bound $W^*_\mu \leq \overline{W}_\mu$. Consider a one-round mechanism that offers $\mu(n)$ to each plaintiff. If all accept, the defendant pays $\overline{W}_\mu = n\mu(n)$. Under any other acceptance profile, suppose s plaintiffs litigate. Each litigant's expected payoff is at most $\mu(s) \leq \mu(n)$. Hence the plaintiffs' joint payoff under any acceptance profile is weakly below $n\mu(n)$, while accepting all offers yields exactly $n\mu(n)$. By tie-breaking, all accept, so $W^*_\mu \leq n\mu(n) = \overline{W}_\mu$.

Step 3: We prove that $p(n)$ sufficiently close to 1 implies $W^*_\mu < \overline{W}_\mu$. Fix $\overline{x}$, k, and $(p(1), \ldots, p(n-1))$. For each $p(n) \in (p(n-1), 1]$, define

$$\hat{t}^* := \max\left\{\mu(n-1),\ \frac{n\mu(n) - kp(n) - \overline{x}}{n-1}\right\}.$$

Note that $\hat{t}^* < \overline{x}$, since $\mu(n-1) < \overline{x}$ and $n\mu(n) - kp(n) - \overline{x} < (n-1)\overline{x}$.

Consider the one-round offer vector $\hat{t}$ with $\hat{t}_1 = \overline{x}$ and $\hat{t}_i = \hat{t}^*$ for each $i \in \{2, \ldots, n\}$. If all plaintiffs accept, the defendant pays $\hat{W} := \overline{x} + (n-1)\hat{t}^*$.

We show that no acceptance profile yields the plaintiffs a joint payoff strictly above $\hat{W}$. Let j be the number of rejecting plaintiffs among $\{2, \ldots, n\}$. There are three cases.

Case 1: Plaintiff 1 accepts. If $j = 0$, then all plaintiffs settle and the plaintiffs' joint payoff is exactly $\hat{W}$. If instead $j \in \{1, \ldots, n-1\}$, then exactly j plaintiffs litigate, each with expected payoff $\mu(j) \leq \mu(n-1) \leq \hat{t}^*$, and the remaining $n-1-j$ plaintiffs settle for $\hat{t}^*$. Thus the plaintiffs' joint payoff is at most

$$\overline{x} + (n-1-j)\hat{t}^* + j\mu(j) \leq \overline{x} + (n-1)\hat{t}^* = \hat{W}.$$

Case 2: Plaintiff 1 rejects and $j \leq n-2$. Then the total number of litigants is $s = j+1 \leq n-1$. Plaintiff 1 has rejected $\overline{x}$, so he bears cost k regardless of the court outcome and his expected payoff is $\overline{x}p(s) - k < \overline{x}$. Each other litigant (among $\{2, \ldots, n\}$) has rejected an

offer strictly below $\overline{x}$, so his expected payoff is $\mu(s) \leq \mu(n-1) \leq \hat{t}^*$. The remaining $n-1-j$ plaintiffs settle for $\hat{t}^*$. Hence the plaintiffs' joint payoff is strictly below $\overline{x} + (n-1)\hat{t}^* = \hat{W}$.

Case 3: All plaintiffs reject. Then $s = n$ and the plaintiffs' joint payoff is

$$(\overline{x}p(n) - k) + (n-1)\mu(n) = n\mu(n) - kp(n).$$

Since $(n-1)\hat{t}^* \geq n\mu(n) - kp(n) - \overline{x}$ by definition, this is weakly below $\overline{x} + (n-1)\hat{t}^* = \hat{W}$.

Thus, no acceptance profile yields joint payoff strictly above $\hat{W}$, while accepting all offers yields exactly $\hat{W}$. By tie-breaking, the plaintiffs accept all offers, so $\hat{t}$ induces full settlement at total cost $\hat{W}$.

Finally, note that

$$\frac{n\mu(n) - kp(n) - \overline{x}}{n-1} < \frac{n\mu(n) - \overline{x}}{n-1}$$

for all $p(n) > 0$, and $\mu(n) \to \overline{x}$ as $p(n) \to 1$. Since $\mu(n-1) < \overline{x}$, it follows that there exists $\varepsilon > 0$ such that if $p(n) \geq 1 - \varepsilon$, then

$$\mu(n-1) < \frac{n\mu(n) - \overline{x}}{n-1}.$$

For such $p(n)$, we have $(n-1)\hat{t}^* < n\mu(n) - \overline{x}$ and hence

$$\hat{W} = \overline{x} + (n-1)\hat{t}^* < n\mu(n) = \overline{W}_\mu.$$

Therefore, if $p(n) \geq 1 - \varepsilon$, the defendant can induce full settlement at a cost strictly below $\overline{W}_\mu$, implying $W^*_\mu < \overline{W}_\mu$.